Supporting Information

Grabenhorst et al. 10.1073/pnas.1212706109

SI Methods

Electrophysiological Recording. A head holder and a recording chamber were fixed to the skull under general anesthesia and aseptic conditions. Before neuronal recordings, we located the amygdala from bone marks on coronal and sagittal radiographs taken with a guide cannula and electrode inserted at a known coordinate in reference to the stereotaxically implanted chamber. The anteroposterior position of the amygdala was between the sphenoid bone (rostral) and the posterior clinoid process at and above the dorsoventral position of the posterior clinoid (1). We recorded activity from single amygdala neurons from extracellular positions during task performance using standard electrophysiological techniques, including online visualization and threshold discrimination of neuronal impulses on oscilloscopes. We recorded from one neuron at a time; this record permitted varied exploratory tests during early experimental phases. We aimed to record representative neuronal samples from the dorsal, lateral, and basal amygdala.

We sampled activity from about 700 amygdala neurons in exploratory tests with the save–spend task. We recorded and saved the activity of neurons that seemed to respond to at least one task event during online inspection of several trials. This procedure resulted in a database of 329 neurons with task-related responses that we analyzed statistically.

After completion of data collection, recording sites were marked with small electrolytic lesions (15–20 μA × 20–60 s). The animals received an overdose of pentobarbital sodium (90 mg/kg i.v.) and were perfused with 4% paraformaldehyde in 0.1 M phosphate buffer through the left ventricle of the heart. Recording positions were reconstructed from 50-μm-thick stereotaxically oriented coronal brain sections stained with cresyl violet. The histological reconstructions also validated the previously radiographically assessed anatomical position of the amygdala in agreement with earlier reports (1–3). Fig. S7 shows the cresyl violet-stained brain sections from monkey A for the left amygdala. For Fig. 3*G*, we collapsed recording sites from both monkeys spanning 3 mm in the anterior–posterior dimension onto the same coronal section. Because our accuracy of recording site reconstructions was likely to be lower than 1-mm resolution, recording positions with respect to individual amygdala nuclei were approximated based on a stereotaxic atlas (4) and typical anatomical landmarks.

Free Choice Task. In each trial (Fig. 1*A*), the monkey chose between saving the reward that was available on that trial, thereby increasing its magnitude by a variable interest rate, and spending the previously accumulated reward for consumption on the present trial. A natural upper limit to the length of save choice sequences was given by the total amount of liquid that each monkey was able to drink on one trial. Animals initiated trials by placing their hands on an immobile, touch-sensitive key. The trial then started with an ocular fixation spot of 1.3° of visual angle at the center of the computer monitor. Animals were required to keep their gaze on the fixation spot at the stimulus center within 2–4°. At 1,500 ms plus mean of 500 ms (truncated exponential distribution) after fixation spot onset, the two save and spend visual stimuli of 7.0° appeared on the left and right sides of the computer monitor (pseudorandomized). In different blocks of typically 40–100 consecutive trials, different stimuli were used as save cues to indicate different interest rates. Animals indicated their choice with a saccade. The chosen stimulus was then replaced by a peripheral fixation spot of 7.0° of visual angle. The monkey could make its choice as soon as the visual cues appeared. After a delay period of 1,500 ms, a color change of the peripheral fixation spot served as a go signal for the monkey to release the touch key. The release of the touch key was followed by the delivery of the reinforcer (an auditory or visual cue on save trials vs. a drop of juice reward on spend trials). Failures of key touch or fixation breaks were considered errors and resulted in trial cancellation. More than three sequential errors led to a pause in behavioral testing. Accumulated saved rewards were retained across error trials.

To provide an example of how rewards were calculated, consider a series of two successive save choices by the monkey with a base rate of reward $b = 0.11$ and interest rate $q = 1.5$. On the second trial of the choice sequence, after the first save choice, reward $R = 0.11 \times (1 + 1.5) = 0.275$ mL is given. On the third trial, after two successive save choices, reward $R = 0.11 \times (1 + 1.5 + 1.5^2) = 0.523$ mL is given.

Each neuron was typically tested with one to two different interest rates. The duration required for testing neurons with statistically sufficient numbers of trials in both tasks usually precluded using more than two interest rates.

The decision task resembled tasks used to study intertemporal decision-making, because it required choosing between immediately available rewards and future rewards of different magnitudes. However, in contrast to standard intertemporal choice tasks, temporal delays in the present task were not imposed by the experimenter but chosen by the monkeys. Thus, monkeys were free to produce save choice sequences of different lengths, which were associated with different reward magnitudes depending on the interest rate. Furthermore, longer delays invariably involved higher numbers of behavioral reactions (ocular saccades).

The collection of reproducible electrophysiological data from many individual neurons required standardized testing during stable and reproducible behavioral performance. We trained each animal for 3–4 mo before neuronal recordings with the different visual stimuli and the different interest rates (300–400 trial/d, 5 d/wk). The animals were overtrained at the time of neuronal recordings and showed no behavioral signs of additional learning.

Control Task with Fixed Reward. To test whether the monkeys kept track of the amount of reward that they had accumulated through consecutive save choices, we offered them, on randomly interspersed trials, a choice between the accumulated reward and fixed amounts indicated by pretrained visual cues.

