

An “As Soon As Possible” Effect in Human Intertemporal Decision Making: Behavioral Evidence and Neural Mechanisms

Joseph W. Kable and Paul W. Glimcher

Center for Neural Science, New York University, New York, New York

Submitted 27 February 2009; accepted in final form 23 February 2010

Kable JW, Glimcher PW. An “as soon as possible” effect in human intertemporal decision making: behavioral evidence and neural mechanisms. *J Neurophysiol* 103: 2513–2531, 2010. First published February 24, 2010; doi:10.1152/jn.00177.2009. Many decisions involve a trade-off between the quality of an outcome and the time at which that outcome is received. In psychology and behavioral economics, the most widely studied models hypothesize that the values of future gains decline as a roughly hyperbolic function of delay from the present. Recently, it has been proposed that this hyperbolic-like decline in value arises from the interaction of two separate neural systems: one specialized to value immediate rewards and the other specialized to value delayed rewards. Here we report behavioral and functional magnetic resonance imaging results that are inconsistent with both the standard behavioral models of discounting and the hypothesis that separate neural systems value immediate and delayed rewards. Behaviorally, we find that human subjects do not necessarily make the impulsive preference reversals predicted by hyperbolic-like discounting. We also find that blood oxygenation level dependent activity in ventral striatum, medial prefrontal, and posterior cingulate cortex does not track whether an immediate reward was present, as proposed by the separate neural systems hypothesis. Activity in these regions was correlated with the subjective value of both immediate and delayed rewards. Rather than encoding only the relative value of one reward compared with another, these values are represented on a more absolute scale. These data support an alternative behavioral–neural model (which we call “ASAP”), in which subjective value declines hyperbolically relative to the soonest currently available reward and a small number of valuation areas serve as a final common pathway through which these subjective values guide choice.

INTRODUCTION

Many decisions involve a trade-off between the quality of an outcome and the time at which that outcome is received. This is because the subjective value of a potential reward falls as the delay to that reward lengthens, a phenomenon called *temporal discounting* (Bickel and Marsch 2001; Frederick et al. 2002; Green and Myerson 2004). Since the 1930s, the standard economic model has assumed that the subjective value of future rewards must, in principle, decline exponentially (at a constant rate) as a function of delay from the present (Samuelson 1937). Exponential discounters remain consistent as they move through time, making the same choices about future gains regardless of when they make those choices (Strotz 1956). They do not make impulsive preference reversals—initially committing to the larger-and-later of two rewards when both options are in the future, but then switching and taking the smaller-but-sooner reward when that option be-

comes immediately available. Nor do they impulsively overweight the desirability of immediately available rewards, always consuming a greater portion of their resources now than later.

Despite its normative appeal, there is overwhelming evidence that people and other animals are not exponential discounters (Bickel and Marsch 2001; Frederick et al. 2002; Green and Myerson 2004). In particular, their choices exhibit decreasing impatience with delay: subjective value declines at a rate that decreases as the delay between two options grows larger, a kind of “overweighting” of immediately available rewards. Starting from this well-described phenomenon, models in psychology and behavioral economics assume that subjective value declines hyperbolically (Mazur 1987) or quasi-hyperbolically (Laibson 1997) from the present. Since these two models have similar behavioral predictions but very different mathematical implications, we use the term “(quasi-)hyperbolic” discounting when making statements that apply to both models. Importantly, (quasi-)hyperbolic models predict both that choosers will overweight the desirability of immediately available rewards and that they will make impulsive preference reversals.

Hyperbolic and quasi-hyperbolic discounting are descriptive—or in the language of economics *positive*—models of behavior, rather than hypotheses about psychological or neural mechanisms. However, perhaps the best known quasi-hyperbolic model, the β - δ model (Laibson 1997), has inspired the mechanistic hypothesis that discounting behavior arises from the interaction of two competing systems: 1) an impatient emotional system (“ β ”) that exclusively or primarily values immediate rewards and 2) a patient rational system (“ δ ”) that values both immediate and delayed rewards. Previous imaging studies have found some support for this hypothesis, associating each of the two systems with specific brain regions (McClure et al. 2004a, 2007). In these studies the β system was associated with ventral striatum, ventromedial prefrontal, and posterior cingulate cortex, since these regions exhibited greater activity when an immediate reward could be chosen, compared with when all potential rewards were delayed. The δ system was associated with dorsolateral prefrontal and posterior parietal cortex, since these regions exhibited increased activity for all choices, with greater activity when the choice was more difficult or when a delayed reward was selected. To distinguish this behavioral–neural model from the behavioral quasi-hyperbolic model, we refer to this specific behavioral–neural model as the “separate neural systems” hypothesis.

We recently reported evidence inconsistent with the separate neural systems hypothesis (Kable and Glimcher 2007). In our study, participants chose between a fixed immediate gain and a second gain that varied in amount and delay from trial to trial.

Address for reprint requests and other correspondence: J. W. Kable, Department of Psychology, University of Pennsylvania, 3720 Walnut St., Philadelphia, PA 19104 (E-mail: kable@psych.upenn.edu).

We found that blood oxygenation level dependence (BOLD) activity in ventral striatum, medial prefrontal cortex, and posterior cingulate cortex (all part of the proposed β system) was correlated with the subject-specific subjective value of the delayed reward, regardless of whether that reward was to be obtained in hours or in months. We hypothesized that rather than representing an impulsive signal that exclusively or primarily values immediate rewards, these regions encode the subjective value of rewards at any delay. In addition to our study, other recent studies have also found that these brain regions are sensitive to changes in the value of delayed rewards, in a manner that is not predicted by the separate neural systems hypothesis (Ballard and Knutson 2009; Pine et al. 2009; Tanaka et al. 2004).

An often-noted drawback of our previous study, however, was that participants always chose between a future gain and a fixed immediate reward. Thus we could not test one of the key behavioral implications of (quasi-)hyperbolic discounting: that people will make impulsive preference reversals. Nor could we test a key neural claim regarding the hypothesized β system: that mean neural activity is greater when an immediate reward is available compared with when only delayed rewards can be chosen. On a more global level, we were also not able to test whether the value signals we observed were related to the *relative* subjective value of the later gain compared with the sooner reward or whether these signals encoded subjective value on a more absolute scale.

Here we report two experiments that allow us to test the key behavioral implications of (quasi-)hyperbolic models and the key neural predictions of the separate neural systems hypothesis. Our behavioral findings are inconsistent with (quasi-)hyperbolic models and our neural findings are inconsistent with the separate neural systems hypothesis (see also Glimcher et al. 2007). Behaviorally, our subjects did exhibit decreasing impatience as the later reward was delayed in time [the first of the two choice anomalies¹ predicted by (quasi-)hyperbolic models]. However, this was a function of the difference in delay between two rewards. Participants overweighted the value of the soonest available reward regardless of whether it was offered immediately or after a delay. The degree of their decreasing impatience was thus anchored to the soonest possible reward, not simply to the present. This meant that participants in our study did not make impulsive preference reversals in a consistent or statistically significant fashion [the second choice anomaly predicted by (quasi-)hyperbolic models].

Neurally, we did not find robust increases in mean activity in the proposed β areas when an immediate reward could be chosen, compared with when only delayed rewards were available, nor did we find evidence consistent with the existence of a δ system. We did, however, replicate our previous finding that BOLD activity in ventral striatum, medial prefrontal, and posterior cingulate cortex encodes the subjective value of immediate and delayed rewards. Further, we found that these regions encode subjective value on an absolute scale that depends on the delay from the present, rather than only encoding the relative subjective value of the later reward compared with the sooner reference, as might have been inferred from behavior. Our results question the generality of both standard

(quasi-)hyperbolic behavioral models (although not necessarily the generality of hyperbolic discounting) and the proposal that separate neural systems value immediate and delayed rewards. We propose an alternative behavioral–neural model that can account for the behavioral and neural data we have gathered. In this “as soon as possible” (ASAP) model (Glimcher et al. 2007), discounting arises from a unitary value signal that declines with the delay from the present and is also modulated by what other options are currently available. This model can account for our behavioral findings that impatience declines as a function of delay from the soonest possible reward and that impulsive preference reversals do not occur, and our neural findings that activity accurately tracks the overall delay to all rewards.

A conference proceeding containing a preliminary description of these conclusions was previously published (Glimcher et al. 2007).

METHODS

Participants

Twenty-five paid volunteers participated in the first experiment (13 women, 12 men, all right-handed; mean age = 22.2 yr). One of these volunteers was unable to return for the scanning session. Data from two participants' scanning sessions were discarded because these individuals' intertemporal choice behavior was not stable across sessions and this led to an underconstrained estimate of the discount function in the scanning session. Thus data from 25 individuals were included in the behavioral analyses, whereas data from 22 individuals were included in the neural analyses. Four of these 22 had previously participated in a similar experiment in our laboratory (Kable and Glimcher 2007). Ten participants (3 women, 7 men; mean age = 22.4 yr) with the highest discount rates from the first experiment were selected for participation in the second experiment. All 10 participants were included in both the behavioral and neural analyses. Two of these 10 had also participated in the previously published study (Kable and Glimcher 2007). We selected only these 10 individuals to participate because high discount rates provide more statistical power in behavioral and neural tests of hyperbolic discounting, since it is only when discount functions are steep that the distinction between exponential and (quasi-)hyperbolic is detectable. All participants gave informed consent in accordance with the procedures of the University Committee on Activities Involving Human Subjects of New York University.

Experimental design

A trial began with a red dot appearing in the center of the screen for 2 s (Fig. 1). Then, the two choice options were presented, one on each side of the red dot. After 4 s, the two options disappeared. After an additional 4 s, the red dot turned green, which cued participants to indicate their choice. Participants indicated their choice by pressing a button with the hand corresponding to the side of the screen on which their choice appeared.

Each trial involved a choice between a smaller amount of money paid at a sooner time and a larger amount paid at a later time. In the first experiment, there were two conditions. In the NOW condition, the choice was between \$20 available immediately and a larger amount available at a later date. The larger-later option was constructed using one of five delays (1 day to 120 days) and one of five amounts (\$21–116). Thus there were 25 unique choices in this condition. In the 60 DAY condition, the choices were constructed by adding a fixed delay of 60 days to *both* options in the NOW condition. Thus choices in this condition were between \$20 in 60 days and

¹ Here we use “anomaly” in the normative sense.




FIG. 1. The sequence of events within a trial is shown. First, a red dot appeared in the center of the screen for 2 s, then the 2 options were presented on either side of the screen for 4 s. The options disappeared and, after another 4 s, the red dot in the center turned green to cue participants to indicate their choice. Participants pushed a button with their left (right) hand to choose the left (right) option and received feedback indicating the option chosen. Each trial was a fixed length of 22 s. The left/right location of the sooner option was randomized across trials and the trials from the 2 experimental conditions were completely intermixed.

\$21–\$116 in 61–180 days. Each choice was presented four times over the course of a session, for a total of 100 choices per condition and 200 choices per session. The side of the screen on which each option was presented was counterbalanced across the four repetitions. Delays were the same for each participant. Amounts were the same for each participant in the initial two behavioral sessions, and chosen individually for each participant in the scanning session based on their previous behavior so that (as much as possible) an approximately equal number of sooner and later options were chosen. The five amounts chosen were the five predicted indifference points for the delays used, based on the participant's previous sessions and rounded to the nearest whole dollar, or those predicted for a discount rate of $k^{ASAP} = 0.005$ if the participant discounted at a lower rate than this (see Eq. 3 in the following text). This important manipulation ensured that the sooner and later options were of roughly equal average value for each subject. Participants were not told how options would change across sessions. Critically, choices from the two conditions (NOW and 60 DAY) were randomly intermixed and participants were not explicitly informed of the structure of the sets of choices offered to them.

In the second experiment, the NOW condition again involved choices between \$20 available immediately and \$21–\$86 available in 1–120 days. The SCALED 60 DAY condition, however, involved choices between a fixed amount greater than \$20 (ranging from \$32 to \$56 across participants) in 60 days and an even larger amount (up to \$224) in 61–180 days. Amounts for the SCALED 60 DAY condition were chosen as follows: First, for each participant, the fixed amount in 60 days was chosen based on that individual's previously acquired behavioral data as the smallest amount that the individual reliably preferred in 60 days over \$20 immediately. In this way, we identified a reward at 60 days that was roughly equivalent in subjective value to a reward of \$20 now. Then, the other five amounts paired with the larger-later rewards were scaled in the same manner; if subjects were roughly indifferent between \$20 immediately and \$35 at a delay of 60 days, all later rewards were increased by 75%. So, for example, the NOW condition might involve a choice between \$20 now and {\$21, \$22, \$26, \$32, \$44} in 1–120 days, whereas the SCALED 60 DAY

condition involved a choice between \$35 and {\$36, \$37, \$46, \$56, \$77} in 61–180 days.

In the first experiment, subjects participated in two 1 h behavioral sessions and one 2 h scanning session. Sessions were completed over 6–197 days (median = 34 days), with consecutive sessions separated by 2–191 days (median = 14 days). The 10 participants in the second experiment completed an additional behavioral session of the first experiment, 8–427 days after the first scanning session (median = 42 days). They then completed the second experiment, which involved one 2 h scanning session, after an additional 2–62 days (median = 8 days). Since people's decisions involving real gains may differ from those involving hypothetical gains (Camerer and Hogarth 1999; Collier and Williams 1999; Smith and Walker 1993), participants were paid according to four randomly selected trials per session (except for the first behavioral session, which subjects were informed involved only hypothetical choices), using commercially available prepaid debit cards (see Kable and Glimcher 2007 for details). The four subjects who had previously participated in a similar experiment in our laboratory skipped the first session of the first experiment, which involved purely hypothetical choices.

Behavioral analysis

In each condition, we estimated a discount function that described how the subjective value of the larger-later reward declined as the delay from the fixed sooner reward increased. We estimated these discount functions individually for each participant. First, for the five amounts measured at each delay interval (in each condition), we estimated an indifference point from that individual's choices at that delay. This indifference point was estimated by fitting a logistic curve to the proportion of choices of the larger-later reward, plotted as a function of the larger-later amount. The indifference point was thus the amount, at that delay interval, for which the individual would be predicted to choose the smaller-sooner and larger-later rewards with equal frequency. To obtain a discount curve, these five delay/amount indifference points were then divided into \$20 (or the amount of the fixed sooner reward in the SCALED 60 DAY data set) and a single-parameter hyperbolic (Eq. 3) or exponential (Eq. 4) function was fit to the resulting curve.

