nature neuroscience

The neural correlates of subjective value during intertemporal choice

Joseph W Kable & Paul W Glimcher

Neuroimaging studies of decision-making have generally related neural activity to objective measures (such as reward magnitude, probability or delay), despite choice preferences being subjective. However, economic theories posit that decision-makers behave as though different options have different subjective values. Here we use functional magnetic resonance imaging to show that neural activity in several brain regions—particularly the ventral striatum, medial prefrontal cortex and posterior cingulate cortex—tracks the revealed subjective value of delayed monetary rewards. This similarity provides unambiguous evidence that the subjective value of potential rewards is explicitly represented in the human brain.

Since its inception, perceptual psychophysics has rested on the notion that the objective intensity of a stimulus is distinct from the subjective intensity that guides behavior[1,2]. If, for example, the number of photons emitted by a light source is doubled, the perceived intensity only increases by a factor of about 0.3 (ref. 3). Since the eighteenth century, economics has rested on a similar observation[4]. If, for example, the gain earned for an action is doubled, people behave as though the desirability of this action has been increased by a much smaller fraction. Economic theory calls the functions that relate objective values and subjective desirabilities 'preference functions', and, as in psychophysics, the form of these perceptual functions is inferred from a person's observed behavior.

Although the notion that choices reflect subjective desirability is central to nearly all economic theories of decision-making, studies of the neurobiological basis of decision-making have usually related neural activity to more objective measures[5–16] (but see refs. 17–25 for exceptions). Thus, although perceptual neurobiologists have accumulated evidence that the nervous system encodes subjective sensory percepts, rather than objective reality[3,26–29], there is less evidence that subjective valuations, or preference functions, are encoded in the brain.

Decisions about monetary gains that will be available after different delays clearly illustrate the role of preference functions. An individual possessing a $20 check that can be cashed in one week might trade that check for as little as $18 in cash, available immediately. If the check was payable in one month, she might accept as little as $15 now in exchange. The subjective value of a $20 gain declines as the imposed delay to its receipt increases, a phenomenon known as temporal discounting[30–34]. This decline in subjective value can differ across individuals—a second person might accept as little as $10 immediately in lieu of $20 in a month. Therefore, the objective amount and delay associated with each gain alone cannot predict choice; an idiosyncratic function that relates delay to subjective value is required. This person-specific function, which describes the decline in subjective value with increasing delays, is called a discount function. Whereas temporal discounting has been widely studied behaviorally, there have only been a few investigations of the neural mechanisms of intertemporal choice in humans[22,35].

One of the most fruitful approaches for relating neural activity to behavior (including behavior that depends on subjective experiences) has been psychometric-neurometric comparisons[26,27]. This method tests whether a particular, externally quantifiable variable influences both psychophysical and neurobiological measurements in a similar manner. One influential example of this approach found a close match between monkeys' perceptual decisions about motion direction and predictions from simultaneously recorded neural activity in motion area MT[29]. This approach, though, need not be restricted to directly observable phenomena such as subjects' choices; it can also be used to relate indirectly revealed perceptual experiences or mental states to neural activity[3,26]. Indeed, the earliest psychometric-neurometric experiments showed that the subjective intensity of a somatosensory stimulus and the activity of somatosensory nerves were both similar power functions of objective stimulus intensity[28]. Used in this manner, a psychometric-neurometric comparison does more than predict behavior — it tests a specific relationship, or linking proposition[36], between psychological and physiological states. It tests the hypothesis that the rate of neural activity (or some other candidate code) is linearly related to, and thus can serve as a representation of, the magnitude of a subjective mental experience.

Here we adapt the psychometric-neurometric technique to determine whether any brain areas encode the subjective value of delayed monetary rewards to individual human decision-makers. Subjects made repeated choices between immediate and delayed rewards while we simultaneously measured neural activity using functional magnetic resonance imaging (fMRI). We found a clear match between

Center for Neural Science, New York University, 4 Washington Place, Room 809, New York, New York 10003, USA. Correspondence should be addressed to J.W.K. (kable@cns.nyu.edu) or P.W.G. (glimcher@cns.nyu.edu).

Received 7 June; accepted 4 October; published online 4 November; corrected online 11 November 2007 (details online); doi:10.1038/nn2007

the subjective preferences of our subjects and neural activity in the ventral striatum, medial prefrontal cortex and posterior cingulate cortex. These data provide unambiguous evidence that idiosyncratic preference functions are part of the neural mechanism for choice, as opposed to being a purely descriptive ('as if') theoretical construct, an idea that has been widely proposed in economics[37–39]. From an economic perspective, these findings indicate the existence of a neural valuation process that is strikingly similar to the representations of subjective values that are employed in revealed preference theories. From a neurobiological perspective, these findings implicate particular neural structures in the subjective valuation of outcomes during decision-making.

RESULTS

Psychometric results

On each trial, subjects chose between a fixed immediate reward of $20 and a larger delayed reward that varied randomly from trial to trial (**Fig. 1**). The value of the larger reward varied from $20.25 to $110, and the delay varied from 6 h to 180 d. Before fMRI, subjects completed three preliminary behavioral sessions. On the basis of the choices in these behavioral sessions, we estimated a discount function for each subject (**Fig. 2a–i**). Note that interpreting **Figure 2g–i** as pure discount functions relies on the assumption that subjective value increases linearly with the objective amount of money. Violations of this assumption would not alter the fact that these curves reflect subjective value, but would mean that they are better described as discounted utility functions as they incorporate effects of amount as well as delay. These individual preference curves, which express how a subjective quantity (subjective value) changes as a function of an objective variable (delay and monetary amount), are directly analogous to sensory psychophysical functions[2,3]. Consistent with previous findings in a range of species[30–34], the discount curves for all subjects were well–characterized by a hyperbolic function (median R^2 from scanning sessions = 0.95, range = 0.84–0.98):

$$SV = \frac{1}{1 + kD} \quad (1)$$

where SV is subjective value (expressed here as a fraction of immediate value), D is delay (in d) and k is a subject-specific constant. The best fitting k parameter (the discount rate) varied widely across subjects. For our most patient subject ($k = 0.0005$, **Fig. 2a**,**d**,**g**), an immediate gain of $20 was less preferable than a gain of $21 in a month, while for our most impulsive subject ($k = 0.1189$, **Fig. 2c**,**f**,**i**), an immediate gain of $20 was more preferable than a gain of $68 in a month. Ten subjects (out of twelve) showed stable discount rates across sessions, even though sessions occurred over a 1–6-month time span (**Supplementary Fig. 1** online). Only these ten subjects, with relatively stable functions over the tested range, were included in the fMRI experiment.

For comparison, we also fit each subject's discount curve with a single exponential function:

$$SV = e^{-cD} \quad (2)$$

as well as a sum of two exponential functions:

$$SV = \frac{\left(e^{-\beta D} + e^{-\delta D}\right)}{2}, \quad (3)$$

where SV is subjective value and D is delay as in equation (1), and c, β and δ are subject-specific constants. Consistent with the literature, the hyperbolic function equation (1) described the discount curves better than the single parameter exponential function equation (2) (**Supplementary Fig. 2** online). The sum of two exponentials is one formulation of an economic model that explains hyperbolic-like discounting by the combined action of two systems[32,35]: one (β) that discounts more steeply than the person's resulting behavior, and a second (δ) that discounts less steeply than the person's resulting behavior. Not surprisingly, as this functional form can account for hyperbolic-like discounting, the sum of exponentials function fit the data better than the single parameter exponential, and as well as the single parameter hyperbolic function (**Supplementary Fig. 2**). Although other forms of the discount function have been proposed[32–34], for simplicity, we considered only these three widely used forms.

Neurometric results

The ten subjects with stable discount functions completed the same choice task as was used in the preliminary behavioral sessions during fMRI. The behaviorally derived preference curves, which estimated how subjective value declined with delay for each individual, were used to find neural activity that correlated with subjective value. The subjective value of the delayed reward on each trial was estimated by multiplying the objective amount of the reward by the discount fraction (estimated behaviorally) for that delay. This inferred subjective value was then used as a regressor for brain activity. As the immediate reward was constant, this analysis finds areas that track the sum of the subjective values of both rewards, the difference in subjective value of the two rewards, or the ratio of subjective values of the two rewards.