Rewards. A computer-controlled solenoid valve delivered juice reward from a spout in front of the animal's mouth (valve opening time of 100 ms, which corresponds to 0.38 mL). For monkey A, the base rate of reward magnitude, b from Eq. 1, was set to 0.11 mL for all sessions; for monkey B, the base rate was set to 0.11 mL for one-half of the sessions and 0.13 mL for one-half of the sessions. The animal's tongue interrupted an infrared light beam below the adequately positioned spout. An optosensor monitored licking behavior with 0.5-ms resolution (STM Sensor Technology), and the summed durations of beam interruptions during specific trials and task periods provided a measure of licking.

Saving Index. To examine the relationship between saving behavior and interest rate, we constructed a save choice index separately for each monkey and interest rate as follows. First, we calculated the average relative probability of observing a save choice sequence of a specific length, where sequence lengths

varied from zero consecutive save choices to the maximal sequence length observed for a monkey. (The maximal observed sequence length was effectively the upper limit of liquid that the animal could consume on one trial, as described in *Methods*.) These probabilities were defined to sum to 1.0 across sequence lengths within a given interest rate. Thus, these relative probabilities reflect the monkey's behavioral preference for a given sequence of save choices relative to all other possible sequences for a specific interest rate. Second, we weighted (multiplied) these relative probabilities with their associated sequence lengths (i.e., with the number of successive save choices for that sequence), thereby giving higher weight to probabilities that were associated with higher sequence lengths. Third, we calculated the mean over all weighted sequence lengths for a given interest rate. This mean defined the saving index for a given interest rate, and it is plotted in Fig. 1*C* for both monkeys and all interest rates. Thus (Eq. **S1**),

$$SI_q = \frac{1}{n}\sum_{i=1}^{n} P_{i,q} L_i, \qquad \textbf{[S1]}$$

where SI_q is the saving index for a given interest rate q, n is the maximal sequence length observed for the monkey, $P_{i,q}$ is the mean (relative) probability of observing a save choice sequence of a specific save sequence length i, and L_i is the number of successive save choices required to obtain sequence length i.

Logistic Regression of Save–Spend Choices on Differential Values. To analyze monkeys' saving behavior on a trial-by-trial basis, we used logistic regression analysis. First, we reasoned that, in analogy to decision-making in other economic tasks, choices would be guided by an internally computed decision variable that is based on the subjective values that monkeys' assigned to the different choice options. To construct a measure of subjective value in the save–spend task, we used the monkeys' relative probabilities of producing save sequences of specific lengths (as described in the preceding paragraph) and weighted (i.e., multiplied) them by the corresponding objective reward magnitude in milliliters that would be available from spending at that point of a save sequence. For both monkeys and all interest rates, this calculation produced a subjective weighting of the objective reward magnitude according to monkeys' behaviorally observed preferences for different save sequence lengths. Thus, the subjective value SV for spending at a given point i in a save sequence for a given interest q was defined as (Eq. **S2**)

$$SV_i = P_i M_i, \qquad \textbf{[S2]}$$

where P_i is the mean (relative) probability of observing a save choice sequence of a specific save sequence length i, and M_i is the objective reward magnitude in milliliters of juice at that point in the save sequence (the task description is discussed above). To obtain unbiased estimates of these subjective values that could be used as regressors for both behavioral choices and neural data, we used one-half of the behavioral data in each monkey (i.e., one-half of the experimental sessions for a given interest rate) to estimate the subjective values, and we used the other half for subsequent analysis. We then used these subjective values to construct a decision variable to model the monkeys' trial-by-trial choices.

The decision variable differential value plotted in Fig. 1*D* and Fig. S2 was constructed in analogy to decision variables commonly used in studies of intertemporal decision-making. It was defined as the difference between the subjective value for choosing to spend on the present trial and the mean subjective value for choosing to spend on any potential subsequent trial of the same save sequence (where the upper limit of potential future trials was given by the maximal observed sequence length for the monkey). For example, if the monkey was in the fourth trial of a save sequence (after having made three consecutive save choices), the differential value for the present trial would be calculated as the difference between the subjective value for spending on the fourth trial of a save sequence and the mean of the subjective values for spending on any of the potential next five trials (with nine consecutive save choices being the observed maximal number of consecutive save choices for the monkey). Thus, the differential subjective value DV on a given trial n for a given interest rate was calculated as (Eq. **S3**)

$$DV_n = SV_n - \frac{1}{m}\sum_{i=n+1}^{m} SV_i, \qquad \textbf{[S3]}$$

where SV_n is the subjective value of choosing to spend on trial n, and the term in the subtrahend reflects the average subjective value of choosing to spend on any of the potential subsequent trials i of the same save sequence, with m defining the upper limit of the save sequence (given by the maximal observed sequence length for the monkey). We used logistic regression analysis to model the monkeys' choices based on this decision variable. The dependent variable was a binary indicator function denoting whether the monkey made a save or spend choice on a given trial. The independent variable was the differential value for the corresponding trial as defined above. The main purpose of this analysis was to test whether the differential value provided an adequate approximation of the decision variables that guided the monkeys' choices to inform our analysis of the neuronal data. The results of this analysis are summarized in Fig. 1*D*, Fig. S2, and Table S1. We found that this differential value (i.e., a decision variable that incorporated the average subjective value of potential future trials in a sequence) provided a better fit than a comparable decision variable that incorporated only the subjective value of choosing to spend on the next trial. Moreover, a differential value based on subjective values provided a better fit compared with decision variables based only on objective reward magnitudes or choice probabilities.