This follows standard practice in the psychological literature on discounting (Green and Myerson 2004; Mazur 1987). We should note that interpreting these curves as pure discount functions relies on the (almost certainly false) assumption that the subjective value, or utility, of money (independent of delay) increases linearly with objective amount. Although this assumption is a standard one in previous experimental studies of discounting, we do not expect that it would hold generally outside a limited range of monetary amounts and is likely violated to some degree in our data set. Violations of this assumption, however, would not alter the fact that these curves reflect the subjective desirability of delayed gains, although such violations would mean that these functions are better described as discounted utility functions because they incorporate effects of amount as well as delay.²

Imaging

Imaging data were collected with a Siemens Allegra 3T head-only scanner equipped with a head coil from Nova Medical (Wakefield, MA). T2*-weighted functional images were collected using an echoplanar imaging sequence (repetition time [TR] = 2 s, time to echo [TE] = 30 ms, 35 axial slices acquired in ascending interleaved order, $3 \times 3 \times 3$ mm, 64×64 matrix in a 192 mm field of view [FOV]). Each scan consisted of 275 images. The first two images were discarded to avoid T1 saturation effects. There were 25 choice trials during each scan; each trial lasted 14 s with an 8 s intertrial interval (Fig. 1). Each experiment involved one session of eight scans. High-resolution, T1-weighted anatomical images were also collected using a magnetization-prepared rapid gradient-echo sequence (TR = 2.5 s, TE = 3.93 ms, inversion time [TI] = 900 ms, flip angle = 8°, 144 sagittal slices, $1 \times 1 \times 1$ mm, 256×256 matrix in a 256 mm FOV).

fMRI analysis

Functional imaging data were analyzed using BrainVoyager QX (Brain Innovation, Maastricht, The Netherlands). 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 detrended and high-pass filtered (cutoff of 3 cycles/scan, or 0.0064 Hz) to remove low frequency drift in the functional magnetic resonance imaging (fMRI) signal. Images were then coregistered with each person's high-resolution anatomical scan, rotated into the anterior commissure–posterior commissure plane, and normalized into Talairach space using piecewise affine Talairach grid scaling. All spatial transformations of the functional data used trilinear interpolation. Before group-level random-effects analyses, data were

² Thus though we refer to these as discount functions, the reader is enjoined to remember that they inherit the axiomatic features of the functionally combined discounted utility form. It is critical to note that two key tests presented in the following text are unaffected by this combination. The analyses of the neural data described later test for effects of subjective value, which combines both the utility and discount functions by design, rather than separating out specific effects of each. The direct test for preference reversals reported in Fig. 3, C and E does not depend on any logical aspect of this combination. Of course, violations of our assumption that the utility function could be described as linear would mean that the subsequent sections misestimate the precise shape of the “true” discount function. However, for these violations to have caused exponential functions to appear hyperbolic (which we believe to be the most theoretically significant distortion that might be imposed by this assumption), the unmodeled utility functions would have to be strongly concave/accelerating, thus requiring that our subjects be risk-seeking. This seems unlikely because almost all experimental studies have found that human utility functions are convex/decelerating (e.g., Holt and Laury 2002). In other words, our assumption that utility is linear in money likely means that we have underestimated the advantage of hyperbolic over exponential discount functions in the sections that follow.

also spatially smoothed with a Gaussian kernel of 8 mm (full-width half-maximum).

Statistical analyses were performed using a general linear model, estimated with ordinary least squares. To address the temporal autocorrelation in the noise, data were “prewhitened” by estimating a first-order autoregressive [AR(1)] model on the residuals. We used a variant of a “finite-impulse response” (FIR) or “deconvolution” model. In our main analysis, the two conditions in each experiment (NOW and 60 DAY or SCALED 60 DAY) were modeled separately. For each condition, the model included ten covariates that fit the mean activity, across all trials, for each of the first ten time points in a trial. Each of these ten covariates was a one-TR impulse of unit height, occurring at a specific time point in each trial. Ten additional covariates fit the deviations from the mean at each time point that were correlated with the relative subjective value ($SV^{ASAP_{g=1}}$ in Eq. 3 in the following text) of the later reward across trials. Each of these ten covariates was a one-TR impulse, occurring at a specific time point in each trial, with a height corresponding to the mean-centered relative subjective value of the later reward on that trial, as determined individually for each subject using Eq. 3. We chose such a FIR model because it could flexibly capture subjective value effects at any point in the trial, without making assumptions about their precise timing a priori (i.e., for the duration of option presentation, during the delay, preceding the motor response, etc.). This model also did not require any assumptions about the shape of the hemodynamic response function.

Figure 4 shows three effects that were of primary interest: the effect of relative subjective value ($SV^{ASAP_{g=1}}$) at time points 4–6 in the NOW condition, the same effect in the 60 DAY condition, and the difference in mean activity at time points 4–6 between the two conditions. We focus on this particular time window based on the results of our previous study; there were no additional significant effects at other points in the trial. Figure 8 reports the effects of mean activity at time points 4–6 for both conditions (NOW + 60 DAY), as well as the results of two additional analyses, which replaced the relative subjective value regressor in the two conditions with either 1) the participant's reaction time on that trial or 2) a binary variable indicating the participant's choice (later or sooner reward). The additional analyses in Figs. 6 and 7 examined both conditions (NOW and 60 DAY) together. These analyses use the same FIR structure (with 20 covariates modeling both conditions combined rather than 40 modeling the two separately), with a subjective value covariate (SV^{ASAP}) constructed assuming different combination rules (see Table 2 in the following text). In each case, this covariate assumes the ASAP discounting function (Eq. 2) with the gain factor described in Eq. 5, although the covariates differ in how the subjective values of the two options are combined (Table 2). Two additional analyses are included for comparison: one that tests for correlations with the amount of the delayed reward alone and one that tests for correlations with relative subjective value alone (i.e., $SV^{ASAP_{g=1}}$, where the gain factor in Eq. 2 is fixed at one). To be able to compare beta coefficients across analyses in Figs. 6 and 7, as well as across experiments in Fig. 9, the values used to construct each covariate were mean-centered and divided by the SD. The beta coefficients on such normalized covariates have been used previously for model comparison in fMRI (Hampton et al. 2008; Hare et al. 2008).

As discussed in the following text, some of our analyses speak to an important issue, which is distinguishing neural signals that track subjective value from those that track only a component of subjective value, such as the reward magnitude or the discount factor associated with the delay. If a region encoded only magnitude, then it should show a similar correlation with relative subjective value in the NOW and 60 DAY conditions (Fig. 4), since the magnitudes are exactly the same in these conditions. Also, activity in such a region should be best fit by the magnitude alone model in the comparison reported in Fig. 6. Another, perhaps stronger, test would be to include magnitude, discount factor, and subjective value in the same statistical model

(Ballard and Knutson 2009; Pine et al. 2009). We also estimated several such combined models: 1) a model including all three variables (magnitude, discount factor, and subjective value) with no orthogonalization; 2) a model in which discount factor was orthogonalized with respect to magnitude and subjective value was orthogonalized with respect to magnitude and discount factor, so that all shared variance was attributed to the magnitude covariate; and 3) a model in which discount factor was orthogonalized with respect to subjective value and magnitude was orthogonalized with respect to subjective value and discount factor, so that all shared variance was attributed to the subjective value covariate. Since the results of these analyses were not conclusive, we do not discuss them in detail here, although the trends observed were consistent with the conclusions drawn here. We suspect that some of the difference between our results and those of other investigators may be due to a higher correlation between the three variables in our data set, driven by the presence of several very patient individuals (if a person does not discount at all, the reward magnitude is the subjective value).

Group random-effects analyses were performed using the summary statistics approach, which tests whether the mean effect at each voxel is significantly different from zero across subjects. We report all effects that were significant at $P < 0.05$ (corrected) using cluster-level inference. Contrast maps were initially thresholded at $P < 0.005$ (uncorrected) and the appropriate spatial extent threshold for corrected cluster-level inference at $P < 0.05$ was determined for each contrast. Spatial extent thresholds were determined using Monte Carlo simulations on noisy data sets that were artificially generated with the same spatial smoothness as the contrast maps. The cluster threshold estimator implemented in BrainVoyager is thus conceptually similar to the AlphaSim program in Analysis of Functional NeuroImages (AFNI). To investigate whether the whole brain analysis might have missed weaker effects, we also defined regions of interest (ROIs) in ventral striatum, medial prefrontal, and posterior cingulate cortex. These ROIs encompassed all of the voxels in the significant cluster in that region for the relative subjective value effect in the NOW condition (ventral striatum ROI = 1,115 mm³, medial prefrontal ROI = 1,065 mm³, posterior cingulate ROI = 8,695 mm³) and we took the average signal across all of these voxels as the time series for that ROI. We used the relative subjective value effect in the NOW condition because this was the only contrast that generated significant effects in these regions and because it was conceptually identical to the effects reported in our previous study. Also, defining the ROIs in this way does not bias the two statistical tests we perform using the same data set—whether there is a significant correlation with relative subjective value in the 60 DAY condition and whether there is a significant difference in mean activity between the NOW and 60 DAY conditions. We also use the same ROIs in analyzing the independent data set from *experiment 2*.

RESULTS

Hyperbolic and “as soon as possible” discounting

(Quasi-)hyperbolic discounting models assume that subjective value declines roughly hyperbolically relative to the *present*. Specifically, the standard hyperbolic function for the subjective value is

$$SV^H = \frac{A}{1 + k^H D} \quad (1)$$

where SV^H is subjective value estimated from this model, A is the objective amount of the reward, D is the delay to the reward, and k^H is a subject-specific constant. However, as several others have noted (Green et al. 2005; Read 2001; Read and Roelofsma 2003; Scholten and Read 2006), most previous

studies have used choices where one reward is always immediately available (similar to our NOW condition). Thus at least in principle, one alternative to (quasi-)hyperbolic models is one in which subjective value declines hyperbolically *relative to the soonest possible reward*, rather than with regard to the present. We call this possibility “as soon as possible” (ASAP) discounting, to distinguish it from (quasi-)hyperbolic discounting (Glimcher et al. 2007). The most general form of an ASAP function for subjective value is

$$SV^{ASAP} = g(D_{ASAP}) \frac{A}{1 + k^{ASAP}(D - D_{ASAP})} \quad (2)$$

where SV^{ASAP} is the subjective value estimated from this model, D_{ASAP} is the delay to the soonest possible reward, and $g(D_{ASAP})$ is a gain factor that is a function of the delay to the soonest possible reward.

This ASAP discounting function is closely related to several previous proposals. Ok and Masatlioglu (2005) described a class of “relative discounting” functions, in which the difference in delay between two options, rather than the delay from the present, is the key variable. The difference in delay is also a critical variable in Read’s “subadditive” (Read 2001; Read and Roelofsma 2003) and “discounting by intervals” (Scholten and Read 2006) models, in the “elimination-by-aspects” and “common-aspect attenuation” models discussed by Green and colleagues (2005) and in the similarity model reported by Rubinstein (2003). The assumption of a hyperbolic form makes ASAP most similar to “elimination-by-aspects.” However, ASAP differs from these previous models in two critical ways. ASAP is less broad in its psychological claims. It is not committed to the discount function arising from a particular psychological process, such as subadditivity or the serial comparison of dimensions. ASAP is also unique in that it makes explicit neural claims and is thus falsifiable on both behavioral and neural grounds. These other related models describe only behavior. As discussed in the following text, this aspect of ASAP means that *Eq. 2* includes novel testable components that would be irrelevant in a behavior-only model.

Figure 2 provides a graphical comparison between the hyperbolic and ASAP models. There are three important things to note about this comparison. First, both models make the same predictions about the NOW condition, where the soonest reward is available immediately ($D_{ASAP} = 0$). Second, for the ASAP model, the subjective value of a larger-later reward in the 60 DAY condition will not necessarily match the subjective value of the same exact reward in the NOW condition.³ This is because in the ASAP model subjective value depends on the other options that are available, particularly on the delay to the soonest reward in the choice set. A subjective value in the ASAP model is consistent only across choices within a single “condition” in which the soonest possible reward is always offered at the same delay. This choice-set dependence has broader implications that we mention in the DISCUSSION and can explain the deviations between the ASAP and hyperbolic predictions for behavior and

³ It should be noted, however, that this difference could not be observed directly in behavioral choices (i.e., it would not be possible, or at least be very strange, to give someone a choice between 1) \$21 in 61 days when the soonest reward is available immediately and 2) \$21 in 61 days when the soonest reward is available in 60 days). It could be indirectly inferred from a pattern of choices or directly detected through neural measurements of subjective values if the ASAP model is taken to be mechanistic.