Early in the trial, there was a significant correlation between neural activity and subject-specific subjective value (random effects group analysis, voxel-wise $P < 0.001$, spatial extent > 100 mm^3) in the ventral striatum, medial prefrontal cortex and posterior cingulate cortex (**Fig. 3a**, **Supplementary Table 1** online). The timing of this effect (6–10 s into the trial) indicated that it arose from neuronal activity during the visual presentation of the delayed reward option, which we confirmed in a second model that assumed a hemodynamic

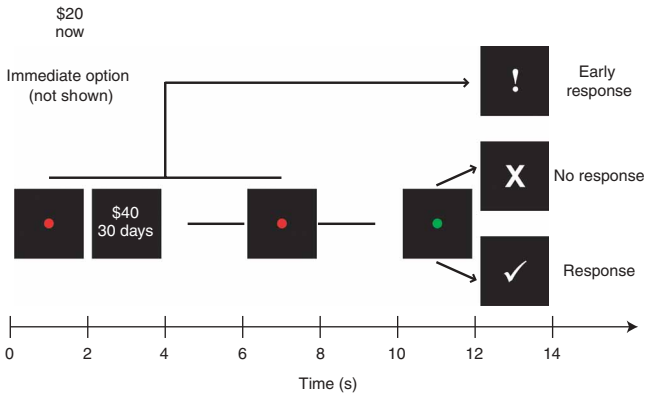

Figure 1 Intertemporal choice task. The sequence of events within a trial is shown. On each trial, subjects chose between an immediate and a delayed reward. The immediate reward was the same ($20) on every trial and was never presented visually. A red dot signaled the beginning of a trial, after which the delayed reward for that trial was presented for 2 s and then replaced again by the red dot. Subjects then had 6 s to consider their choice. Throughout the trial, subjects were required to hold down a button, and they indicated their decision by releasing or continuing to hold the button when the dot turned green. For half of the session, a button release indicated a choice of the delayed reward; for the other half, a button release indicated a choice of the immediate reward. If subjects released the button before the green light appeared, that trial was considered abandoned and removed from the analysis. The inter-trial interval was 2 s for behavioral sessions and 12 s for scanning sessions.

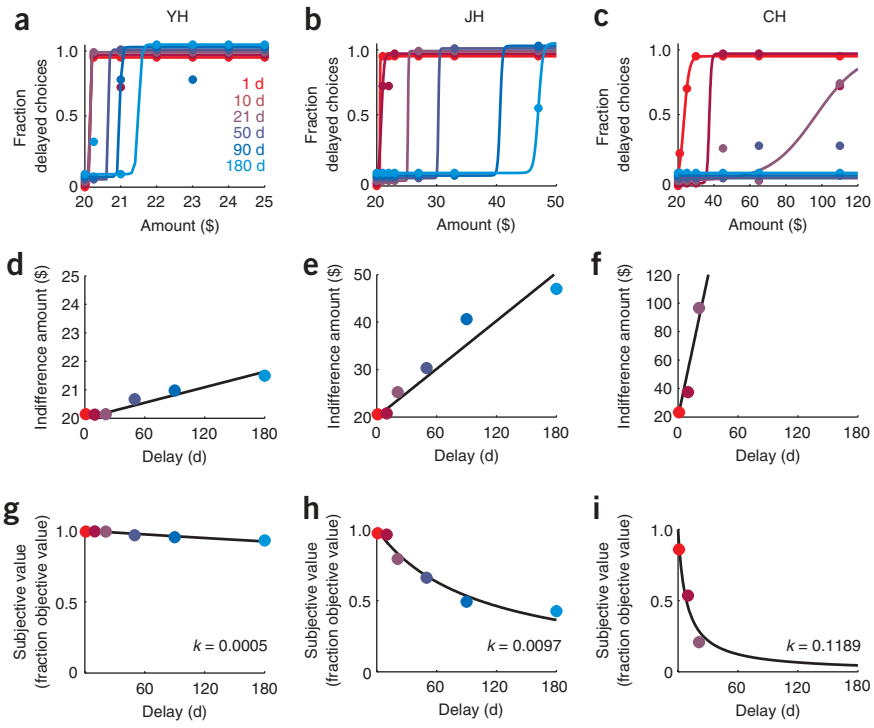

Figure 2 Subject-specific discount functions. (**a**–**c**) Choice data from three subjects during a single scanning session. Points are shaded according to the imposed delay to the delayed reward, and denote the fraction of times the subject chose the delayed reward over an immediate reward of $20 as a function of the objective amount of the delayed reward. The smooth curves are logistic functions fit to these data. Data from different delays are slightly offset so that all data are visible. (**d**–**f**) Indifference points, plotted as a function of the imposed delay to the delayed reward. Indifference points were estimated from the logistic fit in **a**–**c** as the amount for each delay at which the subject would choose the immediate and delayed rewards with equal frequency. The increase in indifference amounts with delay was fit by a line with a fixed intercept at $20. Delays are shaded as in **a**–**c**. (**g**–**i**) Indifference points from (**d**–**f**), divided into $20 to obtain a discount function. The decrease in subjective value with delay was fit with a single-parameter hyperbolic function. Delays are shaded as in **a**–**c**. Data from three subjects are shown (YH in **a**,**d**,**g**; JH in **b**,**e**,**h**; CH in **c**,**f**,**i**) to illustrate the observed heterogeneity in discount functions across subjects. CH was our most impulsive subject ($k = 0.1189$), YH was our most patient subject ($k = 0.0005$), and JH was near the median discount rate ($k = 0.0097$).

response function (**Supplementary Table 2** online). We did not find any significant effects of subjective value during other portions of the trial.

A region in which neural activity is correlated with the subjective value of a delayed reward should show increased activity when the objective amount of this reward increases, decreased activity when the delay to this reward increases, and increased activity when the delayed reward is chosen because it is more valuable. This is exactly what we observed in these three regions (**Fig. 3b–d**, **Supplementary Fig. 3** online and **Supplementary Tables 3–5** online). Subjective value accounted for activity in these regions better than did the objective reward characteristics or the subjects' choices: both the strength (peak z-score) and spatial extent (cluster size in mm^3) of the correlation with subjective value were larger than those of these other variables (**Fig. 3a–d** and **Supplementary Tables 1–5**). Of these other variables, only delay-to-reward reached significance, and only in the medial prefrontal cortex. Furthermore, taking all voxels identified by any of these four variables (subjective value, monetary amount, delay and choice), subjective value showed a stronger effect in pair-wise comparisons with these other variables in 100% of voxels in the ventral striatum, at

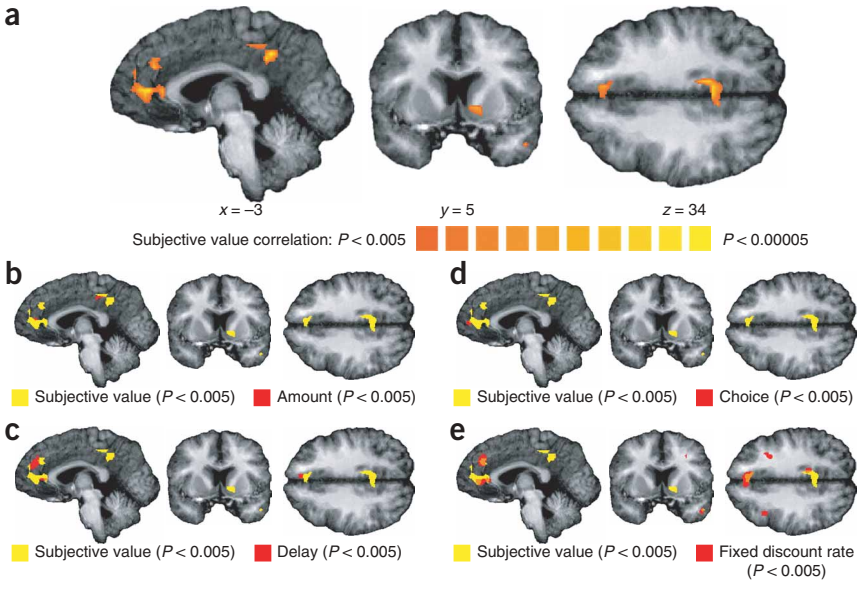

Figure 3 Group analysis showing areas in which activity is correlated with subjective value. (**a**) Areas in which neural activity was correlated with subjective value (during the 6–10-s window) in a random-effects group analysis, overlayed on the mean normalized anatomical image. Areas of correlation can be seen in the medial prefrontal cortex and posterior cingulate cortex (sagittal and axial images) and in the ventral striatum (coronal image). The color scale represents the t-value of the contrast testing for a significant effect of subjective value at time points 4–6 in the trial. (**b**–**e**) Activity in the ventral striatum, medial prefrontal cortex and posterior cingulate cortex was better correlated with subjective value than with (**b**) the objective amount of the delayed reward, (**c**) the inverse delay of the delayed reward, (**d**) the choice of the subject (chose delayed > chose immediate), or (**e**) the value of the delayed reward calculated using a single fixed discount rate for all subjects ($k = 0.01$, near the median for our subjects). Areas in which activity correlated with subjective value are shown in yellow, areas in which activity correlated with the other variables are shown in red, and areas of overlap are shown in orange. All maps are thresholded at $P < 0.005$ (uncorrected), spatial extent > 100 mm^3. Data are shown in radiological convention, with the right hemisphere on the left.

