Research report

Sustained neural activity associated with cognitive control during temporally extended decision making

Tal Yarkoni[a], Jeremy R. Gray[b], Elizabeth R. Chrastil[a], Deanna M. Barch[a],
Leonard Green[a], Todd S. Braver[a],*

[a]Department of Psychology, Washington University, Campus Box 1125, St. Louis, MO 63130, USA
[b]Yale University, New Haven, CT 06520, USA

Abstract

Decision making has both cognitive and affective components, but previous neuroimaging studies in this domain predominantly have focused on affect and reward. The current study examined a decision-making paradigm that placed strong demands on cognitive control processes by making reward payoffs contingent upon decision-making history. Payoffs were maximized by choosing the option that, paradoxically, was associated with a lower payoff on the immediate trial. Temporal integration requirements were manipulated by varying, across conditions, the window of previous trials over which the reward function was calculated. The cognitive demands of the task were hypothesized to engage neural systems responsible for integrating and actively maintaining actions and outcomes over time and the top-down biasing of response selection. Brain activation was monitored with functional magnetic resonance imaging (fMRI) using a mixed-blocked and event-related design to extract both transient and sustained neural responses. A network of brain regions commonly associated with cognitive control functions, including bilateral prefrontal cortex (PFC), bilateral parietal cortex, and medial frontal cortex, showed selectively sustained activation during the task. Increasing temporal integration demands led to a shift from transient to sustained activity in additional regions, including right hemisphere dorsolateral and frontopolar PFC. These results demonstrate the contribution of cognitive control mechanisms to temporally extended decision-making paradigms and highlight the benefits of decomposing activation responses into sustained and transient components.
© 2005 Elsevier B.V. All rights reserved.

1. Introduction

Decision making is a fundamental aspect of human life, involved in everything from choosing what clothes to wear to choosing who to vote for in an election. Psychologically, decision making is intriguing in that it depends on both cognitive and affective factors [14]. Neuroimaging approaches to the study of human decision making are increasingly common and have led to substantial gains in our understanding of the neural mechanisms underlying decision-making processes.

The majority of research in this area has focused on affective or motivational components of decision making. The neural mechanisms of reward have been widely studied in recent years, and studies have demonstrated multiple dissociations in the neural systems underlying different aspects of reward (for a review, see Refs. [6,31,40,53]). For example, different mechanisms have been shown to underlie

---

* Corresponding author.
  E-mail address: tbraver@artsci.wustl.edu (T.S. Braver).

the anticipation of and response to rewards [19,33,43], or the learning [37,52] versus storage [35,42] of stimulus-incentive associations. Such work has contributed in part to the development of theories relating reward and punishment signals, or 'somatic markers,' to the ongoing modification of decision-making behavior on a day-to-day basis [4,5]. In recent years, neuroimaging techniques also have been used to study the influences of reward and affective mechanisms on people's decision-making behavior in economic situations [15,31,47,51].

In contrast to the focus on affect and reward, relatively little attention has been paid to the role of higher cognitive faculties in decision-making behavior. Nevertheless, many decision-making situations require a high degree of self-control. Although self-control is critical for modulating reward-seeking behavior [12,29], neither the neurobiological implementation of such control mechanisms nor the nature of their influence on decision making is well understood. The canonical form of self-control required in such cases is the ability to inhibit pursuit of an immediate reward that is associated with larger long-term costs. Addiction represents a paradigmatic example: the systematic bias drug addicts exhibit towards short-term pleasure at the expense of long-term well-being is often attributed at least in part to a dysfunction of impulse control, that is, of cognitive mechanisms normally responsible for inhibiting pursuit of immediately available rewards when there are unpleasant consequences attached to them [30,58]. Individuals with intact control mechanisms presumably experience some of the same impulses as addiction-prone individuals but are better able to constrain their choice of the immediate pleasure. Cognitive mechanisms that subserve both inhibitory processing and the abstract internal representation of future states may be critical for self-controlled decision-making behavior. There is a growing experimental literature specifically focusing on self-control in decision making [25,38,45], which use a variety of laboratory tasks such as the temporal discounting paradigm [26].