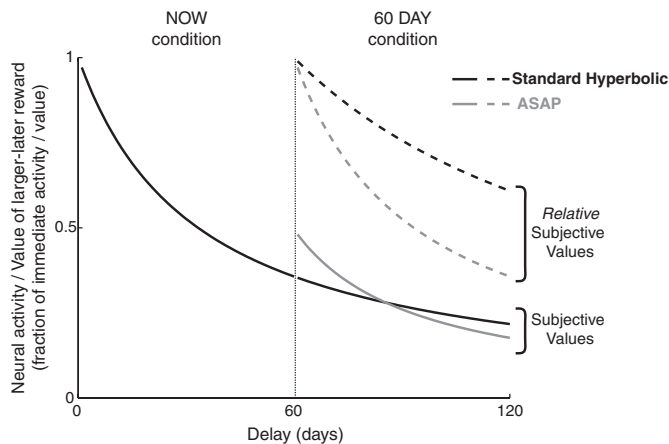


FIG. 2. Schematic illustrating standard hyperbolic discounting vs. “as soon as possible” (ASAP) discounting. Plotted is how the subjective value and relative subjective value of the larger-later reward change in the 60 DAY condition. The 2 models make the same predictions in the NOW condition. Standard hyperbolic discounting predicts decreased impatience in the 60 DAY condition, as evidenced by the shallower decline in relative subjective value, whereas ASAP discounting predicts the same degree of impatience, as evidenced by the same steepness in the decline of relative subjective value. Standard hyperbolic discounting also assumes that subjective values are consistent between the 2 conditions, as evidenced by the continuity of the curves in the 2 conditions, whereas ASAP discounting assumes that subjective values depend on other options available and this choice-set dependence leads to the discontinuity in the 2 curves. All of the features illustrated here for hyperbolic discounting are shared by quasi-hyperbolic discounting as well. The subjective values shown assume $k^H = k^{ASAP} = 0.03$, $g(0) = 1$, and $g(60) = 0.495$.

neural activity in the 60 DAY condition. Third, if the *only* goal is to make behavioral predictions, then the function that serves as the gain factor, $g(D_{ASAP})$, is an unnecessary addition to the model. In our hands, ASAP is meant to be a model of both choices and the neurally encoded subjective values that drive choices. The gain function scales these neurally encoded subjective values within a “condition,” where a “condition” is a set of choices in which the time to the soonest possible option is held constant. The gain function in ASAP thus serves to condition the numerical subjective value of an option on the earliest available option in the choice set. This means that, at a behavioral level, the value of $g(D_{ASAP})$ (which is constant within a choice set with a fixed earliest reward) does not change what is preferred.⁴ There is therefore no way, in principle, to estimate the gain function behaviorally or to distinguish between different gain functions based on the behavioral data alone. Models with different gain functions make the same behavioral predictions.⁵

Although the inclusion of $g(D_{ASAP})$ may for this reason seem odd to purely behavioral scientists, it serves a powerful intuition about both mental state and brain function. Even though in ASAP someone’s preference between two options will

⁴ To see this mathematically, note that preferring option 1 to option 2 requires that $g(D_{ASAP})\{A_1/[1 + k^{ASAP}(D_1 - D_{ASAP})]\} > g(D_{ASAP})\{A_2/[1 + k^{ASAP}(D_2 - D_{ASAP})]\}$ and that the gain factor falls out of this inequality. Thus changing the value of $g(D_{ASAP})$, as long as it is constrained not to be negative, can never reverse the sign of this relation.

⁵ Strictly speaking, different gain functions could make different predictions about the degree of stochasticity in choices, depending on what the assumed choice model is (i.e., logistic vs. Luce’s choice ratio). Testing these predictions, however, would require that the size of these effects is large compared with the other factors influencing stochasticity. Here we stress that different gain factors would not change (mean) preference orderings under standard assumptions.

depend only on the difference in delay between the two, the brain (and the psychological processes it underlies) might still keep track of how much both options are delayed from the present. When choosing \$21 in 61 days over \$20 in 60 days, the internal subjective values of these options may be less than the values when choosing \$21 tomorrow over \$20 today. In other words, subjects may know that a choice between \$21 in 61 days and \$20 in 60 days is worse than a choice between \$21 tomorrow and \$20 today, even when given no opportunity to express that fact. In any case, we include the gain function in the model here because it *does affect the neural predictions*. Neural representations of subjective value would differ for different gain functions in this model.

Subjective value of delayed rewards declines hyperbolically relative to the soonest possible reward (experiment 1)

For the purposes of the behavioral analyses alone, we can make the simplifying assumption that the gain function $g(D_{ASAP})$ is always equal to one, effectively omitting it from the behavior-only predictions of the model

$$SV^{ASAP}_{g=1} = \frac{A}{1 + k^{ASAP}(D - D_{ASAP})} \quad (3)$$

Note that this simplified form of ASAP corresponds to the “elimination-by-aspects” model discussed by Green and colleagues (2005), as well as to one parameterization of the more general “discounting by intervals” model presented by Scholten and Read (2006). One important thing to note about this simplified parameterization of ASAP is that $SV^{ASAP}_{g=1}$ is now the subjective value of a reward, relative to the soonest possible reward. Thus the distinction between an ASAP model in which the gain factor [$g(D_{ASAP})$] varies with the delay to the soonest possible reward versus one in which the gain factor is a constant (i.e., $g = 1$, Eq. 3) corresponds to a distinction between encoding subjective values across choice sets⁶ versus encoding only *relative* subjective values within choice sets. For this reason, we refer in the following text to the subjective values calculated using Eq. 3 ($SV^{ASAP}_{g=1}$) as *relative subjective values*.

Another nice feature of the simplified parameterization of ASAP (Eq. 3) is that it highlights an important general feature of ASAP models. When these models are fit to the NOW and 60 DAY choices separately, the standard hyperbolic and ASAP models make clearly distinguishable predictions about the resulting estimates of the discount rate from the two conditions (k_0^{ASAP} , k_{60}^{ASAP}). It should be clear that ASAP predicts that the discount rates estimated from the two conditions should be the same

$$k_{60}^{ASAP} = k_0^{ASAP}$$

This prediction also reflects the fact that in ASAP the hyperbolic decline in subjective value occurs relative to the soonest possible delay, so people should exhibit the same degree of impatience in the NOW and 60 DAY conditions. An illustra-

⁶ On a more formal note it should be stated that encoding subjective values across choice sets, in the manner that ASAP models do, does *not* ensure consistency in the technical (economic) sense. Indeed, even the notion of a subjective value existing across choice sets is troublingly problematic from a revealed preference perspective.

tion of this prediction can also be seen in Fig. 2, which shows that the decline in relative subjective value for the ASAP model is identical in the NOW and 60 DAY conditions.

In contrast, in (quasi-)hyperbolic discounting, the rate of decline in subjective value decreases as rewards move farther out from the present. (Quasi-)hyperbolic discounting thus predicts that people should be more impatient in the NOW condition compared with the 60 DAY condition. This means that the discount rate estimated in this analysis for the NOW condition (k_0^{ASAP}) should be greater than the discount rate estimated for the 60 DAY condition (k_{60}^{ASAP}). Again, this prediction is illustrated in Fig. 2, which shows that the decline in relative subjective value for the hyperbolic model is less steep in the 60 DAY condition than that in the NOW condition. It can be shown algebraically using *Eqs. 1* and *3* that standard hyperbolic discounting predicts that the two discount rates estimated in this way should be precisely related by the following (see Green et al. 2005 or Supplemental Note for a derivation)⁷

$$k_{60}^{\text{ASAP}} = \frac{k_0^{\text{ASAP}}}{1 + 60k_0^{\text{ASAP}}}$$

where

$$k_0^{\text{ASAP}} = k^H \quad \text{and} \quad k_0^{\text{ASAP}} > k_{60}^{\text{ASAP}}$$

We separated out the choices in the NOW and 60 DAY conditions and fit *Eq. 3* separately for each condition. Consistent with ASAP, but not with (quasi-)hyperbolic discounting, our empirical data indicate that participants discounted at similar rates in both conditions (Fig. 3*B*; see also Glimcher et al. 2007). [These data are thus also inconsistent with models intermediate between ASAP and (quasi-)hyperbolic discounting, such as “common-aspect attenuation”; see Supplemental Note and Supplemental Fig. S1.] The discount rate we observed in the 60 DAY condition was close to the discount rate we observed in the NOW condition for each subject ($k_{60}^{\text{ASAP}} \approx k_0^{\text{ASAP}}$) and the discount rates in the 60 DAY condition were significantly larger than those predicted by standard hyperbolic discounting (Wilcoxon signed-rank test, $P < 1 \times 10^{-8}$). Perhaps surprisingly though, the discount rate we observed in the 60 DAY condition was actually slightly larger than the discount rate we observed in the NOW condition on a subject-by-subject basis ($k_{60}^{\text{ASAP}} > k_0^{\text{ASAP}}$; Wilcoxon signed-rank test, $P = 0.008$). Thus if anything, participants were *more impatient* in the 60 DAY condition than in the NOW condition, an effect that goes in the *opposite* direction from that predicted by (quasi-)hyperbolic models. The size of this effect, however, was very small compared with the discount rates in the two conditions (median difference between k_{60}^{ASAP} and $k_0^{\text{ASAP}} = 0.0002$, median $k_0^{\text{ASAP}} = 0.006$), suggesting this difference might best be considered trivial.

Participants do not make impulsive preference reversals (experiment 1)

In addition to comparing the discount rates estimated from the NOW and 60 DAY data, we also compared the choices in these two conditions directly, to test whether participants made

impulsive preference reversals as predicted by (quasi-)hyperbolic discounting. To do this, we considered pairs of choices across the two conditions that differed only in the addition of a common 60 day delay to both options. (Quasi-)hyperbolic discounting predicts that people will sometimes make impulsive preference reversals between two such choices: taking the larger-later reward from a pair in the 60 DAY condition, while also picking the smaller-sooner reward in the corresponding pair in the NOW condition. ASAP (and the behavioral models like it) differs from (quasi-)hyperbolic models in that it does not predict such reversals. Under ASAP participants should make the same choice in both conditions, on average. Under ASAP only stochasticity in choice could yield changes between the NOW and 60 DAY conditions and these changes should be randomly distributed, rather than systematically distributed toward more impulsive choices in the NOW condition. Across all pairs of choices in the two conditions, we found that participants made the same choice in both conditions >90% of the time (Fig. 3*C*). In addition, even when participants' choices did differ, the differences appeared randomly distributed. Subjects, if anything, were slightly more likely to be more impulsive (choosing smaller-sooner) in the 60 DAY condition, rather than in the NOW condition (paired *t*-test, $P = 0.0001$). Thus the preference reversals we observed in our data set were more likely to go in the *opposite* direction from that predicted by (quasi-)hyperbolic discounting. This effect is completely consistent, however, with the trivially larger discount rates in the 60 DAY condition described earlier.

(Quasi-)hyperbolic discounting, though, predicts preference reversals for only a subset of choice pairs, so including all choice pairs (as we did in Fig. 3*C*) may not be the fairest test of this prediction of more standard models. Therefore using the discount rate estimated for each individual in the NOW condition, we isolated only those choice pairs where standard hyperbolic discounting (*Eq. 1*) specifically predicts that a participant should make an impulsive preference reversal. The predicted number of reversals ranged across participants from 0 to 56% of choice pairs (mean: $8.00 \pm 1.36\%$), with the number of predicted reversals depending on the steepness of the discount function and increasing for more impatient participants. We examined this subset of the behavioral data, which included data from 17 subjects and 36 sessions, for evidence of systematic preference reversals. Critically, even for these choices where the standard hyperbolic model clearly predicts a reversal, participants still made the same choice in both conditions the majority of the time (70% of choices, Fig. 3*D*), as predicted by ASAP. Even when they did make a reversal, these reversals were apparently randomly distributed. Participants were as likely to be more impulsive in the 60 DAY condition as in the NOW condition, even in this subset of the data where they should have been most likely to exhibit preference reversals according to (quasi-)hyperbolic models (paired *t*-test, $P = 0.51$).

Since the absence of preference reversals in this data set seemed odd, given the prominence of (quasi-)hyperbolic discounting in the literature, we also performed an additional analysis to see whether some aspect of our procedure had actively eliminated preference reversals from the behavior of our subjects. The analyses in Fig. 3, *B–D* described earlier included only paid sessions. All participants had made ≥ 200 practice choices by this point. To examine whether practice

⁷ The online version of this article contains supplemental data.