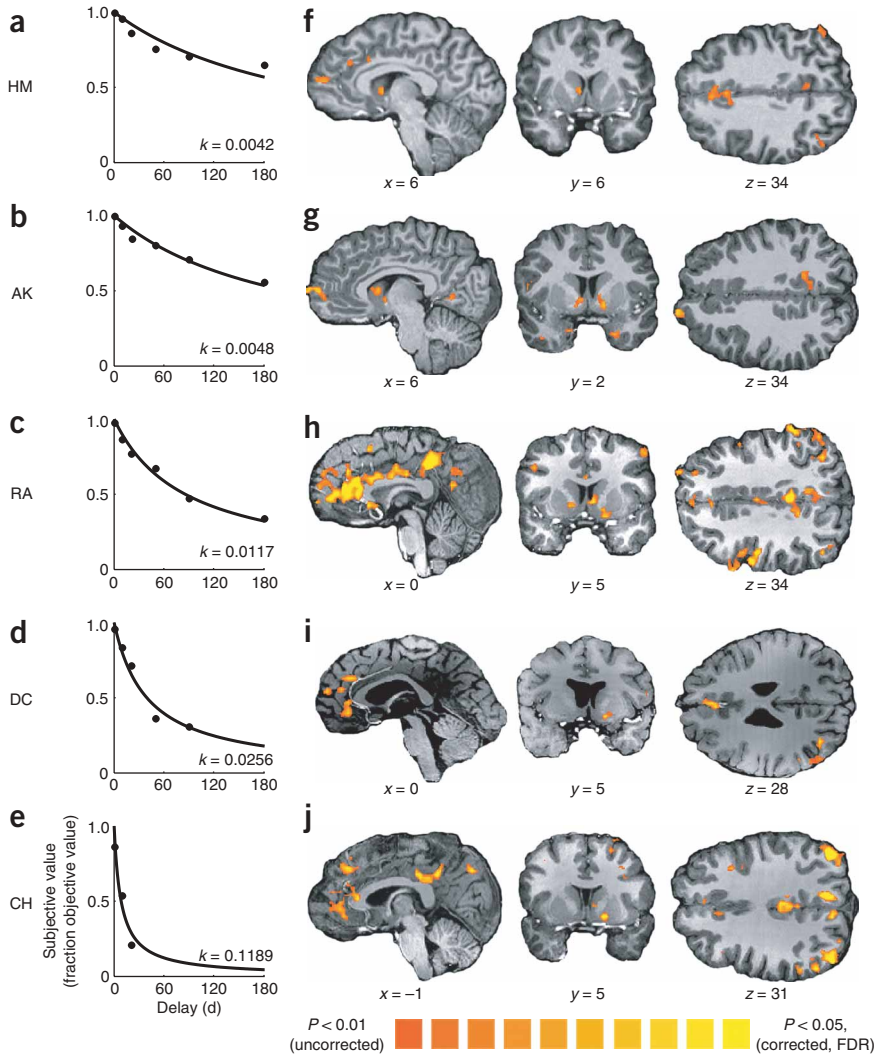

Figure 4 Single-subject analyses showing areas in which activity correlated with subjective value. Data from five subjects are shown to illustrate that subjective value effects were evident for subjects spanning the entire range of behavioral discount rates. (**a**–**e**) Discount functions for five subjects, as measured behaviorally during a scanning session. These subjects include one of our more patient subjects (HM) and our most impulsive subject (CH). (**f**–**j**) Maps showing areas in which neural activity was correlated with subjective value in each subject. The sagittal overlay shows areas of activity in the medial prefrontal cortex, the coronal overlay areas of activity in the ventral striatum, and the axial overlay areas of activity in the posterior cingulate cortex. The color scale represents the *t*-value of the contrast testing for a significant effect of subjective value at time points 4–6 in the trial. These maps are thresholded at $P < 0.01$ (uncorrected), and the color scale ranges from this value to $P < 0.05$ (corrected for false discovery rate). Data are shown in radiological convention, with the right hemisphere on the left.

least 85% of voxels in the posterior cingulate cortex, and at least 62% of voxels in the medial prefrontal cortex. Idiosyncratic subjective value also accounted for activity in these regions better than did assuming a single fixed discount rate for all subjects (**Fig. 3e** and **Supplementary Table 6** online). The effect of value, assuming a single fixed discount rate, only reached significance in the medial prefrontal cortex, and the strength and spatial extent of the correlation with subject-specific subjective value was larger in both the ventral striatum and the posterior cingulate cortex. Finally, these effects of subjective value cannot be explained by choice difficulty or corresponding attentional demands. Using two different indices of choice difficulty, we found no significant effects in the ventral striatum, medial prefrontal cortex or posterior cingulate cortex (**Supplementary Tables 7** and **8** online).

We also found significant effects of subjective value in single-subject analyses, for subjects across the entire spectrum of discount rates. Despite marked differences in the discount functions of the five example subjects shown in **Fig. 4a**–**e**, ranging from patient ($k = 0.0042$) to impulsive ($k = 0.1189$), activity in parts of the ventral striatum, medial prefrontal cortex and posterior cingulate cortex were correlated with subjective

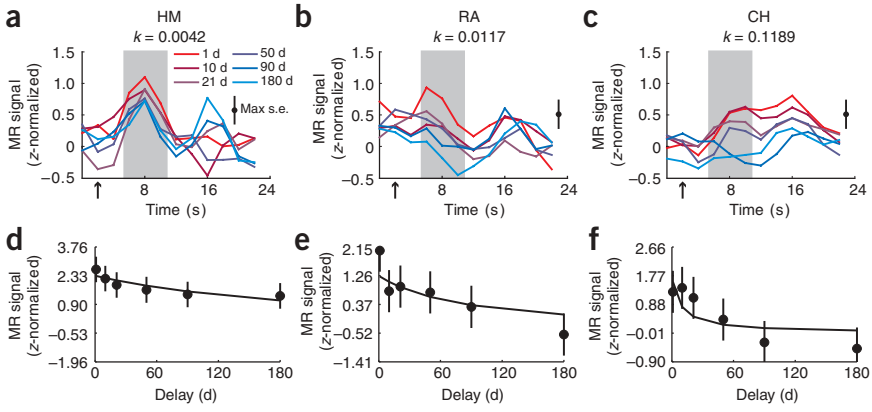

Figure 5 Single-subject time courses and neural discount functions. (**a**–**c**) Data from three subjects (HM, see **Fig. 4a,f**; RA, see **Fig. 4c,h**; and CH, see **Fig. 4e,j**) are shown. Data were averaged over all voxels that showed a correlation between activity and subjective value in the ventral striatum, medial prefrontal cortex and posterior cingulate cortex (from the individual-level analyses shown in **Fig. 4**), and then re-plotted as trial averages. Trial averages are color-coded by the imposed delay to the delayed reward. The 6–10-s window in which we observed significant effects is shown in gray. The largest standard error is shown on the right. The arrows indicate the point in the trial at which the delayed option was presented. (**d**–**f**) Data from **a**–**c**, summed over the 6–10-s window and re-plotted as a function of delay. The solid black line represents average predicted activity at each delay, from the fit of the subjective value regression using a subject-specific discount rate. Predicted activity is simply a scaled and shifted version of each subject's behavioral discount function. This regression is also used to scale the *y*-axis across subjects (see **Supplementary Methods**).

value in each subject (**Fig. 4f–j**). These same regions thus exhibited a different pattern of activity across subjects, with each subject's idiosyncratic pattern of brain activity being predicted by that subject's idiosyncratic preferences (**Fig. 5a–f**, **Supplementary Fig. 4** online). As a function of delay, both the subjective value of delayed gains and the neural activity associated with those gains decreased in a similar manner, whether subjects were more impulsive and showed steeper decreases in both domains (**Fig. 5c**,**f**) or were less impulsive and showed more gradual decreases in both domains (**Fig. 5a**,**d**).