A second factor that suggests the importance of cognitive control mechanisms in decision making is that in many real-world decision-making situations the environment is dynamic and non-stationary. Thus, a given choice may produce a different outcome than it had on a previous occasion because the environment is in a different state. The incentive value of a particular stimulus often is not constant through time, but instead depends on the organism's decision-making history. Foraging for prey provides a good example of a non-stationary sampling process. A predator remains in the same territory so long as there is available prey. As the supply begins to diminish, the incentive to move to a different area increases; however, the predator's own behavior contributes in part to the change in environment. When prey is plentiful, the incremental cost to the predator of eating another rabbit is very small; but the more rabbits eaten, the harder it becomes to find the next one and hence the associated cost increases. Similarly, the first alcoholic drink

that a person consumes at a party typically is rewarding and has little or no associated costs. Each subsequent drink, however, decreases the associated reward while increasing its cost. In such situations, optimal decision-making behavior requires one to detect such changes in the environment (or one's internal state) and dynamically update and flexibly re-adjust decision-making biases accordingly. Such strategic flexibility is a hallmark of cognitive control and depends upon mechanisms that can continuously monitor performance in an on-line manner and use this information to dynamically regulate goal information.

A third aspect to decision making that suggests the involvement of cognitive control processes is that decision-making situations often require an adaptive allocation of choices over time. This is particularly true in tasks involving a non-stationary environment, such as those just discussed. Under these situations, both past choices and their outcomes must be tracked and integrated over time, such that valid predictions can be generated regarding the expected current and future outcomes associated with each choice. This temporal integration component of decision making describes social–economic games, such as matching pennies, rock-paper-scissors, and iterated Prisoner's Dilemma games, in which players update their choice strategy based on integrating over past experiences with an opponent. Thus, one function of cognitive control mechanisms in dynamic decision making may be to enable temporally integrated information about past choices and their outcomes to be translated into predictions regarding the consequences of current choices, which can then serve as biases on action selection [2].

Based on these considerations, the neural mechanisms likely to be associated with self-controlled behavior, dynamic updating of strategy, and temporal integration during decision making should be the same ones associated with cognitive control in other task domains that require such control mechanisms. It is fairly well established that the exertion of cognitive control across a wide variety of tasks engages a reliable network of brain regions including lateral prefrontal cortex (PFC), medial frontal cortex (including the anterior cingulate cortex), and posterior parietal cortex. Together, these brain regions are thought to subserve a variety of control functions, including inhibition, active maintenance, and updating or manipulation of information stored on-line (i.e., in working memory Refs. [8,11,54]). Interestingly, this set of brain regions is quite distinct from those typically highlighted in past lesion and neuroimaging studies of decision making. Such studies have tended to focus on regions more closely associated with affective or reward processing: ventral striatum, amygdala, and ventromedial or orbital PFC [3,27,32,48]. What accounts for the discrepancy between our conceptual analysis and the accumulated neuroimaging results? We suggest that in order to observe decision-making activity in brain regions subserving control functions it is necessary to increase the demands placed on temporal

integration processes during task performance. Few neuro-imaging studies have studied decision-making tasks that involve non-stationary environments and temporal integration requirements (but see Ref. [40]), and none to our knowledge explicitly has manipulated these components of the task situation.

One potential reason for the relative lack of attention to the influence of neural mechanisms of temporal integration and cognitive control in decision making is that such mechanisms may operate at time scales that are not always easy to address with existing neuroimaging techniques. The development of rapid event-related fMRI [10,50] enabled researchers to analyze the neural response to decision-making events on a trial-by-trial basis, providing a powerful tool for looking at transient neural responses tied to particular trials. This approach has proven instrumental in understanding neural mechanisms of reward, which can be studied at short time frames (e.g., an anticipatory period of several seconds, followed by a momentary response to a trial outcome). However, its utility is limited in cases where the neural processes of interest are postulated to span multiple trials or entire task blocks. Older blocked designs provide greater power to analyze temporally extended activation, but the interpretation of results in such designs is problematic because of the difficulty in determining what aspect of the task is associated with the neural activation in specific regions. Thus, investigations focusing on transient, trial-specific processes underlying behavior greatly out-number those focused on sustained neural processes. Yet information about sustained activation is potentially critical to an understanding of decision-making situations that impose a need for temporal integration, because such integration likely requires active maintenance of information across multiple trials. Similarly, neural representations of mental strategies are liable to be reflected more strongly in sustained activity than in transient, event-related activity since participants may carry over the same strategy (e.g., "always select from the left deck of cards") from trial to trial but employ different strategies in different task blocks or different conditions.