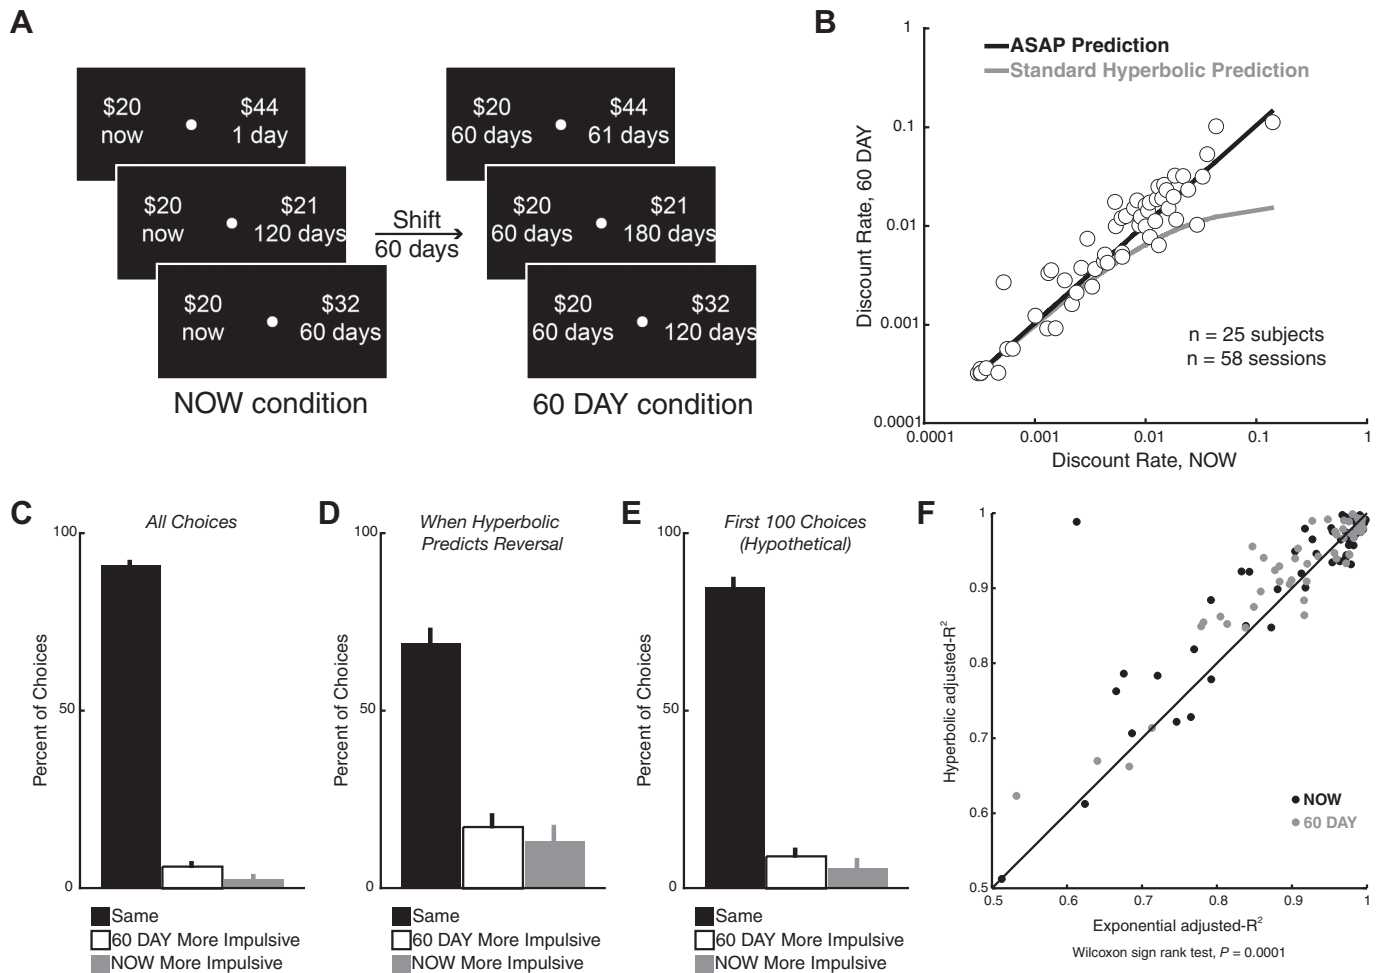


FIG. 3. *A*: schematic illustrating the 2 conditions in *experiment 1*. The NOW condition involved choices between a fixed immediate reward of \$20 and a larger variable reward available after some delay. The 60 DAY condition was constructed from the choices in the NOW condition by adding a fixed delay of 60 days to both options. *B*: the discount rate (k^{ASAP} estimated from Eq. 3) in the 60 DAY condition is plotted as a function of the discount rate in the NOW condition, for all participants and all sessions. Consistent with the ASAP hypothesis, the points lie predominantly on the diagonal, indicating that the discount rates are similar in the 2 conditions. *C*: for pairs of choices differing only by the addition of the fixed 60 day delay, the percentage of choices is shown for which participants made the same choice (both smaller-sooner or larger-later), were more impulsive in the NOW condition (chose smaller-sooner in NOW and larger-later in 60 DAY), or were more impulsive in the 60 DAY condition (chose larger-later in NOW and smaller-sooner in 60 DAY). Participants predominantly made the same choice in both conditions, consistent with ASAP but not standard hyperbolic discounting. *D*: same analysis as in *C*, for only those choices in which standard hyperbolic discounting would predict an impulsive preference reversal. Again participants predominantly made the same choice in both conditions, consistent with ASAP but not standard hyperbolic discounting. *E*: same analysis as in *C* and *D*, for only the first 100 choices the participants faced in the very first session, which involved only hypothetical choices. Even when they did not have extensive experience with the task, participants predominantly made the same choice in both conditions, consistent with ASAP but not (quasi-)hyperbolic discounting. *F*: plotted is the adjusted- R^2 for the ASAP-hyperbolic discount function (Eq. 3) vs. the adjusted- R^2 for the exponential discount function (Eq. 4). Each point represents the fit from a single participant's data from a single session for either the NOW or the 60 DAY condition. Across both conditions, the ASAP-hyperbolic equation provides better fits than the exponential.

with the task could explain our failure to observe preference reversals, we also analyzed the first 100 choices of each participant in the very first practice session involving hypothetical choices. This is the smallest number of choices in which all 50 pairs are represented, thus constituting the smallest data set in which one could systematically search for the existence of preference reversals. (This analysis also excluded the four subjects who had previously participated in a similar experiment in our laboratory, so $n = 21$.) As shown in Fig. 3*E*, the pattern observed in the first 100 hypothetical choices for these subjects is the same as that observed in the full data set of paid sessions. The majority of the time (85% of choices), participants made the same choice in each condition and did not reverse their preference. When the choices did differ, these differences appeared random. Again, there were more in-

stances where participants were more impulsive in the 60 DAY condition, although this difference was not statistically significant (paired t -test, $P = 0.28$).

Participants overweight soonest possible rewards (*experiment 1*)

Of course the ASAP model is not the only model that can explain a failure to observe preference reversals. Another possible explanation for this pattern of results shown in Fig. 3 is that subjective value may decline exponentially with delay for our participants, rather than hyperbolically. Although previous meta-analyses have shown that hyperbolic functions almost always account for intertemporal choice data better than exponential functions (Bickel and Marsch 2001; Frederick et

al. 2002; Green and Myerson 2004; Laibson 1997), our unexpected failure to find preference reversals led us to compare the fit of ASAP to that of the traditional economic model of exponential discounting. The exponential function for subjective value is

$$SV^E = Ae^{-kED} \quad (4)$$

We fit this equation separately in each of the two conditions and therefore obtained different estimates of the exponential discount rate from the NOW and 60 DAY data sets. Since the exponential curve has a constant first derivative, the exponential discounting model predicts that the discount rates estimated in this way should be the same in the two conditions. In other words, exponential discounting, like ASAP, also predicts similar discount rates in the two conditions and no impulsive preference reversals. However, although the exponential *Eq. 4* provided an acceptable fit to the behavioral discount functions in both conditions, the ASAP-hyperbolic *Eq. 3* provided a better fit overall (Fig. 3*F*, Wilcoxon signed-rank test, across both conditions, $P = 0.0001$; for the 60 DAY condition alone, $P = 0.0002$; for the NOW condition alone, $P = 0.086$). Thus the model that best accounted for our participants' choices in our task was the ASAP model, in which choosers hyperbolically discount the value of delayed rewards, but this hyperbolic decline is anchored to the soonest possible reward.

Ventral striatum, medial prefrontal, and posterior cingulate cortex encode subjective value, rather than an immediacy signal (experiment 1)

The behavioral data demonstrate, perhaps surprisingly, that our participants did not make the impulsive preference reversals predicted by (quasi-)hyperbolic models. Our first analysis of the fMRI data were therefore aimed at testing the hypothesis, inspired by the quasi-hyperbolic discounting model, that separate neural systems value immediate and delayed rewards. We constructed a general linear model of the fMRI data that estimated two effects in each condition (NOW or 60 DAY): the mean level of neural activity in that condition, across all trials; and trial-to-trial deviations from that mean that were correlated with the relative subjective value of the larger-later reward ($SV^{ASAP_{g=1}}$, as estimated behaviorally using *Eq. 3*).

The separate neural systems hypothesis proposes that ventral striatum, medial prefrontal, and posterior cingulate cortex encode an immediacy signal, which is high when an immediate reward can be chosen and low in all other situations (McClure et al. 2004a, 2007). This hypothesis predicts that the level of mean activity in the NOW condition should be significantly greater than that in the 60 DAY condition in these regions. Also, according to this hypothesis, neural activity in these regions tracks only whether an immediate reward is available, so changes in the value of the larger-later reward should not affect activity in these regions and there should be no correlation with relative subjective value in either the NOW or the 60 DAY condition.

In contrast, the hypothesis that these regions encode subjective value (as in the ASAP model) predicts a strong correlation with relative subjective value. This hypothesis is also consistent with increased mean activity in the NOW condition compared with the 60 DAY condition, since sooner rewards will

have greater subjective values. However, the absolute predicted size of this difference depends on the parameters of the ASAP model, such as whether the gain factor $g(D_{ASAP})$ is a constant and how impatient each subject is (k^{ASAP}).

Our results were consistent with the hypothesis that ventral striatum, medial prefrontal, and posterior cingulate cortex encode subjective value, but inconsistent with the hypothesis that these regions encode an immediacy signal (see also Glimcher et al. 2007). Replicating our previous findings (Kable and Glimcher 2007), we found that there was a significant effect of relative subjective value ($SV^{ASAP_{g=1}}$) in the NOW condition in ventral striatum, medial prefrontal cortex, and posterior cingulate cortex ($P < 0.05$, corrected; see Fig. 4 and Table 1). However, when we explicitly searched for an "immediacy signal," we found no significant difference in these regions between the mean levels of activity in the NOW and 60 DAY conditions at the whole brain level (at $P < 0.05$, corrected; see Fig. 4, *bottom* and Table 1). Since the whole brain analysis

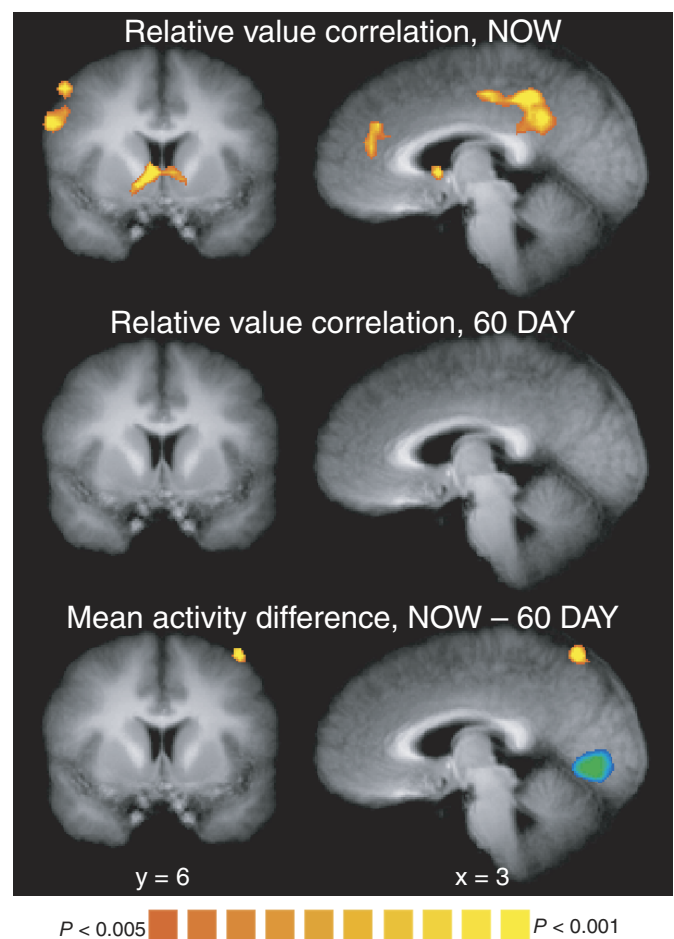


FIG. 4. Group random-effects results are shown for the correlation with relative subjective value (i.e., $SV^{ASAP_{g=1}}$) in the NOW (*top*) or 60 DAY (*middle*) condition and for the difference in mean activity between the NOW and 60 DAY conditions (*bottom*). Each contrast is for time points 4–6 in the trial, which corresponds to a neural effect at the time both options are shown. Maps are thresholded at $P < 0.005$ (uncorrected), with a spatial extent threshold determined for each map so that cluster-level inference is at $P < 0.05$ (corrected). Ventral striatum (coronal) and medial prefrontal and posterior cingulate cortex (sagittal) all exhibited a significant correlation with relative subjective value in the NOW condition, but no significant correlation in the 60 DAY condition and no significant differences in mean activity.

TABLE 1. Location of significant effects in experiment 1 ($P < 0.05$, corrected, cluster-wise inference)

Anatomical Description	Center of Gravity (Talairach Coordinates)			Size, mm ³	Peak Location (Talairach Coordinates)			Peak z -Score
<i>Correlation with relative subjective value, NOW ($P < 0.005$ and cluster size > 924 mm³)</i>								
Posterior cingulate (BA 31/23)	-1	-39	36	8,695	-12	-37	34	4.18
Precentral gyrus/sulcus (BA 6)	47	0	42	5,655	39	-13	49	3.76
Inferior frontal gyrus/lateral orbital (BA 47)	40	23	-10	2,176	33	17	-11	3.54
Precentral gyrus/sulcus (BA 6)	-42	-4	47	1,855	-48	-10	43	3.33
Anterior intraparietal sulcus (BA 40)	-44	-44	52	1,395	-45	-43	49	3.79
Caudate/ventral striatum	5	8	2	1,115	3	5	4	3.55
Medial prefrontal/anterior cingulate (BA 32/24)	2	36	23	1,065	0	35	28	3.13
Inferior frontal gyrus (BA 44)	-50	14	7	1,042	-54	11	7	3.36
<i>Correlation with relative subjective value, 60 DAY ($P < 0.005$ and cluster size > 473 mm³)</i>								
None								
<i>Mean activity difference, NOW - 60 DAY ($P < 0.005$ and cluster size > 782 mm³)</i>								
<i>Medial occipital/calcarine sulcus (BA 17/18)</i>	0	-73	2	14,063	-12	-76	4	-4.36
Posterior middle temporal gyrus/superior temporal sulcus (BA 21/22)	-58	-41	5	2,540	-63	-25	-2	3.57
Posterior parietal (BA 40/7)	-44	-62	46	1,574	-42	-67	46	3.90
Temporal-parietal junction (BA 39/40)	55	-51	27	1,462	54	-46	31	3.53
Posterior parietal (BA 40/7)	40	-48	58	1,387	42	-49	58	3.64
Temporal-parietal junction (BA 39/40)	-63	-44	27	1,080	-63	-43	28	3.68
Anterior intraparietal sulcus (BA 40)	-48	-32	56	933	-57	-25	52	3.45
Angular gyrus (BA 39)	-38	-70	28	901	-39	-70	28	4.45
Precuneus (BA 7)	0	-64	57	882	3	-64	55	3.59
Middle frontal gyrus (BA 6/8)	-34	14	55	808	-30	20	52	3.86

Negative effects are in italics. BA, Brodmann Area.

might have missed weaker effects, we also defined ROIs in the ventral striatum, medial prefrontal, and posterior cingulate cortex. ROIs were defined based on the relative subjective value effect in the NOW condition (see METHODS for details). In each ROI, the difference in mean activity between the NOW and the 60 DAY conditions was positive. This difference was statistically significant in the posterior cingulate cortex ($P = 0.01$, Fig. 5), borderline in medial prefrontal cortex ($P = 0.07$, Fig. 5), and nonsignificant in ventral striatum ($P = 0.25$, Fig. 5; see also Supplemental Fig. S2). As we discuss in more detail in the following text, the size of these differences in mean activity is consistent with what would be predicted if neural activity tracked subjective value. However, the overall pattern of effects fails to support the predictions of the immediacy signal hypothesis. This reinforces our earlier conclusion (Kable and Glimcher 2007) regarding this hypothesis in a new and larger set of participants.