To quantify the degree to which these regions show a different pattern of activity across subjects, we averaged activity from all voxels that showed a subjective value effect in each of the three regions in each subject, and then regressed neural activity on each trial (summed over the 6–10-s window) separately against the imposed delay and objective amount associated with the variable delayed reward. If neural activity tracks subjective value across subjects, then the ratio of these two regression coefficients (delay over amount, reversing the sign on the negative delay coefficient) should be high for steep discounters and low for shallow discounters, as delay has a stronger effect on subjective value for more impulsive discounters. Consistent with this notion, this ratio increased as the subject's discount rate increased (**Fig. 6a**; $\text{slope}_{VS} \pm \text{standard error} = 0.77 \pm 0.19$, t-test, $P = 0.004$; $\text{slope}_{MPFC} = 0.90 \pm 0.21$, $P = 0.003$; $\text{slope}_{PCC} = 0.99 \pm 0.31$, $P = 0.02$; $\text{slope}_{ALL} = 0.87 \pm 0.12$, $P < 10^{-7}$; VS, ventral striatum; MPFC, medial prefrontal cortex; PCC, posterior cingulate cortex).

The preceding analysis demonstrates the correspondence between neural activity and discount rate across subjects without making strong assumptions about the functional form of the discount function. For a more precise test, we assumed that the neural discount function has a hyperbolic form and estimated the neural discount rate for these areas, by refitting the regression model to the averaged time courses from the correlated regions while allowing the discount rate parameter (k) to vary. A psychometric-neurometric match requires both that the neural discount rate increases with the subject's behavioral discount rate and that there is no difference between the two on average. By contrast, if activity in these regions were best explained by the objective amount (delay) of the delayed reward, then the neural discount rate should be

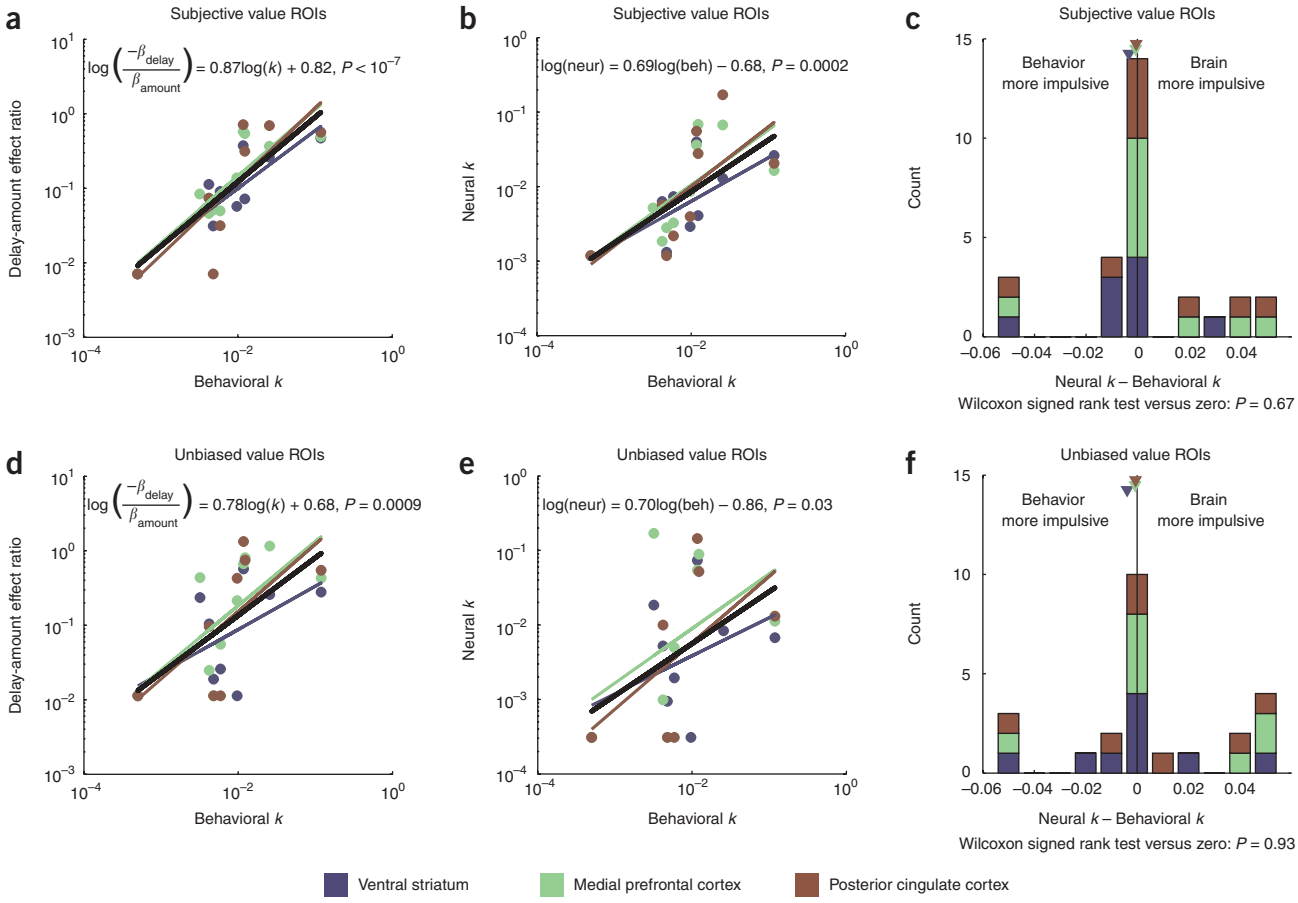

Figure 6 Psychometric-neurometric comparisons. (**a**,**d**) A measure of the neural effect of delay is plotted against the subject's behavioral discount rate for both (**a**) ROIs defined on the basis of the subjective value regression and (**d**) value ROIs defined in an unbiased manner. The neural effect of delay used is the ratio of the regression slope of neural activity against delay compared to that against amount. (**b**,**e**) The neural discount rate is plotted against the subject's behavioral discount rate for both (**b**) subjective value ROIs and (**e**) unbiased value ROIs. The neural discount rate is estimated using a nonlinear version of the subjective value regression where k can vary. Colored lines show the robust linear fit for each ROI; black line shows the fit collapsed across all ROIs. Because of their skewed distribution, the ratio and neural and behavioral discount rates are log-transformed. All three data points from the most patient subject (bottom left) overlap in panels **a**,**b**,**d** and **e**. (**c**,**f**) Difference between neural discount rate and behavioral discount rate for both (**c**) subjective value ROIs and (**f**) unbiased value ROIs. Colored triangles indicate the median differences for each ROI. Panels **a**–**c** exclude two and **d**–**f** exclude one subject-ROI pair where an ROI could not be defined. Panel **d** also excludes two subject-ROI pairs where the correlation with amount was negative and **e**–**f** exclude five subject-ROI pairs where no discount rate accounted for a significant amount of variance in neural activity (see Methods).