Mixed-design approaches to fMRI data recently have been developed that allow neural activity during a task to be decomposed into sustained and transient components [16,57]. In many ways these designs optimize the various trade-offs between block and event-related designs. They are particularly useful when there is reason to suspect that the neural activity associated with a particular task can be attributed to separate sustained and transient components (e.g., Ref. [9]). Most decision-making paradigms employed in neuroimaging studies appear to meet this criterion. There is a clear conceptual difference between neural processes that take place during individual decision-making trials (e.g., selection for action, anticipation, and experience of outcome) and those that must be maintained across many or all trials (maintenance of task instructions, evaluation of multiple outcomes, differing response strategies). It also is

clear that these aspects interact such that changes in strategy affect changes in behavior on individual trials, and vice versa, providing further motivation for such a decomposition. However, owing perhaps to their relatively recent adoption, mixed designs have not yet been applied to the domain of decision making.

The aim of the present paper is to address the two lacunae in the neuroimaging decision-making literature discussed above: (1) the relative absence of neuroimaging studies addressing cognitive influences on decision-making behavior, and (2) the interplay between sustained and transient task-related neural activity. To achieve these goals, we studied a dynamic decision-making task in which the demands for temporal integration of reward information could be manipulated experimentally. We coupled this with a state-item fMRI design that allowed us separately to estimate sustained and transient components of neural activation.

The experimental task was adapted from previous behavioral [23,28,36], computational [17], and neuroimaging studies [40] and involved repeated sampling from two decks of cards, one of which maximized reward in the long run and the other of which minimized it. Crucially, the effects of each deck were delayed, so that repeated sampling from the 'good' deck resulted in a noticeable immediate decrease in reward followed by a gradual increase over succeeding trials, whereas repeated sampling from the 'bad' deck led to the opposite pattern of reward payoff. By varying the size of the window (W) over which the delayed effects accrued, we were able to manipulate the demand for temporal integration, since as window size increases, it becomes increasingly difficult to maintain all of the relevant information (i.e., the rewards accumulated over the last W trials) in mind. For example, a window size of 4 means that the outcomes of the last 4 trials affect the potential outcome of the current trial. On the basis of behavioral pilot studies, we chose one condition in which most participants successfully solve the task and prefer the good deck (W set to 2 trials) and one condition in which most participants develop a strong preference for the bad deck (W set to 12 trials). The contrast between these two conditions allowed us to identify (a) neural mechanisms reflecting the general need for temporal integration in both conditions (i.e., activations common to both conditions), and (b) differences in the pattern of neural activity during optimal (W_2) versus suboptimal (W_{12}) decision making.

Consistent with the notion that temporal integration during decision making depends primarily on general cognitive control mechanisms, we hypothesized that both conditions would be associated with sustained activation in working memory-related brain areas since both require participants to integrate substantial amounts of reward information over time. It was not clear a priori, however, what differences to expect in neural activity between conditions. On the one hand, it seemed plausible that an increase in window size from 2 to 12 trials would translate

into a greater need for temporal integration in the latter condition. This would predict greater activation in cognitive-control regions in the W_{12} condition, since activation in these regions has been shown to scale with increasing task demands [8]. On the other hand, because participants prefer the bad deck in the W_{12} condition, and the magnitude of this preference often is greater than the preference for the good deck in the W_2 condition, an alternative interpretation is that in the W_{12} condition, individuals are insensitive to the delayed temporal effects and respond only to the immediate effect of switching decks (a large immediate increase or decrease). On this account, W_{12} should impose relatively fewer cognitive-control demands because participants may assume that they have solved the task when they settle on the bad deck and subsequently cease to attempt to integrate information. Finally, it also is possible that the differences between the two conditions would be more subtle and emerge not as an overall trend towards greater cognitive control-related activation in one condition than the other, but in differential associations with sustained versus transient modes of processing. To test these various hypotheses, we modeled sustained and transient neural activation separately in each condition.