Ventral striatum, medial prefrontal, and posterior cingulate cortex encode subjective value on an absolute rather than a relative scale (experiment 1)

The same statistical model also allowed us to begin to examine whether ventral striatum, medial prefrontal, and posterior cingulate cortex track the subjective value or the *relative* subjective value of immediate and delayed rewards (see Fig. 2). Single neurons in posterior parietal cortex have been found to track relative subjective values (Dorris and Glimcher 2004; Sugrue et al. 2004), whereas neurons in orbitofrontal cortex and striatum appear to track subjective value on a more absolute scale (Padoa-Schioppa and Assad 2006, 2008; Samejima et al. 2005). However, the issue of absolute versus relative valuation has received less attention in human fMRI studies of the subjective values that drive choice, although the issue has

been studied in the context of the experienced values of outcomes (Breiter et al. 2001; Nieuwenhuis et al. 2005). Remember that, in the ASAP model, the distinction between subjective value and relative subjective value corresponds to a distinction between a model in which the gain factor [$g(D_{ASAP})$ in Eq. 2] varies as a function of the delay to the soonest




FIG. 5. The difference in mean activity between the NOW and 60 DAY conditions is plotted for the ventral striatum, medial prefrontal, and posterior cingulate regions of interest (ROIs), for time points 4–6 in the trial. In each case, this effect is normalized relative to that for the relative subjective value effect in the NOW condition. The dotted lines show the predicted size of this effect for each of the 6 combination rules tested in Fig. 6, illustrating that the size of this effect is consistent with the hypothesis that activity in these regions tracks subjective value [i.e., SV^{ASAP} when $g(D_{ASAP})$ is not a constant]. Solid horizontal lines indicate when the observed effect is significantly different from that predicted by the given combination rule. Note that the “Sum” and “Right” combination rules make almost identical predictions regarding the difference in mean activity and are thus not shown separately.

possible reward and a model in which the gain factor is always a constant (i.e., $g = 1$, as in Eq. 3). This then provides a potential neural parameterization of this aspect of the ASAP model.

If these regions track only the relative subjective value of the later reward compared with the sooner reward [a situation in which $g(D_{ASAP})$ is always equal to a constant], then neural activity should show a similarly sized effect (of relative subjective value, $SV^{ASAP_{g=1}}$, as estimated behaviorally from Eq. 3) in both the NOW and the 60 DAY conditions. However, even though we found a robust effect of relative subjective value in these regions in the NOW condition, this effect did not reach significance in the 60 DAY condition in any region (at $P < 0.05$, corrected; see Fig. 4, *middle* and Table 1). Even when using a more sensitive ROI analysis, the relative subjective value effect in the 60 DAY condition did not reach significance in the ventral striatum ($P = 0.11$), medial prefrontal cortex ($P = 0.52$), or posterior cingulate cortex ($P = 0.28$; see also Supplemental Fig. S2). Thus we failed to find evidence that ventral striatum, medial prefrontal, and posterior cingulate cortex encode only *relative* subjective value. (Note that this finding taken alone is also inconsistent with these regions encoding only the magnitude of rewards and ignoring delays, since the monetary amounts are also exactly the same in the NOW and the 60 DAY conditions.)

However, although these results suggest that ventral striatum, medial prefrontal, and posterior cingulate cortex do not encode *relative subjective value* [i.e., they are incompatible with an ASAP model in which $g(D_{ASAP})$ is equal to a constant], they are consistent with the possibility that these regions encode a more absolute form of *subjective value* [i.e., they are compatible with an ASAP model in which $g(D_{ASAP})$ varies as a function of delay; see Fig. 2]. If these regions track a more absolute form of subjective value, then the effect of a relative subjective value ($SV^{ASAP_{g=1}}$) as a regressor should be greater in the NOW condition than that in the 60 DAY condition, since the neural modulations in the NOW condition would be greater than those in the 60 DAY condition. This is what we observed.

Testing BOLD combination rules: how does the BOLD signal reflect the subjective values of multiple options? (experiment 1)

If neural activity in ventral striatum, medial prefrontal, and posterior cingulate cortex does encode subjective value, as these data imply, this raises an additional question: How does the BOLD signal combine the subjective values of the multiple options presented to a subject on a given trial? Previous studies in which more than one option is presented at the same time have usually assumed that BOLD activity tracks the sum of the values of the two options or the value of the chosen option. There are other possibilities, however. Table 2 lists six possible BOLD combination rules for the current experiment: the sum of the values of both options (“Sum”), the value of the larger-later option alone (“Larger-Later”), the larger-later value minus the smaller-sooner value (“Later Minus Sooner”), the value of the chosen option (“Chosen”), the chosen value minus the unchosen value (“Chosen Minus Unchosen”), and the value of the option presented on the right/left (“Right”).

These six combination rules were selected from a larger list of possible combination rules (Supplemental Table S1), based

TABLE 2. Likely possible combination rules for the subjective values of two options

BOLD Combination Rule	Value
“Sum”	$SV_1 + SV_2$
“Larger-Later”	SV_L
“Later Minus Sooner”	$SV_L - SV_S$
“Chosen”	SV_C
“Chosen Minus Unchosen”	$SV_C - SV_U$
“Right” or “Left”	SV_R or SV_L

The superscript on SV is omitted here for clarity, since the combinations could apply to any subjective value function (Eqs. 1–4). The subscripts refer to how the two options are distinguished. Further possible combinations are listed in Supplemental Table S1 and are used in Supplemental Figs. S1 and S2.

on the numerical predictions of each rule for the modulations in BOLD activity in both the NOW and the 60 DAY conditions. These predictions were calculated using the discount functions estimated for each individual subject and the choices presented to each participant. Only the six combinations of subjective values (i.e., SV^{ASAP}) listed in Table 1 would lead to differential modulations in the NOW and 60 DAY conditions, as we observed. The other combination rules tested predict similar modulations in these two conditions and thus can be discarded for this data set (Supplemental Fig. S3). Importantly, these calculations also showed that *any* combination of relative subjective values (i.e., $SV^{ASAP_{g=1}}$) predicts similar modulations in the NOW and 60 DAY conditions and is therefore inconsistent with the results presented earlier (Supplemental Fig. S4).

Thus six different combinations of ASAP subjective values are broadly consistent with our finding of significant value-related modulations in the NOW condition but not in the 60 DAY condition. All of these combinations predict large trial-to-trial modulations in the NOW condition (in fact, “Sum,” “Larger-Later,” and “Later Minus Sooner” make the *same exact* predictions when restricted to the NOW condition alone, as do “Chosen” and “Chosen Minus Unchosen”; see Supplemental Fig. S3) and smaller trial-to-trial modulations in the 60 DAY condition. Further testing between these combination rules, however, requires making some assumption about the gain function $g(D_{ASAP})$, since it cannot be estimated behaviorally. Here we assume as the simplest possible parameterization that $g(D_{ASAP})$ is the hyperbolic discount factor associated with the soonest possible reward

$$g(D_{ASAP}) = \frac{1}{1 + k^{ASAP} D_{ASAP}} \quad (5)$$

We have examined several related gain functions, such as the exponential discount factor associated with the soonest possible reward. The results discussed in the following text do not hinge on this particular assumption (indeed there may be good theoretical reasons to assume the exponential form). We consider Eq. 5 an ad hoc assumption of a functional form necessary for testing between BOLD combination rules, rather than a critical aspect of the ASAP model.

As discussed earlier, we observed modest differences in mean activity between the NOW and 60 DAY conditions in the ventral striatum, medial prefrontal, and posterior cingulate ROIs. Figure 5 shows that the six BOLD combination rules listed in Table 1 are broadly consistent with this finding. The differences in mean activity presented in Fig. 5 are normalized

for each participant by the relative subjective value effect in the NOW condition, so that the units of these differences can be interpreted as “subjective dollars immediately.” In all three ROIs, the difference in mean activity we observed was within the range of that expected if these regions represented some combination of subjective values (i.e., SV^{ASAP}) and, in each case, was consistent with most of the six combination rules.

We next directly tested which BOLD combination rule best accounted for trial-by-trial modulations in neural activity in the ventral striatum, medial prefrontal, and posterior cingulate ROIs. To do this, we performed an additional set of analyses that tested for correlations with different combinations of ASAP subjective values, combining data across both the NOW and the 60 DAY conditions. The subjective value regressors in this case were normalized so that the beta values from these different analyses could be compared (Hampton et al. 2008; Hare et al. 2008). Figure 6 reports this test for the first five of the BOLD combination rules (the results of the right-only or left-only combination rule are presented in the following text) and includes two control analyses for comparison: 1) one testing for correlations with only the relative subjective value of the larger-later reward (i.e., $SV^{ASAP}_{g=1}$, which we call “Relative”) and 2) one testing for correlations with only the magnitude of the larger-later reward, ignoring the delay (“Amount”). In all three regions, the *sum*, *larger-later*, and *chosen* combination rules provided the best fits and could not be statistically distinguished (at $P < 0.05$, paired *t*-test, one-tailed). In ventral striatum, the best-fitting combination was the subjective value of the larger-later reward, whereas in medial prefrontal, and posterior cingulate cortex the best-fitting combination was the sum of the subjective values. In all three regions, the best-fitting combination provided a significantly better fit than the worst-fitting combination and, in some cases, further differences were statistically significant (at $P < 0.05$, paired *t*-test, one-tailed). The differences between many pairs of these combination rules, however, were small. These tests do reinforce our earlier conclusion that these regions encode subjective value [i.e., $g(D_{ASAP})$ is not constant] rather than relative subjective value [i.e., $g(D_{ASAP})$ is constant]. However, we are unable to draw definitive conclusions about the precise combination rule expressed by the BOLD signal in these areas.

Anterior parietal cortex encodes the subjective value of the contralateral option (experiment 1)

In addition to the three regions discussed earlier, there were five additional areas in which activity was significantly correlated with the relative subjective value of the changing later reward ($SV^{ASAP}_{g=1}$) in the NOW condition ($P < 0.05$, corrected). These were left anterior parietal cortex and dorsal premotor and ventrolateral prefrontal cortices bilaterally (Table 1, Fig. 7). In our previous study (Kable and Glimcher 2007), we did find an effect of subjective value in left anterior parietal cortex in some analyses, but we did not find such an effect previously in any of the prefrontal areas (Fig. 7). In the four frontal areas observed here, neither the relative subjective value effect in the 60 DAY condition nor the difference in mean activity between the two conditions reached significance (Table 1). There was a different pattern of effects in left anterior parietal cortex, in which the difference in mean activity between the two conditions did reach significance (Table 1).

Left anterior parietal cortex was also the only region that displayed a laterality effect: neural activity was significantly more correlated with the subjective value (SV^{ASAP}) of the option presented on the right (and requiring a right-hand motor response), compared with the option presented on the left (Fig. 7). Note that we could not have tested for such a laterality effect in our previous study, in which options were presented centrally and a right-handed “go”–“no go”-style response was intentionally used. Thus the pattern of effects in left anterior parietal cortex we observed here is consistent with activity in the posterior parietal cortex tracking the subjective value of the option presented on the contralateral side or associated with a contralateral motor response (Glimcher 2009; Glimcher et al. 2007).

Is there evidence for a δ system? (experiment 1)

In addition to proposing that ventral striatum, ventromedial prefrontal, and posterior cingulate cortex carry an immediacy signal, the separate neural systems hypothesis also proposes that dorsolateral prefrontal and posterior parietal regions carry an opponent signal that more equally values immediate and delayed rewards. We did find prefrontal and parietal regions in which neural activity scaled with subjective value. However, the hypothesized δ system was defined based on three properties (McClure et al. 2004a, 2007): increased activity for all choices compared with rest, increased activity for difficult compared with easy choices, and greater activity when a delayed reward is chosen compared with when an immediate reward is chosen. In our data set, we saw increased activity across much of the brain, including prefrontal and parietal regions, for all choices compared with rest ($P < 0.05$, corrected, Fig. 8A). Using reaction time as an index of difficulty, we also saw a correlation between activity and difficulty in prefrontal, parietal, and anterior cingulate regions, although this occurred at the time of the response rather than when the choice options appeared ($P < 0.05$, corrected; Fig. 8, B and C). However, we found no regions in which activity was significantly greater on trials where the later reward was chosen, compared with when the sooner reward was chosen (at $P < 0.05$, corrected; Fig. 8D). We consider this last contrast the key prediction regarding the δ system—the first two effects have many possible interpretations. We therefore conclude that there is not strong evidence in our data set for a δ system or a set of regions that shows uniformly increased activity for all choices of delayed rewards over immediate rewards; however, this does not preclude a role for these regions in a more specific subset of choices (Hare et al. 2009).