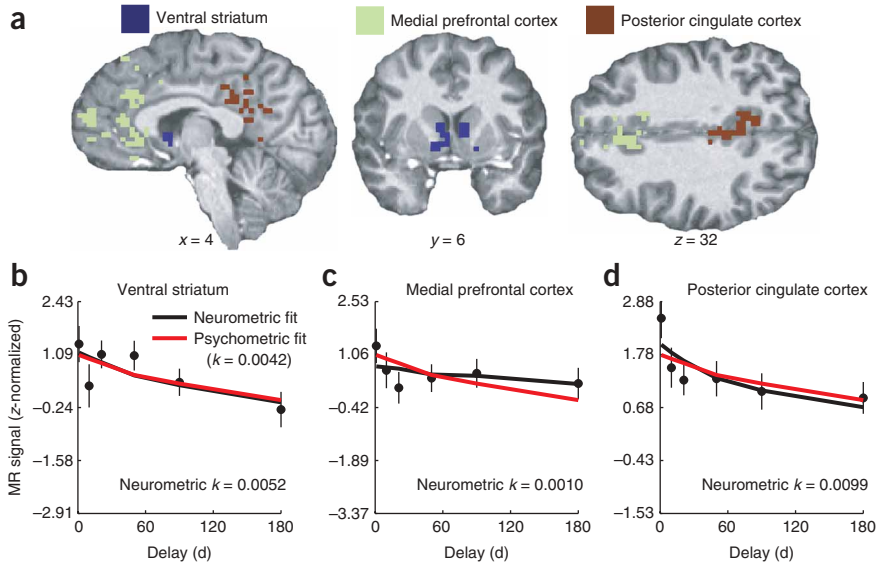

Figure 7 Single-subject example of unbiased value ROIs and resulting neural discount functions. (**a**) ROIs in the ventral striatum, medial prefrontal cortex and posterior cingulate cortex are shown for one subject (HM, same subject shown in **Fig. 4a**,**f** and **Fig. 5a**,**d**), who demonstrated a relatively close psychometric-neurometric match in each region. Voxels were selected within these anatomically defined regions which showed either greater activity for trials involving the largest objective amount of the delayed reward than for trials involving the smallest amount, or greater activity for trials involving the shortest delay to the delayed reward than for those involving the longest delay. (**b**–**d**) A psychometric-neurometric comparison is shown for each ROI for this subject. Mean neural activity and standard error, summed over the 6–10-s window, are plotted as a function of the imposed delay to the delayed reward. The red line represents average predicted activity at each delay, from the fit of the subjective value regression using the subject's behavioral discount rate ($k = 0.0042$). This regression is also used to scale the *y*-axis across ROIs (see Methods). The black line represents average predicted activity at each delay, from the fit of a nonlinear version of the subjective value regression where the discount rate is allowed to vary. The discount rates estimated from the neural data are 0.0052 in the ventral striatum, 0.0010 in the medial prefrontal cortex and 0.0099 in the posterior cingulate cortex.

consistently smaller (greater) than the behavioral discount rate. If activity were best explained by a single fixed discount rate across all subjects, then the neural discount rate should show no correlation with the behavioral discount rate. Consistent with a psychometric-neurometric match, the neural discount rate in these three regions increased with the subject's behavioral discount rate (**Fig. 6b**, $\text{slope}_{VS} \pm$ standard error $= 0.58 \pm 0.23$, *t*-test, $P = 0.04$; $\text{slope}_{MPFC} = 0.75 \pm 0.28$, $P = 0.03$; $\text{slope}_{PCC} = 0.80 \pm 0.37$, $P = 0.06$; $\text{slope}_{ALL} = 0.69 \pm 0.16$, $P = 0.0002$), and the difference between the neural discount rate and the subject's behavioral discount rate was centered on zero (**Fig. 6c**; $\text{median}_{VS} = -0.0035$, Wilcoxon signed rank test versus zero, $P = 0.30$; $\text{median}_{MPFC} = -0.0008$, $P = 1$; $\text{median}_{PCC} = -0.0006$, $P = 1$; $\text{median}_{ALL} = -0.0009$, $P = 0.67$).

As the previous analyses were conducted in regions-of-interest (ROIs) that were defined by the subjective value regression, they merely confirm that the regions identified exhibit a psychometric-neurometric match. It would be even more compelling if such a match existed in value-sensitive voxels in these regions that were selected in an unbiased manner. We defined value-sensitive ROIs in these areas in a manner that was unbiased with respect to the neural discount rate (see Methods). In these ROIs, the negative effect of delay, relative to the positive effect of amount, again increased with the subject's behavioral discount rate (**Fig. 6d**; $\text{slope}_{VS} \pm$ standard error $= 0.58 \pm 0.36$, *t*-test, $P = 0.15$; $\text{slope}_{MPFC} = 0.86 \pm 0.36$, $P = 0.046$; $\text{slope}_{PCC} = 0.90 \pm 0.46$, $P = 0.10$; $\text{slope}_{ALL} = 0.78 \pm 0.21$, $P = 0.0009$). In addition, consistent with a psychometric-neurometric match, the neural discount rate in these ROIs increased with the subject's behavioral discount rate (**Fig. 6e**; $\text{slope}_{VS} \pm$ standard error $= 0.51 \pm 0.48$, *t*-test, $P = 0.32$; $\text{slope}_{MPFC} = 0.73 \pm 0.67$, $P = 0.31$; $\text{slope}_{PCC} = 0.89 \pm 0.66$, $P = 0.23$; $\text{slope}_{ALL} = 0.70 \pm 0.31$, $P = 0.03$), and the difference between the neural and behavioral discount rates was centered on zero (**Fig. 6f**; $\text{median}_{VS} = -0.0039$, Wilcoxon signed rank test versus zero, $P = 0.50$; $\text{median}_{MPFC} = -0.0008$, $P = 0.95$; $\text{median}_{PCC} = -0.0006$, $P = 0.94$; $\text{median}_{ALL} = -0.0009$, $P = 0.93$; see **Fig. 7a–d** for a single-subject example).

Finally, to evaluate an alternative hypothesis about the role of these regions in intertemporal choice[32,35], we repeated the above analyses using the β (steep exponential) and δ (shallow exponential) value functions from the sum of exponentials discount function fit to each subject's behavior. Perhaps not surprisingly, as the β and δ value functions are correlated to some degree with subjective value, the brain regions in which activity is correlated with these three value functions are largely overlapping (**Supplementary Fig. 5** online, **Supplementary Tables 9–11** online). A psychometric-neurometric comparison, though, provides a more precise test of which value function best accounts for neural activity. As β and δ are exponential functions, we estimated the neural discount rates using a single exponential (rather than hyperbolic) function, and compared this rate to the behavioral discount rate estimated using a single exponential and the exponential β and δ discount rates estimated behaviorally. In both sets of ROIs, the difference between the neural discount rate and the behavioral discount rate was again centered on zero (**Fig. 8a**,**b**; subjective value ROIs, median $= -0.0006$, Wilcoxon signed rank test versus zero, $P = 0.91$; unbiased value ROIs, median $= -0.0006$, $P = 0.77$). By contrast, the difference between the neural discount rate and β is centered to the left of zero, indicating that the theoretically defined β term discounts more steeply than neural activity in these areas (**Fig. 8c**,**d**; subjective value ROIs, median $= -0.0081$, $P = 0.001$; unbiased value ROIs, median $= -0.0088$, $P = 0.23$), while the difference between the neural discount rate and δ is centered to the right of zero, indicating that the theoretically defined δ term discounts less steeply than neural activity (**Fig. 8e**,**f**, subjective value ROIs, median $= 0.0024$, $P < 0.0001$; unbiased value ROIs, median $= 0.0016$, $P = 0.003$). Thus, neural activity in the ventral striatum, medial prefrontal cortex and posterior cingulate cortex tracks the subjective value of rewards as determined from behavior, rather than tracking a theoretically defined component of value that is more impulsive (β) or more patient (δ) than the person's behavior.