Data Analysis. We counted neuronal impulses in each neuron on correct trials relative to different task events with time windows that were fixed across all neurons: 1,000 ms before fixation spot (PreFP), 1,775 ms after fixation spot but before cues (FP; starting 25 ms after fixation spot onset), 300 ms after cues (Cue; starting 20 ms after cue onset), and 500 ms during the reward/outcome period of the preceding trial (Out-1; starting 50 ms after reward onset). We first identified task-related responses by comparing activity in the FP, Cue, and Out-1 periods with a control period (PreFP) using the Wilcoxon test ($P < 0.05$). Because the PreFP period served as the control period, we did not select for task-relatedness in this period and included all neurons with observed impulses in the analysis. We then used the following multiple regression model to assess relationships to trial-by-trial save–spend choices, different measures of reward value, left–right actions, left–right cue positions, and saccadic reaction times ($P < 0.05$) (Eq. **S4**):

$$Y = \beta_0 + \beta_1 SS + \beta_2 RM + \beta_3 DV + \beta_4 LR + \beta_5 SL + \beta_6 RT + \varepsilon, \qquad \textbf{[S4]}$$

with SS as the save vs. spend choice, RM as the sum of objective reward magnitudes available for save and spend choices, DV as the subjective differential value used for behavioral modeling, LR as left vs. right action, SL as spatial cue position (save cue left vs. right), and RT as saccadic reaction time; β_{1-6} are corresponding regression coefficients, β_0 is the intercept, and ε is error. On average, intercorrelations between these regressors

were low (Table S2). Because reward magnitudes for both save and spend choices increased monotonically over save trials, different potential reward magnitude regressors (e.g., sum of reward magnitudes and chosen/not chosen reward magnitude) were highly correlated and produced similar results. The same model was used to analyze responses in the imperative task.

Standardized regression coefficients (β values) in Figs. 1*D* and 3*E* were defined as x_i (s_i/s_y); x_i is the raw slope regression coefficient for regressor i, and s_i and s_y are the SDs of independent variable i and the dependent variable, respectively (5).

Decoding of Choices from Neuronal Data. Here, we provide more details about the decoding analysis. We used a biologically plausible classifier to decode choices from neuronal data on a trial-by-trial basis (6). The decoding procedure used by the classifier was based on a nearest neighbor algorithm. The neuronal activity measured in impulses per second on a single trial in an individual neuron was used as input to the classifier. For each individual neuron, every trial was represented in the space spanned by the distribution of its impulse rates on save and spend choice trials and decoded by assigning it to the class of its nearest neighbor using the Euclidean distance (6). This type of classification is biologically plausible in that a real downstream neuron could perform the classification in a similar way by comparing the input on a given trial with a stored vector of synaptic weights (7). We used a leave-one-out cross-validation procedure, in which every trial was decoded based on the distribution of impulse rates from all other trials. To investigate population coding of choices, we considered neurons as simultaneously recorded in the sense that the trial-specific responses of all neurons were grouped together and that decoding proceeded using the cross-validation procedure just described (6). Classification performance was measured as the percentage of correctly decoded individual trials, which we averaged across responses. We repeated the analysis using linear discriminant analysis, which also used leave-one out cross-validation procedures. To produce the graphs in Fig. 3*F* and Fig. S5, we randomly selected a given number of responses at each step and then determined the percentage correct. For each step, this procedure was repeated 10 times. We used a permutation test with 1,000 iterations to define statistical significance of the classification. Statistical significance was defined as the probability that the observed percentage correct was below a given percentile of the probability distribution of classification results based on randomly shuffled data.

Sliding Window Regression Analysis. We used sliding window multiple regression analysis (using the regression model described above) with a 200-ms window, and then, we moved the window in steps of 25 ms across each trial. Coefficients of partial determination (CPDs) (5) were defined as (Eq. **S5**),

$$\mathrm{CPD}(\mathrm{X_i}) = [\mathrm{SSE}(\mathrm{X_{-i}}) - \mathrm{SSE}(\mathrm{X_{-i}}, \mathrm{X_i})]/\mathrm{SSE}(\mathrm{X_{-i}}), \qquad \textbf{[S5]}$$

with SSE(X) indicating the sum of squared errors in a regression model that includes a set of regressors X, and $\mathrm{X_{-i}}$ indicating the set of regressors that includes all regressors except $\mathrm{X_i}$. Using methods from previous studies (5), latencies of choice and value coding were defined as the first window in which a CPD was 3 SDs above the mean CPD obtained from a permutation test (1,000 iterations) for three consecutive steps.