Ventral striatum, medial prefrontal, and posterior cingulate cortex do carry a value signal during choices between delayed rewards (experiment 2)

To summarize, the results of the first experiment suggest that ventral striatum, medial prefrontal, and posterior cingulate cortex track the subjective value of immediate and delayed rewards, on a more absolute rather than purely relative scale. However, this hypothesis predicts that there should be some measurable value signal in these regions during choices involving only delayed options, yet the relative subjective value effect was not significant in any of these three regions in the 60

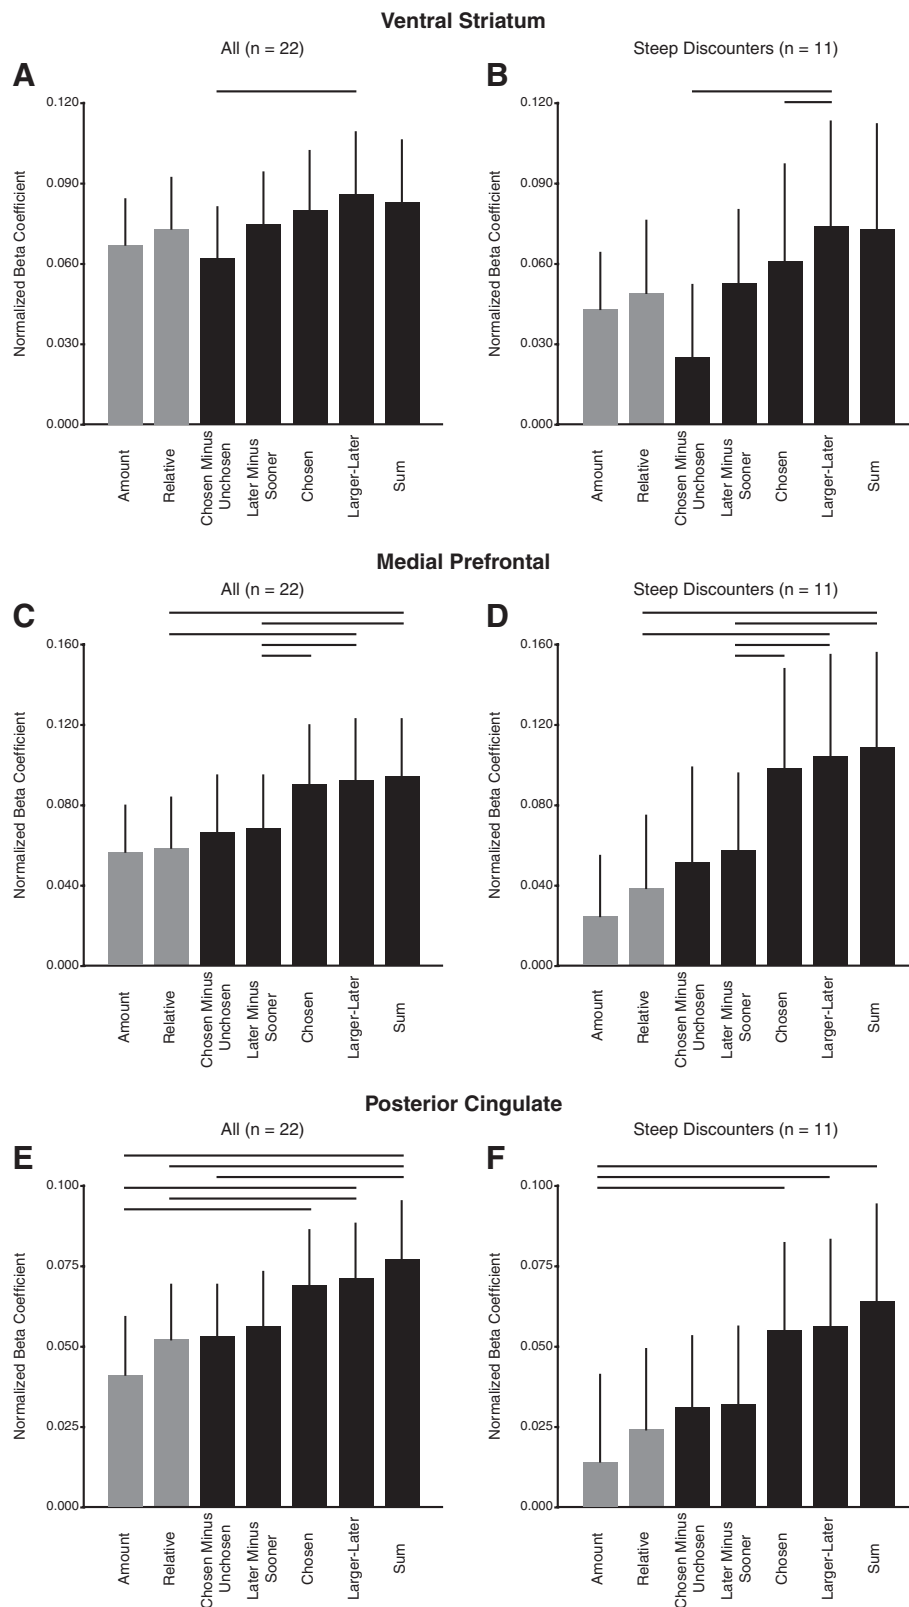


FIG. 6. A model comparison is shown for the ventral striatum (*A* and *B*), medial prefrontal (*C* and *D*), and posterior cingulate (*E* and *F*) ROIs. Normalized beta coefficients for the subjective value covariate [specifically, SV^{ASAP} with $g(D_{ASAP})$ equal to Eq. 5] are plotted for 7 different combination rules, which differ in how they assume the subjective values of 2 options are combined. These coefficients are for the contrast involving time points 4–6 in the trial, which corresponds to a neural effect at the time both options are shown. Comparisons including all participants are shown in *A*, *C*, and *E*. Since the correlation between the predictions of different rules is high in patient participants, comparisons for the 11 most impatient participants are shown in *B*, *D*, and *F*. Lines indicate paired comparisons where one combination rule provides a significantly better fit ($P < 0.05$, paired *t*-test, one-tailed).

DAY condition alone. This observation suggests either that no signal occurs in these areas for the 60 DAY condition alone or that the signal exists but is beneath our measurement threshold. Although the latter seems likely and is consistent with our findings and reasoning so far, it would be reassuring to have

some positive evidence for a value signal during choices between delayed options. The search for this positive evidence thus motivated our second experiment. Ten subjects from the first experiment participated in this second experiment, which involved one session in the MRI scanner. (The

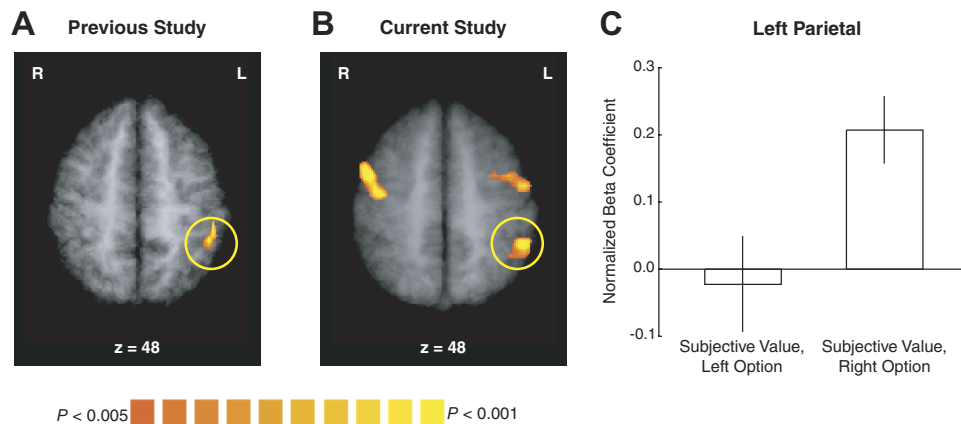


FIG. 7. **A**: group random-effects results from our previous study are shown, illustrating a correlation with the subjective value of the changing delayed reward in left anterior parietal cortex. These results are from the analysis assuming a hemodynamic response function reported in the Supplemental Materials of Kable and Glimcher (2007). The map is thresholded at $P < 0.005$ (uncorrected), with a spatial extent threshold of 100 mm^3 . **B**: group random-effects results from the current study are shown, for the correlation with relative subjective value in the NOW condition. This contrast is for time points 4–6 in the trial, which corresponds to a neural effect at the time both options are shown. The map is thresholded at $P < 0.005$ (uncorrected), with a spatial extent threshold so that cluster-level inference is at $P < 0.05$ (corrected). Correlations in left anterior parietal cortex as well as premotor areas bilaterally can be seen in this slice. **C**: results from an analysis in which the subjective values of the left and right options were modeled separately. As in Fig. 6, we calculate the subjective values of the left and right options using the ASAP model [specifically, SV^{ASAP} with $g(D_{ASAP})$ equal to Eq. 5] and model the NOW and 60 DAY conditions together. In the left anterior parietal region shown in **B**, only the correlation with the subjective value of the right option is significant and this correlation is significantly greater than that with the left option. This contrast is for time points 4–6 in the trial.

behavioral data for this experiment, which by design do not provide as clean a test between standard hyperbolic discounting and ASAP as that in *experiment 1*, are shown in Supplemental Fig. S5.)

The parameterization of the ASAP model that we found support for in the first experiment assumes that the subjective value of the later reward varies with delay differences in a similar manner in the NOW and 60 DAY conditions, only that the absolute size of these neural signals is reduced in the 60 DAY condition by some gain function [$g(D_{ASAP})$]. If the small absolute size of these signals is the reason that the relative subjective value correlation failed to reach significance in the 60 DAY condition, then increasing the gain on these signals by scaling the amounts of the rewards should recover a significant correlation, further supporting the ASAP model. Accordingly, to construct the options for the SCALED 60 DAY condition, we not only added a delay of 60 days to both options from the NOW condition, but also increased the reward amounts proportionally by a scaling factor. This scaling factor was chosen individually for each participant based on previous sessions, so that the sooner reference option (\$X in 60 days) would be weakly preferred (almost identical in subjective value) to \$20 now. By scaling-up all of the amounts in the SCALED 60 DAY set of choices in this way, we should make the overall subjective values of the SCALED 60 DAY choice set roughly

equivalent to the subjective values of the NOW choice set. Thus the critical question in this experiment is then whether scaling the amounts in this way improves the correlation between BOLD and relative subjective value in the SCALED 60 DAY condition, as would be predicted by the ASAP model with a nonconstant gain function.

Given the smaller number of participants in this experiment, we did not expect that this effect would reach a corrected level of significance in a whole brain random-effects analysis. However, as seen in Fig. 9, increasing the amounts improved the effect of relative subjective value in the SCALED 60 DAY condition to a level that is similar to that in the NOW condition in the first experiment for these same ten participants. The subjective value regressors were normalized for these analyses to allow comparison of the beta values. If these regions track subjective value [i.e., $g(D_{ASAP})$ is not a constant], the prediction is that the normalized betas will be (nearly) identical in the NOW and SCALED 60 DAY conditions and lower in the 60 DAY condition from *experiment 1*, which is exactly the pattern we observed. [The prediction, if these regions track relative subjective value $g(D_{ASAP})$ is equal to a constant, is that the normalized betas in all three conditions, including the 60 DAY condition, will be the same. This possibility is ruled out by the neural data from *experiment 1*.] At the ROI level, the SCALED 60 DAY effect reached significance in ventral striatum ($P =$




FIG. 8. Group random-effects results are shown, for increased activity compared with rest in both the NOW and 60 DAY conditions (**A**), for the correlation between activity and reaction time (**B** and **C**), and for greater activity when the later reward was chosen (**D**). The contrasts in **A**, **B**, and **D** are for time points 4–6 in the trial, which corresponds to a neural effect at the time both options are shown. The contrast in **C** is for time points 8–10 in the trial, which corresponds to a neural effect at the time when the participant responds. All maps are thresholded at $P < 0.005$ (uncorrected), with a spatial extent threshold so that cluster-level inference is at $P < 0.05$ (corrected).

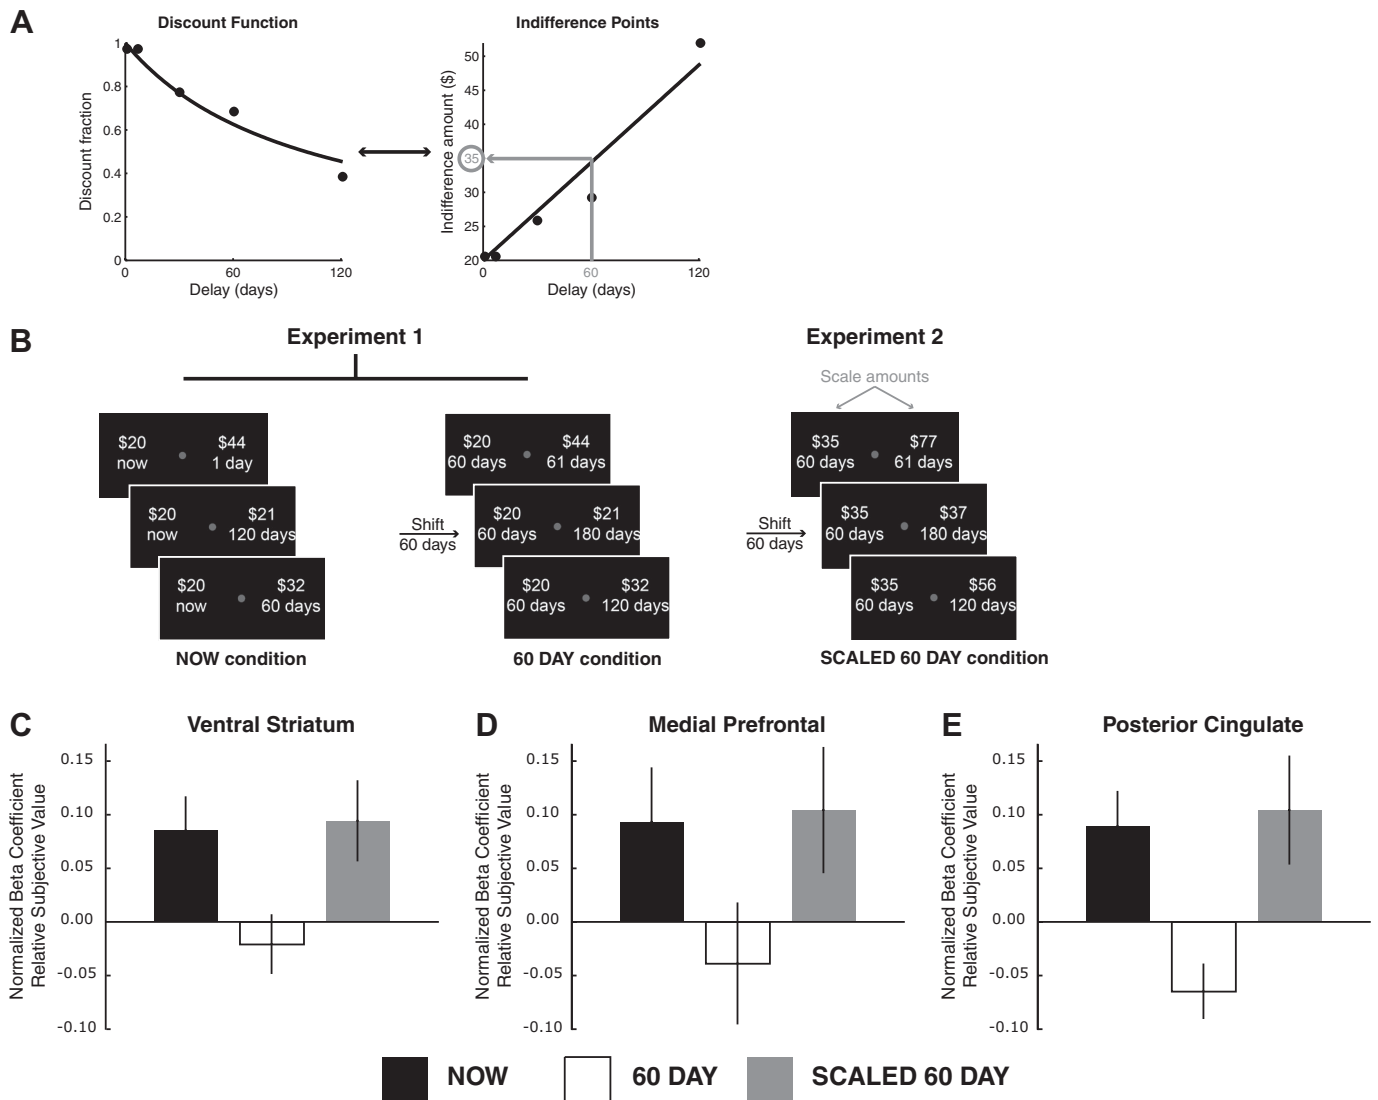


FIG. 9. *A*: schematic illustrating how the rewards in the SCALD 60 DAY condition were determined in *experiment 2*. For each participant, a discount function was estimated in the NOW condition in *experiment 1*, which describes how the subjective value of a reward declines with delay compared with an immediate reward reference. This hyperbolic discount function corresponds to a line in indifference point space, which connects the amount of money at each delay that is subjectively equivalent to \$20 immediately. In the second experiment, the amount of the fixed 60 day reward was chosen so that that reward would be subjectively equivalent to \$20 now, according to that participant's discount function. In the example, this amount was \$35 in 60 days. *B*: schematic illustrating the conditions in *experiment 1* and *experiment 2*. *C–E*: for the 10 subjects that participated in both experiments, the normalized beta coefficients are shown for the correlation with relative subjective value in the NOW and 60 DAY conditions in the first experiment and for the SCALD 60 DAY condition in the second experiment. Each contrast is for time points 4–6 in the trial, which corresponds to a neural effect at the time both options are shown. Scaling the amounts in *experiment 2* recovers a correlation with relative subjective value in ventral striatum (*C*), medial prefrontal (*D*), and posterior cingulate cortex (*E*).