DISCUSSION

Revealed preference theories in economics posit that decision-makers behave as though different options have different subjective values[4,37–39]. Here we have shown that neural activity in several brain

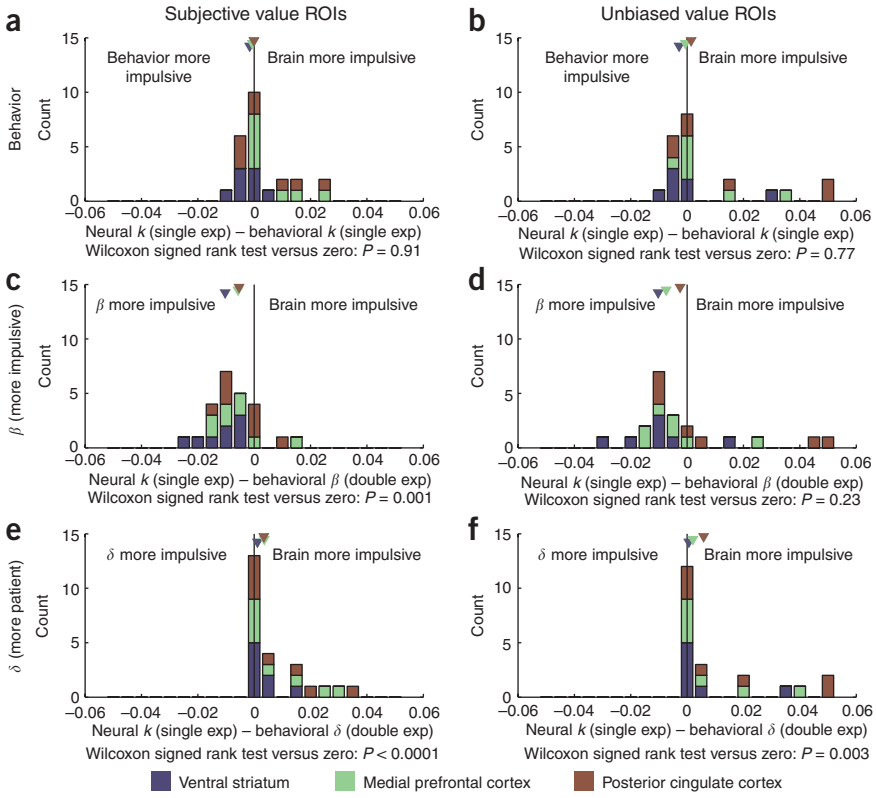

Figure 8 Neural activity tracks subjective value, and not a more impulsive (β) or more patient (δ) estimate of value. (**a**,**b**) Difference between the neural (single exponential) discount rate and the behavioral (single exponential) discount rate, calculated separately for each ROI in each subject. The difference between neural and behavioral discount rates is centered on zero for both (**a**) ROIs defined on the basis of the subjective value regression and (**b**) value ROIs defined in an unbiased manner. (**c**,**d**) Difference between the neural (single exponential) discount rate and β, the steeper exponential from the sum of exponentials discount function estimated behaviorally. On average, β is larger than the neural discount rate for both (**c**) subjective value ROIs and (**d**) unbiased value ROIs. (**e**,**f**) Difference between the neural (single exponential) discount rate and δ, the shallower exponential from the sum of exponentials discount function estimated behaviorally. On average, the neural discount rate is larger than δ for both (**e**) subjective value ROIs and (**f**) unbiased value ROIs. Colored triangles in each panel indicate the medians for each ROI separately. These data exclude two subjects (six subject-ROI pairs) for which the fit of the β-δ model collapsed to a single exponential function (that is, $\beta = \delta$). Panels **b**,**d** and **f** also exclude two subject-ROI pairs for which no discount rate accounted for a significant amount of variance in neural activity (see Methods).

regions—particularly the ventral striatum, medial prefrontal cortex and posterior cingulate cortex—tracks the revealed subjective value of delayed monetary rewards. Activity in these regions increases as the objective amount of a reward increases and decreases as the imposed delay to a reward increases. Furthermore, across subjects, the neural tradeoffs between amount and delay that are captured by the neurometric discount functions match the behavioral tradeoffs between these variables that are captured by the psychometric discount functions. These results indicate that choosing between immediate and delayed monetary rewards involves, at a physical level, comparing neurally encoded subjective values.

Our suggestion that the ventral striatum, medial prefrontal cortex and posterior cingulate cortex help to value immediate and delayed outcomes is consistent with a large body of existing evidence. The medial prefrontal and posterior cingulate cortices are part of a 'baseline' cortical system that is important for behavior mediated by internal goals[40,41]. Furthermore, all three regions show increased activity after subjects receive a reward or immediately before an expected reward. This reward-related activity correlates with reward size and probability[6–16], and is greater for more preferred rewards than for less preferred rewards[17,18,20,21]. In reinforcement learning tasks, activity in the medial prefrontal cortex is correlated with the predicted reward value estimated from behavioral models[22,24]. In choice tasks, these three areas show greater activity during choices that involve large gains than during choices that involve small gains[13,23,42,43], as well as greater activity during choices that involve small losses than during choices that involve large losses[44]. Our results extend these findings in two ways. First, we show that the ventral striatum, medial prefrontal and posterior cingulate cortex also participate in the subjective valuation of delayed rewards, rather than just of immediate gains and losses. Second, by quantitatively estimating both behavioral and neural discount functions, we show that activity in these regions precisely tracks the subjective value of possible rewards during choice — in other words, these regions exhibit a psychometric-neurometric match.

Our findings falsify a hypothesis regarding the neurobiological basis of intertemporal choice[35]. In ref. 35, the authors hypothesized that these same three regions — the ventral striatum, medial prefrontal cortex and posterior cingulate cortex — form an impulsive neural system that exclusively or primarily values immediate rewards. Drawing an analogy to a valuable economic model[32], they suggested that hyperbolic-like discounting resulted from the action of this impulsive (β) system working with a more patient (δ) system, based in the lateral prefrontal cortex and posterior parietal cortex. This conclusion was based principally on the finding that these three areas showed greater activity for choices that involved an immediate reward than for choices that involved only delayed rewards. However, this empirical finding, which our current results do not challenge, is also compatible with the hypothesis that activity in these regions represents the subjective value of rewards at all delays, as the subjective value of immediate rewards is greater than that of delayed rewards. Our observation that activity in these regions varies when only the delayed reward changes falsifies the hypothesis that these regions exclusively value immediate rewards. Furthermore, our finding that neural activity in these regions tracks changes in the subjective value of the delayed reward (determined by subject-specific behavioral discount functions) shows that these regions do not even primarily value immediate rewards, as the value implied by neural activity is not more impulsive than the person's behavior, as the β–δ hypothesis requires.

The fact that the ventral striatum, medial prefrontal cortex and posterior cingulate cortex encode the subjective value of delayed monetary rewards, whether those rewards are selected by pressing a button or by taking no action, indicates that these regions are involved in the valuation of reward outcomes rather than in movement control. If this is correct, then these areas might provide input to structures that

encode the subjective value of particular actions, such as the movement-related areas of the posterior parietal cortex[5,19,45,46]. Consistent with this notion, activity in the parietal cortex was correlated with subjective value in one analysis (**Supplementary Table 2**). Future studies should more closely examine the interaction between outcome (or goods-based[23]) and action values, as well as the process by which value signals affect choice. The three regions we have focused on might carry out different roles in this process, and such investigations might also shed light on the function of other regions (such as the temporal-parietal junction and insula) that were identified in some of our analyses.

A few issues suggest some caution when interpreting our results. First, unlike ref. 35, we did not include an experimental condition that involved choices between two delayed rewards. Future studies should include such choices to verify that these regions do not encode subjective value in a categorically different manner when an immediate reward is not available. By varying both rewards, such studies could also address how activity in these regions reflects the combination of subjective values from multiple rewards presented simultaneously. In this experiment, as one option was always constant, areas that implemented different combination rules (addition, subtraction, division) or that encoded the subjective value of the chosen option would all show similar patterns of activity in our analyses. In addition, we only scanned individuals with stable discount functions. It remains unknown what activity in these regions looks like for individuals who show a different kind of choice behavior — for example, those who follow a simple heuristic or cut-off rule.

In the current study we adapted the method of psychometric-neurometric comparisons, which has been widely used in perceptual neurophysiology, for use with fMRI. Although this approach is related to previous model-based fMRI studies[22,24,25], few studies have directly estimated neurometric functions from fMRI data. Our results indicate that such psychometric-neurometric comparisons should prove fruitful in studying the neural mechanisms of decision-making, and hold promise for contributing to both neuroscience and economics. For neuroscience, this method can link the function of different neural systems with the hidden processes that are posited in quantitative models of decision-making. For economics, this method provides preliminary evidence that the physical mechanism that underlies intertemporal choice is strikingly similar to the most general class of revealed preference models, resting on a continuous underlying scale for subjective value. Although this similarity does not prove or disprove theoretical models that make no predictions about the physical mechanisms of choice, it does underline the conviction of many that theoretical models resting on a demonstrable physiological foundation will probably prove more robust than those that lack such a foundation.