We were interested in identifying two types of brain regions. First, to test the hypothesis that temporally extended decision making requires sustained neural activation, we sought to identify regions showing a consistent increase in sustained task-related activation relative to event-related activity across task conditions. Thus, an important constraint was that activation in these regions could not interact with window size. Second, to examine subtler changes in state and event-related activity between conditions, we also identified regions where state or event-related activity was significantly modulated by window size.

2. Materials and methods

2.1. Participants

Participants were 28 right-handed undergraduates at Washington University in St. Louis or individuals from the surrounding community (mean age = 22.4 years, SD = 3.6). All participants consented to participate in return for financial remuneration ($25/h plus task earnings). Potential participants with a neurological or psychiatric history were excluded from the experiment, which was approved by the Washington University IRB. Due to technical malfunctions during neuroimaging, two participants did not complete both conditions. Results are therefore reported only for the 26 participants with intact data.

2.2. Experimental task

The computerized decision task was adapted from previous experiments (e.g., Ref. [28]) and was instantiated as a card game in which participants made repeated selections from one of two decks. Participants were provided with the following instructions: "The task is to draw cards from one of two decks to earn money. Each time you see a question mark, you should choose a card from the left or right deck. Some cards pay more money and others less money; money is never lost. The two decks are different, and there is a way to earn significantly more or less points. You will be paid real money for each point. Try to earn as much as possible." Participants were given practice trials to familiarize them with the task, but not told anything further and were left to figure out how to earn as much as possible.

Choice of one deck always paid out a larger immediate reward than the other deck, but decreased the amount of the reward to be received from both decks over the next several trials. In contrast, choice of the other deck paid a smaller amount immediately following its selection on a given trial but increased the amount to be received from both decks over several following trials. Thus, repeated choice of the locally optimal deck, what we call the bad deck (i.e., the one with the highest immediate reward), leads to the global minimum in terms of total earnings. Conversely, the globally larger reward is earned from repeatedly choosing the deck that pays off less on any given trial, which we call the good deck because it leads to the global maximum.

The amount of points earned on a given trial was computed to be the sum of four separate quantities: (a) a constant base amount (awarded regardless of choice); (b) a long-term effect (dependent on choice history); (c) a smaller immediate effect (dependent only on current selection); and (d) a small amount of random noise. A formal reward equation is provided in Appendix A; however, two features of the reward computation are important to note here. First, the long-term effect is greater than the short-term effect (in the current study, the short-term effect was 2 points and the long-term effect was 4 points; see Appendix A for further details). Second, a selection from one of the two decks, hereafter referred to as the good deck, served to increase the long-term effect but at the same time also decreased the short-term effect. Because of the greater impact of the long-term effect on reward, repeated selection from the good deck would lead to the global maximum of reward payoff. Yet because of the paradoxical reduction in the short-term effect, this benefit of good deck selection could take a number of trials to accrue. In contrast, selection from the other deck, hereafter referred to as the bad deck, did not positively impact either the long- or short-term payoff. Consequently, repeated selection from the bad deck would lead to only the baseline amount being awarded (i.e., the global minimum payoff). The reward function can also be conceptualized in terms of deck-switching effects. A switch in deck selection from the good deck to the bad deck will cause an immediate increase in reward but also a delayed reduction in the amount that can be accrued. Conversely, a switch from the bad deck to the good deck

will not cause an immediate increase in reward payoff (and can even cause an immediate decrease) but will have a substantial delayed beneficial effect on the amount of reward that can be accrued.

Another variable that also impacted the reward payoff was the number of trials over which the long-term effect of choice accrued. This variable, hereafter referred to as the window size or W, determines how many of the previous trials (up to and including the current one) must be taken into account when calculating the amount awarded on the current trial. For example, if W is set to 2, only choices made on the current trial and the one before it affect the amount earned. As W grows, the delay between choice and change in reward increases. If W is very large – say, 12 – decisions made a long time ago (i.e., more than 10 trials back) still influence the amount earned on the current trial.