SI Results

Licking Durations. We measured anticipatory licking durations on save and spend trials before the cues were presented (Fig. S3). In the free choice task, significant differences in licking durations between save and spend trials would likely indicate a difference in reward expectation between these trials, because immediate rewards were only delivered if the monkeys chose to spend. Indeed, for both monkeys, licking times were significantly different between save and spend trials in the free choice task, despite individual differences between animals (both $P < 0.001$, Mann–Whitney test). We also examined licking durations in the imperative control task before cue appearance on every trial. If licking durations in the imperative task also differed between save and spend trials, even before cue appearance, this result might indicate that the monkeys anticipated these trial types, similar to the free choice task. Indeed, for both monkeys, licking times in the imperative task were significantly different between save and spend trials (both $P < 0.001$, Mann–Whitney test). No significant differences in licking patterns were found between the free choice and imperative tasks. Fig. S3 shows this pattern of licking durations for monkey A.

Reaction Times. We analyzed the reaction times of the saccades with which monkeys indicated their choices on save and spend trials. This analysis helped to test whether choice-differential neuronal activity could be explained by task difficulty as measured with reaction times (Fig. S3). For both monkeys, saccadic reaction times were longer on save compared with spend trials in both the free choice task ($P < 0.001$, Mann–Whitney test) and the imperative control task ($P < 0.001$, Mann–Whitney test). If choice-differential neuronal activity reflected differences in task difficulty between save and spend trials [or any secondary variable resulting from differences in task difficulty (for example, differential attention or arousal levels) on save vs. spend trials], then neuronal activity should differ between save and spend trials on both the free choice task and the imperative task. By contrast, as reported in the text, neuronal responses showed differences between these trial types only in the free choice task and not in the imperative task. These observations would argue against an explanation of our effects in terms of task difficulty.

Comparison of Behavioral Models. We evaluated whether monkeys' choices were better explained by the differential subjective value model compared with a simpler model that only incorporated the monkeys' average choice probabilities for different save sequences. If choices were explained by a model based solely on the monkeys' save sequence distributions, this finding might suggest that the animals developed a simple counting strategy and did not incorporate trial-by-trial changes in differential subjective value. To test this directly, we compared a logistic regression model that incorporated only the monkey's choice probabilities with one that also incorporated differential value as a covariate. The models were fit separately for different interest rates and the two animals. Across animals and interest rates, the differential value regressor remained significant ($P < 0.002$ in all cases), even if choice probabilities were included as an additional regressor. The same result was obtained if differential value was orthogonalized with respect to choice probabilities (i.e., the shared variance between regressors was assigned to choice probabilities) using Gram–Schmidt orthogonalization (8). This result suggested that the differential value regressor explained a significant proportion of variance not accounted for by simple choice probabilities. Indeed, a direct comparison of standardized regression coefficients showed significantly higher coefficients for differential value compared with choice probability across monkeys and interest rates ($P = 0.001$, paired t test). To test whether the differential value model provided a better fit, even if the number of model parameters was taken into account, we used the Akaike information criterion (AIC) and Bayesian information criterion (BIC), which penalize models with higher numbers of free parameters (9). AIC is defined as $-2 \ln \mathrm{L} + 2 \mathrm{k}$, in which L is the likelihood of the model, and k is the number of model parameters. BIC is defined as $-2 \ln \mathrm{L} + \mathrm{k} \ln \mathrm{N}$, in which N is the number of observations. Across animals and interest rates, AIC and BIC comparisons consistently fa-

vored the differential value model over the simpler choice probability model. Together, these results indicate that monkeys' choices were better explained by the differential value model, and thus, they were influenced by trial-by-trial variations in differential subjective value.

Anatomical Location of Recorded Neurons. Of all 94 choice-selective neurons, 45 neurons were recorded in the dorsal amygdala (Fig. 3*G*) (including central and medial nuclei; of 168 neurons in total recorded in this area), 3 neurons were recorded in the lateral amygdala (of 23 recorded neurons in this area), 20 neurons were recorded in the basomedial amygdala (of 67 recorded neurons in this area), 16 neurons were recorded in the dorsal basolateral amygdala (of 47 recorded neurons in this area), and 10 neurons were recorded in the basoventral amygdala (of 24 recorded neurons in this area). No systematic differences were found between recording sites (nonsignificant χ^2 test).

Reward Expectation. To test whether choice-predictive activity could be explained by differences in reward expectation between save and spend trials, we tested specifically for changes in neuronal activity across successive save trials. Previous studies have shown that reward expectancy-coding neurons show activity changes across trials to reward receipt (10). By contrast, many of our choice-predictive responses showed no significant value coefficients, which modeled systematic changes over successive save trials [85 of 127 choice-predictive responses (67%); 37 of 57 with choice-predictive activity in the free choice but not imperative task (65%)]. Moreover, all of our choice-predictive responses showed a significant choice regressor, although several measures of value were included as covariates in the multiple regression model. This finding suggested that choice coding in these responses cannot be explained in terms of reward value coding or related expectancy. Furthermore, the imperative control task served to explicitly control for simple reward expectation effects. In both tasks, reward expectation between save and spend trials was similar, which was indicated by patterns of differences in licking durations and saccadic reaction times (Fig. S3). By contrast, 80% of choice-predictive responses tested in both tasks failed to show significant choice coefficients in the imperative task, suggesting that choice coding in these neurons is unlikely caused by simple reward expectation.