METHODS

Subjects. Twelve paid volunteers participated in the experiment (7 women, 5 men, all right-handed, mean age = 21.7 years). Ten (6 women, 4 men, mean age = 21.2 years) participated in both behavioral and scanning sessions. The remaining two participated only in behavioral sessions, as their intertemporal preferences were not stable. One of these subject's choices reversed more than once as the delayed amount increased, precluding estimation of an indifference point; the other had significantly different discount rates in their last two behavioral sessions. All participants gave informed consent in accordance with the procedures of the University Committee on Activities Involving Human Subjects of New York University.

Task. On each trial, subjects chose between a smaller amount of money paid immediately and a larger amount paid at a later time (**Fig. 1**). The smaller immediate amount was $20 on all trials. To simplify the display, only the variable delayed option was presented on each trial. This option was constructed using one of six delays (6 h–180 d) and one of six amounts ($20.25–$110). Thus there were 36 unique choices in each session. Each of these was presented four times, for a total of 144 trials. Delays were the same for every subject, but changed across sessions; amounts were chosen individually for each subject on the basis of their behavior in previous sessions (and thus also varied across sessions) to ensure that an approximately equal number of immediate and delayed options were chosen. Subjects were not told how options would change across sessions, and we did not observe any systematic shifts in subject's choices in response to these changes (**Supplementary Fig. 1**). Subjects indicated their choices by releasing a button. For half of each session, subjects released the button to select the fixed immediate reward, and for the other half subjects released the button to select the variable delayed reward. A button release was used, rather than a button press, to keep subjects alert on trials in which their choice did not require a response. All subjects participated in three 1-h behavioral sessions, and ten participated in one or two 2-h scanning sessions. Sessions were completed over 1–6 months, with consecutive sessions separated by at least 3 d. As people's decisions involving real gains might differ from those involving hypothetical gains[47–49], subjects were paid according to four randomly selected trials per session (except for the first behavioral session, which involved only hypothetical choices), using commercially available pre-paid debit cards (see **Supplementary Methods** online for details).

Imaging. Imaging data were collected with a Siemens Allegra 3T head-only scanner equipped with a head coil from Nova Medical. T2*-weighted functional images were collected using an EPI sequence (TR = 2 s, TE = 30 ms, 35 axial slices acquired in ascending interleaved order, 3 × 3 × 3 mm, 64 × 64 matrix in a 192-mm field of view). Each scan consisted of 236 images. The first two images were discarded to avoid T1 saturation effects. There were eighteen choice trials during each scan; each trial lasted 14 s with a 12-s inter-trial interval (**Fig. 1**). Each subject completed 8–16 scans over 1–2 sessions, with most subjects ($n = 7$) completing eight scans in one session. High-resolution, T1-weighted anatomical images were also collected using an MPRAGE sequence (TR = 2.5 s, TE = 3.93 ms, TI = 900 ms, flip angle = 8°, 144 sagittal slices, 1 × 1 × 1 mm, 256 × 256 matrix in a 256-mm field of view).

Data analysis. Functional imaging data were analyzed using BrainVoyager QX (Brain Innovation), with additional analyses performed in MATLAB (MathWorks). Functional images were sinc-interpolated in time to adjust for staggered slice acquisition, corrected for any head movement by realigning all volumes to the first volume of the scanning session using six-parameter rigid-body transformations, and de-trended and high-pass filtered (cutoff of 3 cycles per scan, or 0.0064 Hz) to remove low-frequency drift in the fMRI signal. Images were then co-registered with each subject's high-resolution anatomical scan, rotated into the AC-PC plane, and normalized into Talairach space using piecewise affine Talairach grid scaling. All spatial transformations of the functional data used trilinear interpolation. For group-level random effects analyses only, data were also spatially smoothed with a gaussian kernel of 8 mm (full-width half-maximum).

Single-subject analyses were performed using multiple linear regression, estimated with ordinary least-squares. As we did not want to assume a priori where in a trial subjective value might have effects, we used a variant of a 'finite-impulse response' or 'deconvolution' model. This model included twelve covariates that fit the mean activity, across all trials, for each of the first twelve time points in a trial, and twelve covariates that fit the deviations from the mean at each time point that were correlated with the subjective value of the delayed reward across trials. Additional models (**Supplementary Tables 3–11**) replaced the subjective value terms with an alternative variable (see **Supplementary Methods**).

We performed group random-effects analyses using the summary statistics approach, which tests whether the mean effect at each voxel is significantly different from zero across subjects. On the basis of previous reports[22,35], we hypothesized a priori that the ventral striatum, medial prefrontal cortex and posterior cingulate cortex would contribute to valuation during intertemporal choice. Accordingly, we considered any activation in these regions significant if it exceeded $P < 0.001$ (uncorrected), spatial extent > 100 mm^3. For display purposes only, contrast maps were interpolated to 1 × 1 × 1 mm, and thresholded at $P < 0.005$ (uncorrected), spatial extent > 100 mm^3.

We defined anatomical ROIs in individual subjects in the medial prefrontal cortex, posterior cingulate cortex and ventral striatum. Voxels were selected within each ROI according to their functional responsiveness in single-subject analyses. For the subjective value ROIs, all voxels were selected that exhibited a subjective value effect at time points 4–6 of $P < 0.05$ (uncorrected). For the unbiased value ROIs, all voxels were selected that exhibited greater activity (time points 4–6, $P < 0.05$, uncorrected) for either the largest amount compared to the smallest amount of the delayed reward across all delays, or the shortest delay compared to the longest delay to the delayed reward across all amounts. ROIs where no voxels met these functional criteria ($n = 3/60$) were omitted from further analyses.

Activity was averaged across all selected voxels within an ROI and summed over time points 4–6 for each trial. The resulting data were regressed separately against the objective amount of the larger delayed reward or the imposed delay to the larger delayed reward, and also fit with a nonlinear version of the subjective value regression to estimate the neural discount rate. For the regression analyses in **Figure 6**, both the ratio of the delay and amount regression coefficients and the neural discount rates were log transformed because of their skewed distribution. Because the log transformation only permits positive values, any negative values were replaced with the smallest observed positive value. All negative values were close to zero, so constraining values to be positive in this manner only reduced the range of ratios by less than 2.5% and the range of neural discount rates by less than 1%. ROIs for which the amount regression coefficient was negative ($n = 2/57$) or where no discount rate accounted for a significant amount of the variance in neural activity ($n = 5/57$) were omitted from further analyses. Fits to the ROI data and the regression analyses in **Figure 6** both used robust regression to minimize the effect of outliers. Robust regression was performed in MATLAB, which uses an iteratively re-weighted least-squares algorithm with the weights at each iteration determined by applying a bi-square function to the residuals from the previous iteration. Greater detail regarding the statistical analyses of the fMRI data are provided in the **Supplementary Methods**.

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We thank A. Bisin, A. Caplin, M. Grantner, D. Heeger and I. Levy for their comments on earlier versions of this paper. This work was supported by grants from the National Institutes of Health (F32-MH75544 to J.W.K. and R01-NS054775 to P.W.G.), the McDonnell Foundation (to P.W.G.) and the Seaver Foundation.