0.03) and a similar nonsignificant trend was seen in medial prefrontal ($P = 0.11$) and posterior cingulate ($P = 0.07$). These significance levels are similar to those seen in the NOW condition when the data set is limited to only these ten subjects (ventral striatum, $P = 0.02$; medial prefrontal, $P = 0.10$; posterior cingulate, $P = 0.02$). Thus the results of this second experiment support the idea that these three regions do encode an absolute subjective value signal of some kind [$g(D_{ASAP})$ is not equal to a constant] during choices between only delayed rewards and that this signal was just too small to reliably detect in the original 60 DAY condition.

DISCUSSION

Behaviorally, we found that in our experiments the subjective value of a larger-later reward declined hyperbolically at a

similar rate when it was compared with either an immediate reward or a reward in 60 days (Glimcher et al. 2007). Our subjects therefore did not make impulsive preference reversals in our task, as predicted by (quasi-)hyperbolic discounting. We have called this pattern of choices “as-soon-as-possible” (ASAP) discounting and these findings (with several others in the literature) raise the possibility that (quasi-)hyperbolic models may not *always* provide the best description of human intertemporal choice behavior. Neurally, we found that BOLD activity in ventral striatum, medial prefrontal, and posterior cingulate cortex tracked the subjective values (in a more absolute sense) of immediate and delayed rewards, rather than only the *relative* subjective values of rewards [as would be the case if $g(D_{ASAP})$ were equal to a constant]. These regions did not respond to immediate rewards alone and activity did not

increase for immediate rewards separately from modulations in subjective value. Finally, neural activity also scaled with subjective value in dorsal premotor and anterior parietal areas, suggesting that value signals may influence motor planning through pathways related to those previously implicated in visuosaccadic decision making (Dorris and Glimcher 2004; Glimcher 2009; Glimcher et al. 2007).

(Quasi-)hyperbolic models predict two technical “irrationalities” (in the economic sense) in intertemporal choice behavior. First, these models predict that people overvalue sooner rewards in comparison with later rewards—this reflects the hyperbolic curvature of the discount function or the quasi-hyperbolic functions that emulate it. Second, these models predict impulsive preference reversals because the discount function is both hyperbolic *and* this hyperbola is anchored to the immediate present. Although we did observe overvaluation of sooner rewards (hyperbolic curvature of the discounted utility function), we did not observe impulsive preference reversals. Our failure to observe the second of these two properties predicted by (quasi-)hyperbolic models was unexpected. These impulsive preference reversals are a “stylized fact” in the intertemporal choice literature and have been offered as the explanation for impulsive behaviors, varying from low savings rates to drug addiction (Ainslie and Haendel 1983; Frederick et al. 2002; Laibson 1997). Since our behavioral finding was therefore quite unexpected, we believe that resolving this discrepancy with the previous literature merits an extended discussion.

Although dozens of studies have found greater support for hyperbolic compared with exponential discounting, most of these studies have used choices in which one option is always available immediately. These studies, like ours, do demonstrate an overvaluation of soonest rewards but can, in principle, provide no evidence about the existence of impulsive preference reversals. Like our previous work on this topic (Kable and Glimcher 2007), these studies cannot distinguish between (quasi-)hyperbolic and ASAP models and are silent about the existence of preference reversals. Many fewer studies have systematically changed the delay to the soonest reward, as is needed to directly measure the impulsive preference reversals that distinguish the two hypotheses. Among the smaller number of studies that specifically tested for impulsive preference reversals (in humans making choices about monetary rewards, as in the current experiment), it is important to acknowledge that the empirical findings are actually quite mixed.

Impulsive preference reversals of the kind predicted by (quasi-)hyperbolic models have been found in some studies (Ainslie and Haendel 1983; Casari 2006; Green et al. 1994; Kirby and Herrnstein 1995). One notable feature of these experiments, however, is that they all used a procedure that explicitly searched for (typically a single) preference reversal in each subject. In those experiments, subjects started with a choice between a smaller immediate option and a larger delayed option, under conditions in which it had previously been determined that the subjects would choose the immediate option. On consecutive trials, a gradually increasing common delay was added to both options. The procedure then terminated either when the subject chose the larger-later option or when a limit on the number of trials was reached. Most subjects in these procedures make at least one impulsive reversal, as predicted by (quasi-)hyperbolic models. Unfortunately, because none of these experiments measured individual

discount rates, the stronger prediction of (quasi-)hyperbolic models that the impulsive preference reversal should occur at a specific delay for each subject was not examined. Also, because subjects were not offered the opportunity to make choices that would result in patient preference reversals, the other strong prediction of (quasi-)hyperbolic models that this kind of reversals should not occur has not been verified.

It is also interesting to note that, although less widely cited, almost as many studies as report preference reversals have reported choice behavior that is generally consistent with ASAP (Ahlbrecht and Weber 1997; Baron 2000; Holcomb and Nelson 1992; Read 2001; Read and Roelofsma 2003). These studies used designs similar to ours, involving two sets of decisions in which the only difference between them was the addition of a common delay to both options. These studies differed in whether indifference points were inferred from choice or directly elicited, in whether the sooner or later reward varied across trials, and in how the options were presented visually. In each case, the measured discount rates decreased with an increasing delay between the two options, but did *not* depend on the delay to the smaller-sooner option.

Other studies have found an intermediate pattern (between standard hyperbolic and ASAP), although using much longer (hypothetical) delays (Green et al. 2005) or only in matching paradigms (direct elicitation of indifference points) and not in choice (Ahlbrecht and Weber 1997; Read and Roelofsma 2003). One additional study found greater support for a hyperbolic than for an ASAP model, although this study was not designed to directly test for preference reversals (Pine et al. 2009).

Overall, this review of the literature shows that experimental protocols designed to test the stronger assumptions of (quasi-)hyperbolic models may actually elicit classes of behavior at variance with these models. Although we do not doubt that impulsive preference reversals occur under some conditions, these results highlight the conclusion that impulsive preference reversals are not a *universal* aspect of human intertemporal decision making and raise the possibility that more than one model may be required to account for human discounting behavior.

Although these previously cited experiments suggest that impulsive preference reversals, of the kind predicted by (quasi-)hyperbolic models, are not always observed, our particular failure to observe impulsive preference reversals might still be due to some unique aspect of the procedure we used. There are two specific aspects of our design that differ from the other studies reviewed earlier and the impact of these procedures must be considered in evaluating our results. First, our subjects had more experience and practice making choices in the task. For example, the studies that have found evidence for impulsive preference reversals tested either nearly 100 subjects on about 10 choices (Ainslie and Haendel 1983; Casari 2006), or about 20–30 subjects on nearly 100–200 choices (Green et al. 1994; Kirby and Herrnstein 1995). In our data set, 25 participants made 400–800 choices each.

However, since our failure to observe impulsive preference reversals was evident even in the first half of the very first session (the first 100 choices), it seems unlikely that extensive practice with the task alone accounts for our failure to observe preference reversals. Second, we used a novel payment mechanism in our experiment. In studies that found evidence for impulsive preference reversals, either all choices were hypothetical (Ainslie and Haendel 1983; Green et al. 1994) or there

was a stated probability that one of a participant's choices would be paid (Casari 2006; Kirby and Herrnstein 1995). In our experiment, each participant was paid 4–12 choices using a bank-sponsored prepaid debit card and each subject therefore gained extensive experience receiving delayed payments from us. Although we also did not observe preference reversals in our hypothetical choice sessions, we cannot rule out the possibility that our novel payment mechanism made preference reversals even less likely in the paid sessions. This finding raises the possibility that the preference reversals that have been observed previously may involve an interaction between risk and delay-to-payment that has not been widely discussed.

Importantly, although our subjects did not make impulsive preference reversals, our findings do not imply that these subjects would be completely consistent in their intertemporal trade-offs, as they would be if they had used an exponential discount function. Our subjects did show an overvaluation of the soonest rewards characteristic of all hyperbolic-like discount functions. It is important to realize that subjects whose behavior was well modeled with ASAP would exhibit intransitive choice cycles (Ok and Masatlioglu 2005; Read and Roelofsma 2003). These subjects would thus be inconsistent in their choices, like (quasi-)hyperbolic choosers, but the inconsistency would take a very specific form that differs from the impulsive preference reversals made by (quasi-)hyperbolic choosers. For example, an ASAP chooser might select \$45 in 120 days over \$20 now and \$20 now over \$30 in 60 days. Based on these two choices, a transitive decision maker would obviously have to prefer \$45 in 120 days over \$30 in 60 days. An ASAP chooser, however, overweights the value of the soonest possible reward (whenever that reward may arrive) and thus could choose \$30 in 60 days over \$45 in 120 days instead—a clear violation of the principle of transitivity. Models such as ASAP, in which choices display the property of stationarity (choices are not affected by changes in a common delay) but violate transitivity (formally that there is no strict ordering of preferences over outcomes), have been previously characterized in the economic literature (Ok and Masatlioglu 2005). Although our design did not allow us to test for the intransitivities predicted by ASAP, similar intransitive intertemporal choice patterns have previously been observed (Roelofsma and Read 2000).

Whereas the behavioral data we present here are inconsistent with (quasi-)hyperbolic discounting, the neural data are inconsistent with the proposal, inspired by the quasi-hyperbolic behavioral model, that separate neural systems value immediate and delayed rewards (McClure et al. 2004a, 2007). In that view, ventral striatum, medial prefrontal, and posterior cingulate cortex carry an impulsive signal that exclusively or primarily values immediate rewards. That hypothesis predicts that in the present experiment: 1) there should be significantly greater mean activity in these areas in the NOW condition, which involves an immediate reward, compared with the 60 DAY condition, which involves only delayed rewards; and 2) there should be only a weak correlation, if any, with the relative subjective value of the delayed reward in the NOW condition. However, we found the opposite pattern in these regions, a strong correlation with relative subjective value, and more modest differences in mean activity. These findings are consistent with the idea that these regions encode a subjective value signal for both immediate and delayed rewards and are

inconsistent with the notion that these brain areas carry an impulsive signal that primarily values immediate rewards, as required by the separate neural systems hypothesis. This extends a conclusion we have drawn previously (Glimcher et al. 2007; Kable and Glimcher 2007) and is also in line with other recent studies implicating these regions in the valuation of delayed rewards (Ballard and Knutson 2009; Pine et al. 2009; Tanaka et al. 2004).

We should note two important caveats about this conclusion. First, there are numbers of design differences across these different studies (Ballard and Knutson 2009; Kable and Glimcher 2007; McClure et al. 2004a, 2007; Pine et al. 2009; Tanaka et al. 2004) that could account for the differing conclusions drawn in these studies. Importantly, as Fig. 5 demonstrates, our hypothesis that these regions encode subjective value does not preclude the observation that these regions exhibit greater activity when an immediate reward is available. Indeed, we observed modestly greater activity in these areas for immediate rewards in the first experiment, presumably because subjects valued the immediate rewards more under these conditions. Second, it is important to note that our conclusions are meant to be specific to the separate neural systems hypothesis outlined in the INTRODUCTION and should not be generalized to all neural models of intertemporal choice that include multiple valuation components. Our data are inconsistent with models of monetary intertemporal choices that propose a specialized neural system for valuing immediate rewards alone, but we cannot extend this conclusion to other multiple-systems models that we have not tested (Daw et al. 2005; Dayan and Balleine 2002; Rangel et al. 2008) or other types of intertemporal choice tasks that we have not examined.

We also found that ventral striatum, medial prefrontal, and posterior cingulate cortex encode subjective value on an absolute temporal scale, rather than encoding only the *relative* subjective value of the later reward compared with the sooner reward. [Within the framework of ASAP, we found that the function $g(D_{ASAP})$ could not be treated as a constant.] In the first experiment, activity in ventral striatum, medial prefrontal, and posterior cingulate cortex was significantly correlated with the relative subjective value of the larger-later reward in the NOW condition, but not in the 60 DAY condition. If neural activity in these regions encodes relative subjective value alone, then there should have been a similar-sized effect in the 60 DAY condition. The second experiment further explores this negative result regarding choices between delayed rewards. In that experiment, scaling the reward amounts so that the options in the NOW and SCALED 60 DAY conditions are similar in subjective value leads to similar-sized neural modulations, thus demonstrating that there are indeed value signals in these regions for choices between delayed rewards. This conclusion is also consistent with data demonstrating that single neurons in these regions encode subjective value on a more absolute scale (Padoa-Schioppa and Assad 2006, 2008; Samejima et al. 2005). This contrasts with our previous proposal that single neuron activity in the lateral intraparietal area (LIP) and areas within the proposed frontoparietal choice network generally track relative subjective values (Dorris and Glimcher 2004; Glimcher 2009; Glimcher et al. 2007). This difference in the way the subjective value is encoded may then represent a fundamental difference between valuation and choice circuits, as we have discussed in more detail elsewhere (Glimcher 2009).