Neuronal Coding of Value, Action, Visual Cue Features, and Reaction Times. In addition to choice coding, we confirmed the known value (reward magnitude and differential value) coding in the amygdala (3, 11–14) in 225 of 846 task-related responses (27%). Furthermore, 32 of 169 task-related responses in the Cue period (19%) were modulated by the spatial arrangement of the cues, consistent with known visual feature responses in the amygdala (11). Few neurons in the cue period showed a significant regression coefficient for left/right eye movements (15 neurons; 9%) or saccadic reaction times (10 neurons; 6%). In all other task periods, less than 5% of responses (our significance threshold) were modulated by eye movement direction, spatial cue position, or reaction time.

Control Analyses Testing for Different Forms of Value Coding. The main analysis reported in the paper identified 127 neuronal responses with significant choice coefficients. A subgroup of these responses coded choice without also coding value, which was indicated by nonsignificant coefficients for reward magnitude and differential value (Table 1). To further examine whether significant choice coefficients might be explained by other forms of value coding, we performed supplementary analyses. Previous studies found that neurons in the orbitofrontal cortex and striatum code different types of value signals (15–19), including the value of specific choice options, irrespective of whether the option is chosen (offer value or action value signals), the value of the chosen option, irrespective of its identity (chosen value signals), and the value of a specific choice option if that option is chosen (subtype of chosen value signals). We tested whether the choice-predictive responses described in the present study can be explained in terms of such value signals. We used the following three supplementary regression models. In (Eq. **S6**)

$$\text{Model S1: } Y = \beta_0 + \beta_1 SS + \beta_2 RM + \beta_3 SVspend + \beta_4 SVsave + \beta_5 LR + \beta_6 SL + \beta_7 RT + \varepsilon, \quad \textbf{[S6]}$$

SVspend is the subjective value of the spend choice option (irrespective of whether it is chosen), and *SVsave* is the subjective value of the save choice option (irrespective of whether it is chosen); all other regressors are defined as in our main regression model. This model, thus, tests for coding of offer value. In (Eq. **S7**)

$$\text{Model S2: } Y = \beta_0 + \beta_1 SS + \beta_2 RM + \beta_3 CV + \beta_4 UCV + \beta_5 LR + \beta_6 SL + \beta_7 RT + \varepsilon, \quad \textbf{[S7]}$$

CV is the subjective value of the chosen option (regardless of whether it is a save or spend choice), and *UCV* is the subjective value of the not chosen option. This model, thus, tests for coding of chosen value. In (Eq. **S8**)

$$\text{Model S3: } Y = \beta_0 + \beta_1 SS + \beta_2 RM + \beta_3 CVspend + \beta_4 CVsave + \beta_5 LR + \beta_6 SL + \beta_7 RT + \varepsilon, \quad \textbf{[S8]}$$

CVspend is the subjective value of the spend choice option only if it is chosen (taking a value of zero if it is not chosen), and *CVsave* is the subjective value of the save choice option only if it is chosen (taking a value of zero if it is not chosen). This model, thus, tests for coding of a subtype of chosen value, which combines information about the chosen value with information about the identity of the chosen option.

Model S3 is of particular interest, because coding of the chosen value for a specific option could appear very similar to coding of the categorical choice. Thus, if neuronal responses showed a significant choice regressor in this model, despite the inclusion of the chosen value covariates, this result would support our conclusion of choice coding beyond value. Using these additional regression models to estimate coefficients for choice signals simultaneously with different value signals, we found that the majority of our choice-predictive responses was not accounted for by value coding: Model S1 resulted in 110 responses (84 neurons) with significant choice coefficients (compared with 127 such responses in 94 neurons with our main regression model), model S2 resulted in 101 responses (77 neurons) with significant choice coefficients, and model 3 resulted in 106 responses (88 neurons) with significant choice coefficients. The percentages of neurons with significant choice coefficients but nonsignificant value coefficients were 56%, 68%, and 52% for models S1–S3, respectively (compared with 63% with our main regression model).

Thus, the majority of choice-predictive responses found in amygdala neurons could not be explained in terms of different types of value coding and rather, seemed to reflect the monkeys' categorical choices. We acknowledge that some choice-predictive responses could be interpreted as special types of value coding; indeed, this result may be expected from our main analysis, because some of the choice-predictive responses identified with our main regression model had both significant choice and value coefficients (Table 1).