Published online at http://www.nature.com/natureneuroscience
Reprints and permissions information is available online at http://npg.nature.com/reprintsandpermissions

1. Weber, E.H. *De Pulsu, Resorptione, Audita et Tactu. Annotationes Anatomicae et Physiologicae.* (C.F. Koehler, Lipsiae, 1834).
2. Fechner, G.T. *Elemente der Psychophysik* (Breitkopf & Hartel, Leipzig, 1860).
3. Stevens, S. Neural events and the psychophysical law. *Science* **170**, 1043–1050 (1970).
4. Bernoulli, D. *The Works* (Birkhauser, Boston, 1982).
5. Platt, M.L. & Glimcher, P.W. Neural correlates of decision variables in parietal cortex. *Nature* **400**, 233–238 (1999).
6. Abler, B., Walter, H., Erk, S., Kammerer, H. & Spitzer, M. Prediction error as a linear function of reward probability is coded in human nucleus accumbens. *Neuroimage* **31**, 790–795 (2006).
7. Cromwell, H.C. & Schultz, W. Effects of expectations for different reward magnitudes on neuronal activity in primate striatum. *J. Neurophysiol.* **89**, 2823–2838 (2003).
8. Delgado, M.R., Nystrom, L.E., Fissell, C., Noll, D.C. & Fiez, J.A. Tracking the hemodynamic responses to reward and punishment in the striatum. *J. Neurophysiol.* **84**, 3072–3077 (2000).
9. McCoy, A.N., Crowley, J.C., Haghighian, G., Dean, H.L. & Platt, M.L. Saccade reward signals in posterior cingulate cortex. *Neuron* **40**, 1031–1040 (2003).
10. O'Doherty, J., Critchley, H., Deichmann, R. & Dolan, R.J. Dissociating valence of outcome from behavioral control in human orbital and ventral prefrontal cortices. *J. Neurosci.* **23**, 7931–7939 (2003).
11. Breiter, H.C., Aharon, I., Kahneman, D., Dale, A. & Shizgal, P. Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron* **30**, 619–639 (2001).
12. Nieuwenhuis, S. *et al.* Activity in human reward-sensitive brain areas is strongly context dependent. *Neuroimage* **25**, 1302–1309 (2005).
13. Rogers, R.D. *et al.* Distinct portions of anterior cingulate cortex and medial prefrontal cortex are activated by reward processing in separable phases of decision-making cognition. *Biol. Psychiatry* **55**, 594–602 (2004).
14. Kuhnen, C.M. & Knutson, B. The neural basis of financial risk taking. *Neuron* **47**, 763–770 (2005).
15. Knutson, B., Taylor, J., Kaufman, M., Peterson, R. & Glover, G. Distributed neural representation of expected value. *J. Neurosci.* **25**, 4806–4812 (2005).
16. Bowman, E.M., Aigner, T.G. & Richmond, B.J. Neural signals in the monkey ventral striatum related to motivation for juice and cocaine rewards. *J. Neurophysiol.* **75**, 1061–1073 (1996).
17. Cromwell, H.C., Hassani, O.K. & Schultz, W. Relative reward processing in primate striatum. *Exp. Brain Res.* **162**, 520–525 (2005).
18. McCoy, A.N. & Platt, M.L. Risk-sensitive neurons in macaque posterior cingulate cortex. *Nat. Neurosci.* **8**, 1220–1227 (2005).
19. Dorris, M.C. & Glimcher, P.W. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* **44**, 365–378 (2004).
20. McClure, S.M. *et al.* Neural correlates of behavioral preference for culturally familiar drinks. *Neuron* **44**, 379–387 (2004).
21. Gottfried, J.A., O'Doherty, J. & Dolan, R.J. Encoding predictive reward value in human amygdala and orbitofrontal cortex. *Science* **301**, 1104–1107 (2003).
22. Tanaka, S.C. *et al.* Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nat. Neurosci.* **7**, 887–893 (2004).
23. Padoa-Schioppa, C. & Assad, J.A. Neurons in the orbitofrontal cortex encode economic value. *Nature* **441**, 223–226 (2006).
24. Daw, N.D., O'Doherty, J.P., Dayan, P., Seymour, B. & Dolan, R.J. Cortical substrates for exploratory decisions in humans. *Nature* **441**, 876–879 (2006).
25. O'Doherty, J.P., Buchanan, T.W., Seymour, B. & Dolan, R.J. Predictive neural coding of reward preference involves dissociable responses in human ventral midbrain and ventral striatum. *Neuron* **49**, 157–166 (2006).
26. Johnson, K.O., Hsiao, S.S. & Yoshioka, T. Neural coding and the basic law of psychophysics. *Neuroscientist* **8**, 111–121 (2002).
27. Parker, A.J. & Newsome, W.T. Sense and the single neuron: probing the physiology of perception. *Annu. Rev. Neurosci.* **21**, 227–277 (1998).
28. Werner, G. & Mountcastle, V.B. Quantitative relations between mechanical stimuli to the skin and neural responses evoked by them. in *The Skin Senses* (ed. Kenshalo, D.R.) 112–138 (Charles C Thomas Publisher, Springfield, Illinois, 1968).
29. Newsome, W.T., Britten, K.H. & Movshon, J.A. Neuronal correlates of a perceptual decision. *Nature* **341**, 52–54 (1989).
30. Frederick, S., Loewenstein, G. & O'Donoghue, T. Time discounting and time preference: a critical review. *J. Econ. Lit.* **40**, 351–401 (2002).
31. Mazur, J.E. An adjusting procedure for studying delayed reinforcement. in *Quantitative Analysis of Behavior: The Effects of Delay and Intervening Events on Reinforcement Value* (eds. Commons, M.L., Mazur, J.E., Nevin, J.A. & Rachlin, H.) 55–73 (Lawrence Erlbaum Associates, Publishers, Hillsdale, New Jersey, 1987).
32. Laibson, D. Golden eggs and hyperbolic discounting. *Q. J. Econ.* **12**, 443–477 (1997).
33. Green, L. & Myerson, J. A discounting framework for choice with delayed and probabilistic rewards. *Psychol. Bull.* **130**, 769–792 (2004).
34. Benhabib, J., Bisin, A. & Schotter, A. Present-bias, quasi-hyperbolic discounting and fixed costs. *Working Paper* (New York University Department of Economics, New York, 2006) <http://www.econ.nyu.edu/user/benhabib/pshype0106.pdf>.
35. McClure, S.M., Laibson, D.I., Loewenstein, G. & Cohen, J.D. Separate neural systems value immediate and delayed monetary rewards. *Science* **306**, 503–507 (2004).
36. Teller, D.Y. & Pugh, E.N. Linking propositions in color vision. in *Colour Vision: Physiology and Psychophysics* (eds. Mollon, J.D. & Sharpe, L.T.) 577–589 (Academic Press, New York, 1983).
37. von Neumann, J. & Morgenstern, O. *The Theory of Games and Economic Behavior* (Princeton University Press, Princeton, New Jersey, 1944).
38. Samuelson, P.A. *Foundations of Economic Analysis* (Harvard University Press, Cambridge, Massachusetts, 1947).
39. Gul, F. & Pesendorfer, W. The case for mindless economics. *Working Paper* (Princeton University Department of Economics, Princeton, New Jersey, 2005) <http://www.princeton.edu/~pesendor/mindless.pdf>.
40. Gusnard, D.A. & Raichle, M.E. Searching for a baseline: functional imaging and the resting human brain. *Nat. Rev. Neurosci.* **2**, 685–694 (2001).
41. Zysset, S., Huber, O., Ferstl, E. & von Cramon, D.Y. The anterior frontomedian cortex and evaluative judgment: an fMRI study. *Neuroimage* **15**, 983–991 (2002).
42. Hsu, M., Bhatt, M., Adolphs, R., Tranel, D. & Camerer, C.F. Neural systems responding to degrees of uncertainty in human decision-making. *Science* **310**, 1680–1683 (2005).
43. Paulus, M.P. & Frank, L.R. Anterior cingulate activity modulates nonlinear decision weight function of uncertain prospects. *Neuroimage* **30**, 668–677 (2006).
44. Tom, S.M., Fox, C.R., Trepel, C. & Poldrack, R.A. The neural basis of loss aversion in decision-making under risk. *Science* **315**, 515–518 (2007).
45. Sugrue, L.P., Corrado, G.S. & Newsome, W.T. Matching behavior and the representation of value in the parietal cortex. *Science* **304**, 1782–1787 (2004).
46. Shadlen, M.N. & Newsome, W.T. Motion perception: seeing and deciding. *Proc. Natl. Acad. Sci. USA* **93**, 628–633 (1996).
47. Collier, M. & Williams, M.B. Eliciting individual discount rates. *Exp. Econ.* **2**, 107–127 (1999).
48. Camerer, C.F. & Hogarth, R.M. The effects of financial incentives in experiments: a review and capital-labor-production framework. *J. Risk Uncertain.* **19**, 7–42 (1999).
49. Smith, V.L. & Walker, J.M. Monetary rewards and decision cost in experimental economics. *Econ. Inq.* **31**, 245–261 (1993).