Because of this dynamic, value of W has a substantial impact on the difficulty of the task. The larger the value of W, the greater the number of trials over which the long-term effects accrue. As a result, the magnitude of increment or decrement in the long-term effect per trial becomes increasingly smaller. However, the magnitude of the immediate effect remains constant independent of W. As a result, the balance between immediate and long-term rewards is shifted and participants become more sensitive to the immediate effect because the long-term effect is harder to detect. For example, when W is small, the reward amount due to the long-term effect changes in large increments (e.g., for $W = 2$, the delayed effect increases by 2 points per trial, when using the long-term value chosen in the current study). Under such conditions it is relatively easy for participants to detect this effect and come to preferentially choose the good deck even though such choices do not always provide an immediate increase in reward. If, however, W is set to a larger value, say 12, then the contribution of the long-term effect changes in very small increments (e.g., only one-third of a point on each trial, given the values used in the current study). Moreover, because the reward amount participants receive is always rounded to the nearest integer, in some cases changes in participants' behavior may not show up as an increase in reward payoff for several trials. Consequently, participants are much more likely to fail to discern the underlying pattern of the long-term effect when W is large than when it is small and therefore less likely to exercise appropriate cognitive control. Thus, increasing W should lead directly to poorer decision-making performance (i.e., reduced frequency of selection from the good deck and a lower amount of total reward amount received in the experiment).

In the present study, W was set at 2 in one condition and at 12 in the other condition (W_2 and W_{12}, respectively). Performance was defined as the percentage of choices from the good deck (range: 0–100%; 50% = no consistent preference or chance performance). Because trials on which participants completely failed to respond resulted in no reward, omissions (1.4% of all trials across all participants)

were grouped with 'bad' selections for the purpose of computing choice performance.

Stimuli were presented centrally on a Macintosh visual display, using Psyscope software [13]. Each trial began with the presentation of two decks of cards faced down, with a question mark appearing midway between them (see Fig. 1, top). The display was presented until the participant pressed one of two response keys (to indicate deck choice), from a specially constructed button box. If a response was not made within 2500 ms, the trial was scored as "no response." Immediately after a choice was made, the corresponding deck was "turned over" to reveal a number, representing the amount earned on that trial (Fig. 1, bottom). This outcome display lasted for a variable period no greater than 5000 ms (participant's response time + outcome display = 5000 ms). The response time, deck chosen, and amount earned were recorded, and cumulative earnings were updated on each trial. A variable inter-trial interval then occurred which ranged from 0 to 5 s (in steps of 2500 ms). The variability (jittering) of the inter-trial interval enabled event-related responses to be estimated for each trial type.

2.3. Neuroimaging procedures

To accommodate fMRI scanning, participants performed the task as 8 scanning runs of 40 card choice trials per run, yielding a total of 160 trials per task condition (W_2 and W_{12}). All 8 runs were performed consecutively on the same day. Winnings were updated across runs, and participants were explicitly instructed that they were doing the same
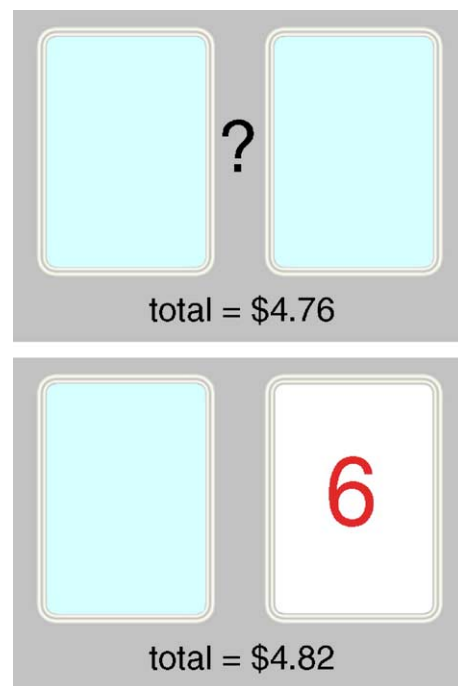

Fig. 1. Decision-making task: Participants select from 1 of 2 decks, and the number of points earned or lost is revealed after a brief delay. Cumulative earnings also are displayed.

task. After four runs, they were informed that they now were doing a task that looked the same but was, in fact, a different task, and that they should approach it "starting from scratch." The order in which the two conditions were performed was counterbalanced across participants.