1. Aggleton JP, Passingham RE (1981) Stereotaxic surgery under X-ray guidance in the rhesus monkey, with special reference to the amygdala. *Exp Brain Res* 44(3):271–276.
2. Bermudez MA, Schultz W (2010) Reward magnitude coding in primate amygdala neurons. *J Neurophysiol* 104(6):3424–3432.
3. Bermudez MA, Schultz W (2010) Responses of amygdala neurons to positive reward-predicting stimuli depend on background reward (contingency) rather than stimulus-reward pairing (contiguity). *J Neurophysiol* 103(3):1158–1170.
4. Paxinos G, Huang X-F, Toga AW (2000) *The Rhesus Monkey Brain in Stereotaxic Coordinates* (Academic, San Diego).
5. Cai X, Kim S, Lee D (2011) Heterogeneous coding of temporally discounted values in the dorsal and ventral striatum during intertemporal choice. *Neuron* 69(1):170–182.
6. Quian Quiroga R, Snyder LH, Batista AP, Cui H, Andersen RA (2006) Movement intention is better predicted than attention in the posterior parietal cortex. *J Neurosci* 26(13):3615–3620.
7. Rolls ET, Treves A (1998) *Neural Networks and Brain Function* (Oxford Univ Press, Oxford).
8. Draper NR, Smith H (1998) *Applied Regression Analysis* (Wiley Interscience, New York), 3rd Ed.
9. Kutner MH, Nachtsheim CJ, Neter J, W L (2004) *Applied Linear Statistical Models* (McGraw-Hill, New York).
10. Sugase-Miyamoto Y, Richmond BJ (2005) Neuronal signals in the monkey basolateral amygdala during reward schedules. *J Neurosci* 25(48):11071–11083.
11. Paton JJ, Belova MA, Morrison SE, Salzman CD (2006) The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* 439(7078):865–870.
12. Rolls ET (2000) Neurophysiology and functions of the primate amygdala, and the neural basis of emotion. *The Amygdala: A Functional Analysis*, ed Aggleton JP (Oxford Univ Press, Oxford), 2nd Ed, pp 447–478.
13. Schoenbaum G, Chiba AA, Gallagher M (1999) Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. *J Neurosci* 19(5):1876–1884.
14. Tye KM, Stuber GD, de Ridder B, Bonci A, Janak PH (2008) Rapid strengthening of thalamo-amygdala synapses mediates cue-reward learning. *Nature* 453(7199):1253–1257.
15. Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. *Nature* 441(7090):223–226.
16. Kim S, Hwang J, Lee D (2008) Prefrontal coding of temporally discounted values during intertemporal choice. *Neuron* 59(1):161–172.
17. Lau B, Glimcher PW (2008) Value representations in the primate striatum during matching behavior. *Neuron* 58(3):451–463.
18. Roesch MR, Singh T, Brown PL, Mullins SE, Schoenbaum G (2009) Ventral striatal neurons encode the value of the chosen action in rats deciding between differently delayed or sized rewards. *J Neurosci* 29(42):13365–13376.
19. Samejima K, Ueda Y, Doya K, Kimura M (2005) Representation of action-specific reward values in the striatum. *Science* 310(5752):1337–1340.

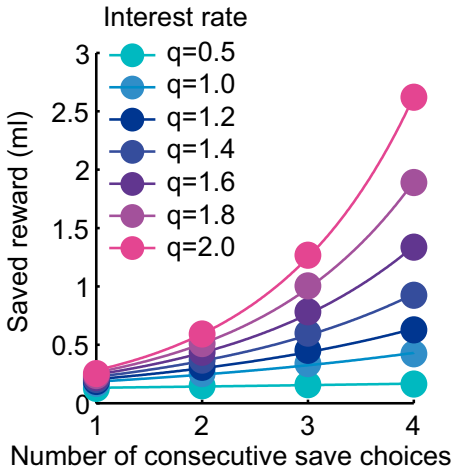

Fig. S1. Increase of reward magnitude for consecutive save choices as a function of interest rate *q*. (In this example, reward magnitude was calculated for a base volume of 0.09 mL.)

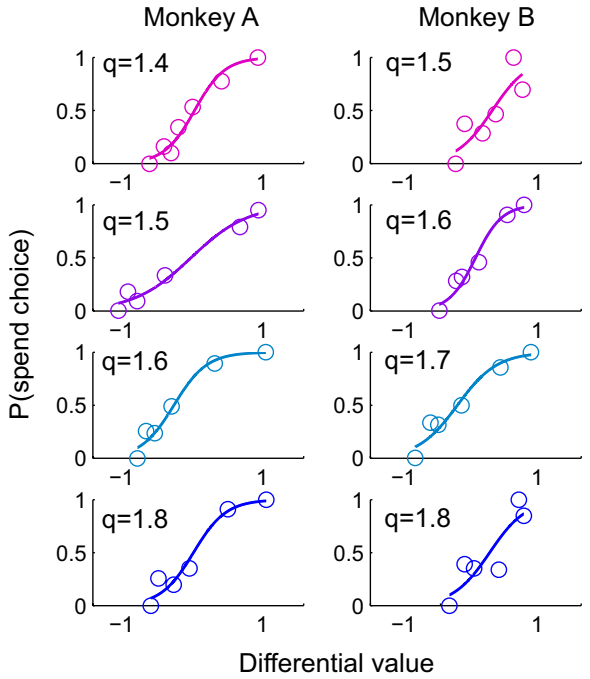

Fig. S2. Spend choice probability as a function of differential subjective value for save and spend choice options shown for different interest rates *q* (curves represent logistic fits to choice data; all logistic regressions were significant at $P < 0.001$).