Here we have proposed that ventral striatum, medial prefrontal, and posterior cingulate cortex encode the subjective value of immediate and delayed rewards and that these subjective values can directly account for the ASAP effect in choices. We have argued that these regions encode SV^{ASAP} from Eq. 2. Another possibility, which we have raised elsewhere and which is not ruled out by these data, is that these regions encode subjective values that are *exponentially discounted* (SV^E) and that the ASAP effect arises because of a choice process that involves divisive normalization (Dorris and Glimcher 2004; Glimcher 2009; Glimcher et al. 2007) in frontoparietal choice networks. Such a hypothesis would suggest that valuation circuits compute a more nearly rational discounted subjective value that is then induced to violate transitivity (or, more formally, to violate the axiom of Ordering derived by Fishburn and Rubinstein 1982) in the frontoparietal choice mechanism. In either case, our findings add to the growing literature indicating that BOLD activity, particularly in ventral striatum or ventromedial prefrontal cortex, scales with the values of particular options during decision making (Daw et al. 2006; Hampton et al. 2006; Hare et al. 2008; McClure et al. 2004b; Plassmann et al. 2007; Tanaka et al. 2004; Tom et al. 2007). The idea that these regions encode subjective value is also consistent with previous neurophysiological studies in these areas (Lau and Glimcher 2008; McCoy and Platt 2005; Padoa-Schioppa and Assad 2006, 2008; Samejima et al. 2005).

Our findings also raise a question that has not been systematically addressed in previous studies: if BOLD activity is correlated with value, what combination of values does it track when there is more than one option on offer during decision making? Most previous fMRI studies showing such correlations either had only one option (Hare et al. 2008; McClure et al. 2004b; Plassmann et al. 2007), fixed the value of one of only two options (Kable and Glimcher 2007; Tom et al. 2007), or assumed a particular combination rule (Daw et al. 2006; Hampton et al. 2006; Tanaka et al. 2004). Although the design of our study did not permit us to uniquely identify a combination rule, activity in ventral striatum, medial prefrontal, and posterior cingulate cortex was best correlated with the sum of values, the chosen value, and the value of the larger-later option alone. This finding is reassuring, since the sum or chosen value is usually what is assumed in fMRI studies.

An important goal for future research should be to distinguish the precise roles of these three brain regions in valuation. Figures 5 and 6 suggest that the precise form of the value signal may be different in ventral striatum, compared with medial prefrontal and posterior cingulate cortex. One recent study suggested that the ventral striatal signals represent prediction errors concerning the value gained from making each decision, whereas the medial prefrontal signals represent subjective values that drive choice (Hare et al. 2008). In our study, such prediction errors and subjective value signals are highly correlated, so we cannot distinguish between the two signals. Another recent study suggested that striatum contains independent signals related to utility, discount, and subjective value (the combination of utility and discount) (Pine et al. 2009). In our study, we did not have the statistical power to distinguish whether there were three distinct value-related signals. In any case, our data suggest that striatal, medial prefrontal, and posterior cingulate regions play complementary roles in valuing both immediate and delayed rewards. Future experiments should build on the work of Hare et al. (2008)

and Pine et al. (2009) to try to characterize exactly what these complementary functions may be.

BOLD activity was also modulated by subjective value in left anterior parietal cortex, as well as dorsal premotor and ventral prefrontal regions bilaterally. Additionally, activity in anterior parietal cortex was best correlated with the value of the option associated with a contralateral motor response. We did not find these effects as prominently in our previous study (Kable and Glimcher 2007); however, our previous design minimized the influence of motor planning, since no motor response was required in half of the trials. Thus these regions, and particularly anterior parietal cortex, may encode the subjective value of potential actions during motor planning. This interpretation is consistent with neurophysiological evidence that neurons in area LIP encode the value of particular actions (Dorris and Glimcher 2004; Platt and Glimcher 1999; Shadlen and Newsome 2001; Sugrue et al. 2004). The fact that we observed this response significantly only in the left hemisphere may be explained by a lack of sensitivity or it may represent a hemispheric difference in humans in the encoding of motor planning signals, which mirrors the hemispheric difference observed in humans for attentional signals (Coulthard et al. 2008). In either case, such representations provide one mechanism whereby the signals we observed in ventral striatum, medial prefrontal, and posterior cingulate could be translated into action (Glimcher 2009; Glimcher et al. 2007).

In conclusion, this study extends our previous work (Kable and Glimcher 2007), which found that ventral striatum, medial prefrontal, and posterior cingulate cortex encode the subjective value of immediate and delayed rewards on a common scale. The behavioral evidence from these experiments is not consistent with standard hyperbolic and quasi-hyperbolic models of intertemporal choice. The neural evidence we have gathered is not consistent with the proposal that separate neural systems value immediate and delayed rewards. Taken together, these data are consistent with a behavioral and neural model that we have called ASAP discounting, in which subjective value declines hyperbolically relative to the soonest reward and a small number of valuation areas serve as a final common pathway through which subjective values guide choice.

ACKNOWLEDGMENTS

We thank J. Lawrence for assistance in piloting the behavioral experiment and D. Burghart, M. Grantner, K. Louie, and J. Stump for comments on earlier versions of this work.

GRANTS

This work was supported by National Institutes of Health Grants F32-MH-75544 to J. Kable and R01-NS-054775 to P. Glimcher and a McDonnell Foundation grant to P. Glimcher.

DISCLOSURES

No conflicts of interest are declared by the authors.

REFERENCES

- Ahlbrecht M, Weber M.** An empirical study of intertemporal decision making under risk. *Manage Sci* 43: 813–826, 1997.
- Ainslie G, Haendel V.** The motives of the will. In: *Etiologic Aspects of Alcohol and Drug Abuse*, edited by Gottheil E, Druley KA, Skoloda TE, Waxman HM. Springfield, IL: Thomas, 1983, p. 119–140.
- Ballard K, Knutson B.** Dissociable neural representations of future reward magnitude and delay during temporal discounting. *NeuroImage* 45: 143–150, 2009.

- Baron J.** Can we use human judgments to determine the discount rate? *Risk Anal* 20: 861–868, 2000.
- Bickel WK, Marsch LA.** Toward a behavioral economic understanding of drug dependence: delay discounting processes. *Addiction* 96: 73–86, 2001.
- Breiter HC, 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.
- Camerer CF, Hogarth RM.** The effects of financial incentives in experiments: a review and capital-labor-production framework. *J Risk Uncertainty* 19: 7–42, 1999.
- Casari M.** *Pre-commitment and Flexibility in a Time Decision Experiment.* West Lafayette, IN: Purdue Univ. Department of Economics, 2006.
- Collier M, Williams MB.** Eliciting individual discount rates. *Exp Econ* 2: 107–127, 1999.
- Coulthard EJ, Nachev P, Husain M.** Control over conflict during movement preparation: role of posterior parietal cortex. *Neuron* 58: 144–157, 2008.
- Daw ND, Niv Y, Dayan P.** Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci* 8: 1704–1711, 2005.
- Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ.** Cortical substrates for exploratory decisions in humans. *Nature* 441: 876–879, 2006.
- Dayan P, Balleine BW.** Reward, motivation, and reinforcement learning. *Neuron* 36: 285–298, 2002.
- Dorris MC, Glimcher PW.** Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* 44: 365–378, 2004.
- Fishburn PC, Rubinstein A.** Time preference. *Int Econ Rev* 23: 677–694, 1982.
- Frederick S, Loewenstein G, O'Donoghue T.** Time discounting and time preference: a critical review. *J Econ Lit* 40: 351–401, 2002.
- Glimcher PW.** Choice: towards a standard back-pocket model. In: *Neuroeconomics: Decision Making and the Brain*, edited by Glimcher PW, Camerer CF, Fehr E, Poldrack RA. New York: Academic Press, 2009, p. 501–509.
- Glimcher PW, Kable J, Louie K.** Neuroeconomic studies of impulsivity: now or just as soon as possible? *Am Econ Rev* 97: 142–147, 2007.
- Green L, Fristoe N, Myerson J.** Temporal discounting and preference reversals in choice between delayed outcomes. *Psychon Bull Rev* 1: 383–389, 1994.
- Green L, Myerson J.** A discounting framework for choice with delayed and probabilistic rewards. *Psychol Bull* 130: 769–792, 2004.
- Green L, Myerson J, Macaux EW.** Temporal discounting when the choice is between two delayed rewards. *J Exp Psychol Learn Mem Cogn* 31: 1121–1133, 2005.
- Hampton AN, Bossaerts P, O'Doherty JP.** The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *J Neurosci* 26: 8360–8367, 2006.
- Hampton AN, Bossaerts P, O'Doherty JP.** Neural correlates of mentalizing-related computations during strategic interactions in humans. *Proc Natl Acad Sci USA* 105: 6741–6746, 2008.
- Hare TA, Camerer CF, Rangel A.** Self-control in decision-making involves modulation of the vmPFC valuation system. *Science* 324: 646–648, 2009.
- Hare TA, O'Doherty J, Camerer CF, Schultz W, Rangel A.** Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *J Neurosci* 28: 5623–5630, 2008.
- Holcomb JH, Nelson PS.** Another experimental look at individual time preference. *Rational Soc* 4: 199–220, 1992.
- Holt CA, Laury SK.** Risk aversion and incentive effects. *Am Econ Rev* 92: 1644–1655, 2002.
- Kable JW, Glimcher PW.** The neural correlates of subjective value during intertemporal choice. *Nat Neurosci* 10: 1625–1633, 2007.
- Kirby KN, Herrnstein RJ.** Preference reversals due to myopic discounting of delayed reward. *Psychol Sci* 6: 83–89, 1995.
- Laibson D.** Golden eggs and hyperbolic discounting. *Q J Econ* 112: 443–477, 1997.
- Lau B, Glimcher PW.** Value representations in the primate striatum during matching behavior. *Neuron* 58: 451–463, 2008.
- Mazur JE.** An adjusting procedure for studying delayed reinforcement. In: *Quantitative Analysis of Behavior: The Effects of Delay and Intervening Events on Reinforcement Value*, edited by Commons ML, Mazur JE, Nevin JA, Rachlin H. Hillsdale, NJ: Erlbaum, 1987, p. 55–73.
- McClure SM, Ericson KM, Laibson DI, Loewenstein G, Cohen JD.** Time discounting for primary rewards. *J Neurosci* 27: 5796–5804, 2007.
- McClure SM, Laibson DI, Loewenstein G, Cohen JD.** Separate neural systems value immediate and delayed monetary rewards. *Science* 306: 503–507, 2004a.
- McClure SM, Li J, Tomlin D, Cypert KS, Montague LM, Montague PR.** Neural correlates of behavioral preference for culturally familiar drinks. *Neuron* 44: 379–387, 2004b.
- McCoy AN, Platt ML.** Risk-sensitive neurons in macaque posterior cingulate cortex. *Nat Neurosci* 8: 1220–1227, 2005.
- Nieuwenhuis S, Heslenfeld DJ, von Geusau NJ, Mars RB, Holroyd CB, Yeung N.** Activity in human reward-sensitive brain areas is strongly context dependent. *NeuroImage* 25: 1302–1309, 2005.
- Ok EA, Masatlioglu Y.** *A General Theory of Time Preferences.* New York: New York Univ. Department of Economics, 2005.
- Padoa-Schioppa C, Assad JA.** Neurons in the orbitofrontal cortex encode economic value. *Nature* 441: 223–226, 2006.
- Padoa-Schioppa C, Assad JA.** The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nat Neurosci* 11: 95–102, 2008.
- Pine A, Seymour B, Roiser JP, Bossaerts P, Friston KJ, Curran HV, Dolan RJ.** Encoding of marginal utility across time in the human brain. *J Neurosci* 29: 9575–9581, 2009.
- Plassmann H, O'Doherty J, Rangel A.** Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *J Neurosci* 27: 9984–9988, 2007.
- Platt ML, Glimcher PW.** Neural correlates of decision variables in parietal cortex. *Nature* 400: 233–238, 1999.
- Rangel A, Camerer C, Montague PR.** A framework for studying the neurobiology of value-based decision making. *Nat Rev Neurosci* 9: 545–556, 2008.
- Read D.** Is time-discounting hyperbolic or subadditive? *J Risk Uncertainty* 23: 5–32, 2001.
- Read D, Roelofsma PHMP.** Subadditive versus hyperbolic discounting: a comparison of choice and matching. *Org Behav Hum Decision Process* 91: 140–153, 2003.
- Roelofsma PH, Read D.** Intransitive intertemporal choice. *J Behav Decision Making* 13: 161–177, 2000.
- Rubinstein A.** “Economics and psychology”? The case of hyperbolic discounting. *Int Econ Rev* 44: 1207–1216, 2003.
- Samejima K, Ueda Y, Doya K, Kimura M.** Representation of action-specific reward values in the striatum. *Science* 310: 1337–1340, 2005.
- Samuelson P.** A note on the measurement of utility. *Rev Econ Studies* 4: 155–161, 1937.
- Scholten M, Read D.** Discounting by intervals: a generalized model of intertemporal choice. *Manage Sci* 52: 1424–1436, 2006.
- Shadlen MN, Newsome WT.** Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 86: 1916–1936, 2001.
- Smith VL, Walker JM.** Monetary rewards and decision cost in experimental economics. *Econ Inquiry* 31: 245–261, 1993.
- Strotz RH.** Myopia and inconsistency in dynamic utility maximization. *Rev Econ Studies* 23: 165–180, 1956.
- Sugrue LP, Corrado GS, Newsome WT.** Matching behavior and the representation of value in the parietal cortex. *Science* 304: 1782–1787, 2004.
- Tanaka SC, Doya K, Okada G, Ueda K, Okamoto Y, Yamawaki S.** Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nat Neurosci* 7: 887–893, 2004.
- Tom SM, Fox CR, Trepel C, Poldrack RA.** The neural basis of loss aversion in decision-making under risk. *Science* 315: 515–518, 2007.