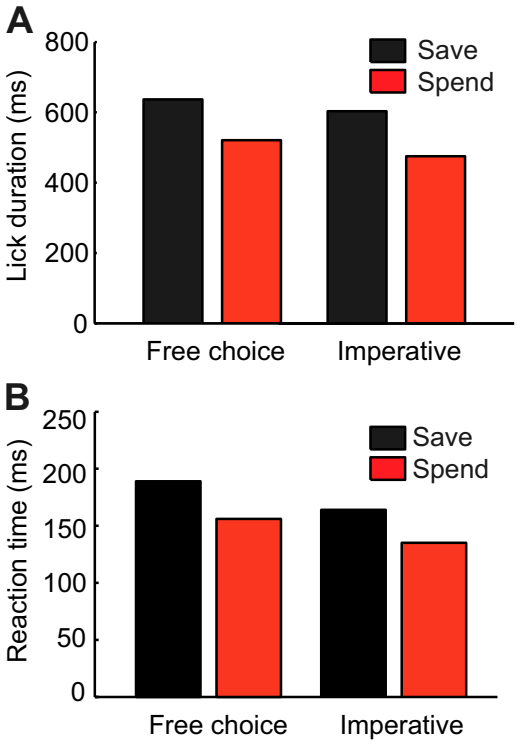

Fig. S3. Licking durations and saccadic reaction times. (*A*) Medians of anticipatory licking durations in the free choice and imperative control tasks measured before cue appearance. Significant differences were found between save and spend trials for both tasks ($P < 0.001$, Mann–Whitney test). (*B*) Medians of saccadic reaction times in the free choice and imperative control tasks. Significant differences were found between save and spend trials for both tasks ($P < 0.001$, Mann–Whitney test). Similar results were found for both monkeys.

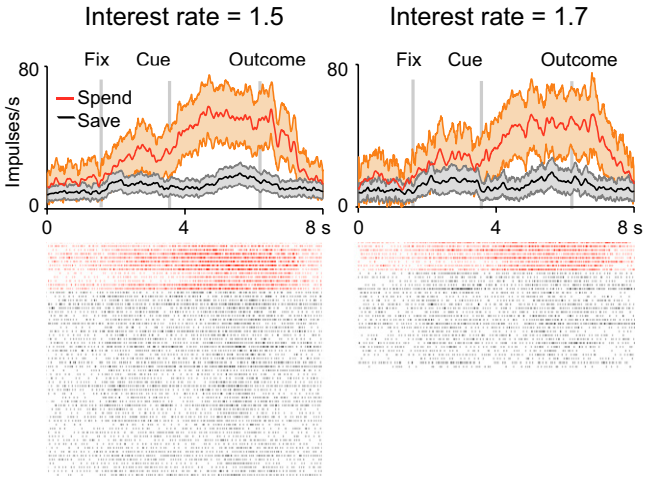

Fig. S4. Choice-predictive responses irrespective of interest rate. (*Left*) Example neuron showing a significant choice regression coefficient when tested with an interest rate of q = 1.5. (*Right*) The save–spend difference (significant choice regression coefficient) persisted when the same neuron was tested with a different interest rate of q =1.7.

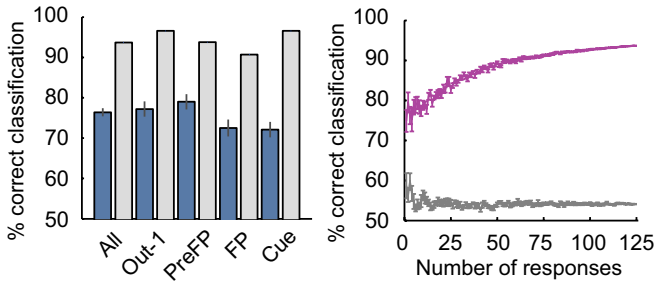

Fig. S5. (*Left*) Classification accuracy (percent correct classification ± SEM) of a biologically plausible classifier (1) decoding save–spend choices from neuronal impulse rates using all 127 choice-predictive responses. Blue bars, average accuracy for classification based on individual responses; gray bars, accuracy when data were combined across responses. Task periods: Out-1 (n = 19), PreFP (n = 27), FP (n = 38), Cue (n = 43), all periods (All; n = 127). Chance performance is 50%. (*Right*) Increase in classification accuracy (mean ± SEM) as a function of the number of responses from which data were combined (n = 127). Gray trace, classification based on randomly permuted data.

1. Quian Quiroga R, Snyder LH, Batista AP, Cui H, Andersen RA (2006) Movement intention is better predicted than attention in the posterior parietal cortex. *J Neurosci* 26(13):3615–3620.

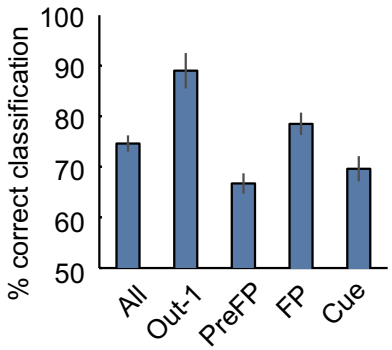

Fig. S6. Classification accuracy (percent correct classification ± SEM) of a classifier based on linear discriminant analysis using 57 choice-predictive responses. Task periods: Out-1 (n = 8), PreFP (n = 16), FP (n = 20), Cue (n = 13), All (n = 57). Chance performance is 50%.

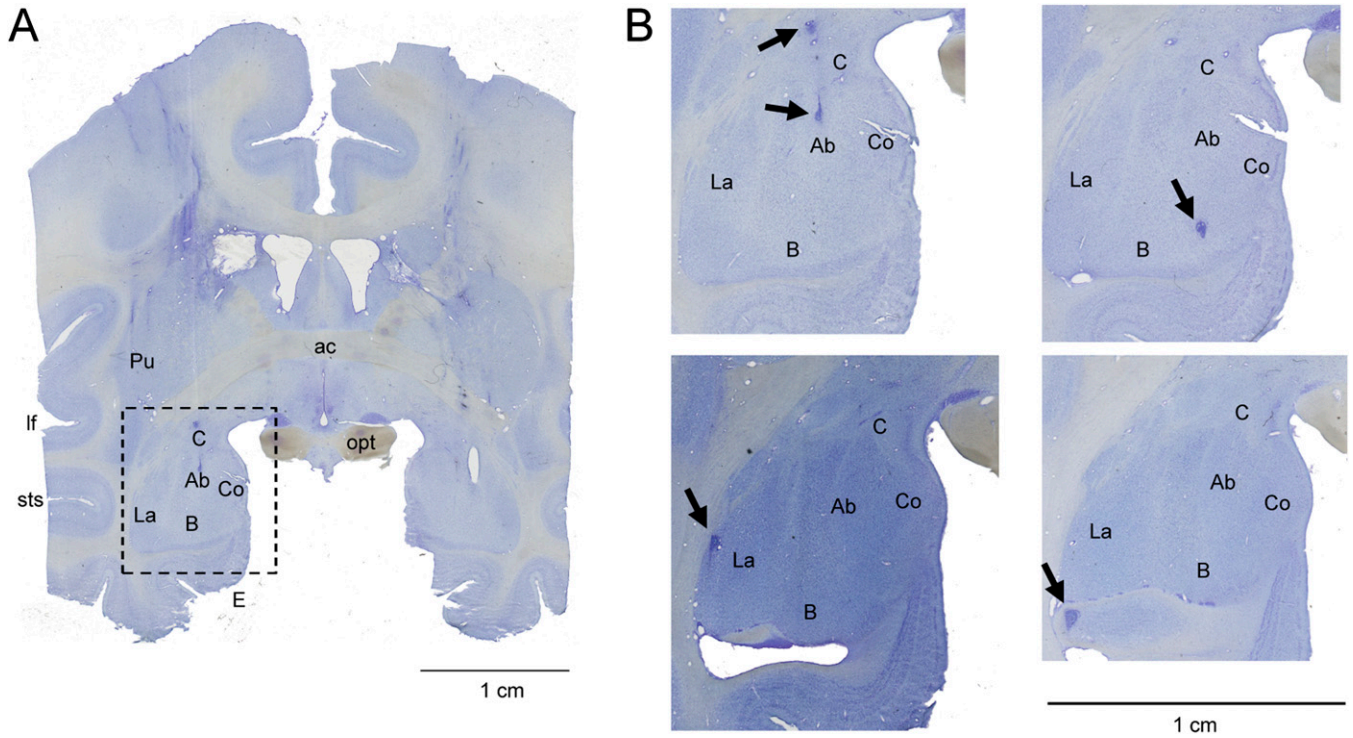

Fig. S7. Histological reconstruction of recording sites. (*A*) Photomicrograph of a cresyl violet-stained coronal section showing the amygdala and surrounding structures in monkey A. The dashed box indicates the location of the region magnified in *B*, *Upper Left*. (*B*) Arrows mark electrolytic lesions made after recordings. Lesions were placed to indicate typical recording sites in the amygdala (*Upper*) and the estimated boundaries of the amygdala (*Lower*). Location of individual nuclei were approximated based on a stereotaxic atlas (1) and typical anatomical landmarks. Ab, accessory basal nucleus; ac, anterior commissure, B, basal nucleus; C, central nucleus; Co, cortical nucleus; E, entorhinal cortex; lf, lateral fissure; La, lateral nucleus; opt, optic tract; Pu, putamen; sts: superior temporal sulcus.

1. Paxinos G, Huang X-F, Toga AW (2000) *The Rhesus Monkey Brain in Stereotaxic Coordinates* (Academic, San Diego).

Table S1. Logistic regression of save–spend choices on differential value

Regressor	Monkey A (differential value)	Monkey B (differential value)
β (SE)	1.65 (0.09)*	1.12 (0.14)*
N (trials)	3,859	2,939
χ^2	850.42†	476.84†
−2 log likelihood	3,266.33	3,086.82
Cox and Snell R^2	0.20	0.19
Nagelkerke R^2	0.30	0.21
Percent correct	83.4	76.0

χ^2 denotes result of omnibus test for significance of model coefficients.
*Significant at $P < 1 \times 10^{-8}$.
†Significant at $P < 0.001$.

Table S2. Correlations between regressors for the main multiple regression model averaged across all sessions

	SS	RM	DV	LR	SL	RT
SS	1					
RM	0.24	1				
DV	0.28	0.09	1			
LR	−0.02	−0.02	−0.03	1		
SL	0.02	−0.03	−0.03	−0.44	1	
RT	−0.24	−0.06	−0.1	0.11	0.02	1

DV, differential subjective value; LR, left–right action choice; SL, save cue left–right position; RM, reward magnitude; RT, saccadic reaction time; SS, save–spend choice.