Whole-brain images were acquired on a 3-T head-only Allegra System (Siemens, Erlangen, Germany), with a standard circularly polarized head coil. Head movement during scanning was minimized using pillows and tape. Headphones dampened scanner noise and enabled communication with participants. Structural images were acquired using an MP-RAGE T1-weighted sequence. Functional images were acquired using an asymmetric spin-echo echo-planar sequence (TR = 2520 ms, TE = 25 ms, flip = 90°), sensitive to blood oxygen level-dependent (BOLD) magnetic susceptibility. Each volume contained 40, 3.75-mm thick slices (in-plane resolution 3.75 × 3.75 mm). Due to a change in acquisition software, the last 12 participants were erroneously scanned with 1.875 mm spacing between slices (however, this change did not affect data quality, given the imaging resolution and spatial smoothing used in pre-processing).

Each scanning run consisted of alternating cycles of task blocks (including task trials and jittered inter-trial intervals) and fixation blocks. The inclusion of fixation blocks was an important feature of the scanning design to enable us to conduct state-item analyses. Task blocks were 20 trials in duration. Including interleaved fixations, each block was approximately 150 s long. Fixation blocks (denoted by a centrally presented crosshair) were 40 s in duration. The first four images of functional runs were discarded to allow for stabilization of longitudinal magnetization (indicated by T1 times). Each run lasted approximately 5.5 min.

Functional imaging data were pre-processed prior to statistical analyses according to the following procedures. All functional images were first temporally aligned across the brain volume, corrected for movement using a rigid-body rotation and translation correction [22,55], and then registered to the participant's anatomical images (in order to correct for movement between the anatomical and function scans). The data then were scaled to achieve a whole-brain mode value of 1000 for each scanning run (to reduce the effect of scanner drift or instability), resampled into 3 mm isotropic voxels, and spatially smoothed with a 9 mm FWHM Gaussian kernel. Participants' structural images were transformed into standardized atlas space [56] using a 12-dimensional affine transformation [60,61]. The functional images were registered to the reference brain using the alignment parameters derived for the structural scans.

A general-linear model approach [21] was used to estimate parameter values for both event-related responses (item effects) and for sustained activity associated with the entire task block (state effects). State effects can be independently coded into the GLM, using an assumption of a fixed-shape response of long duration (i.e., boxcar convolved with a gamma function). The logic of this approach is that the event-related effects should be decaying back to baseline during inter-trial intervals, whereas state-related effects should remain relatively constant and of increased amplitude relative to blocks of fixation. In recent work, this approach to GLM coding of sustained and transient responses has been validated via both simulation and empirically based methodological studies [57]. Event-related effects were analyzed by estimating values for the various time points within the hemodynamic response. The duration of this epoch was taken to be 20 s (8 scanning frames). The 8-point event-related estimates were then cross-correlated with a standard hemodynamic response function to compute an estimated event-related response magnitude. Finally, the magnitude estimates for state and event-related effects for each individual participant were submitted to a group analysis using random-effects model ANOVAs or t tests.

To identify brain regions showing either sustained or transient activation during the decision-making task, we used a conjunction approach that involved the application of multiple contrast tests, with each set at a relatively low threshold. Such procedures have been used in previous studies [7,9,46], and we believe they optimize the trade-off between sensitivity/power and false-positive protection (i.e., type I vs. type II error). In order for a brain region to be accepted as selective for a particular effect, all voxels within the region were required to be statistically significant ($P < 0.05$) in all tests for that effect (described below). Note that the $P < 0.05$ threshold was applied to each map individually and was not the overall family-wise error rate for the entire analysis. The estimated overall level of protection is about 0.05^N, where N is the number of maps entered into a given conjunction (for a similar approach, see Ref. [24]). Moreover, a region was considered significant only if it contained a cluster of 8 or more contiguous voxels. The cluster-size requirement ensured an additional overall image-wise false-positive rate of $P < 0.05$ [20,39]. Finally, to increase interpretability, only positive activations were considered in all of these analyses.

We conducted two distinct sets of analyses. First, to test the hypothesis that different neural networks were selectively associated with sustained versus transient processing during the decision-making task, we separately identified regions that showed either significantly greater state-related activity than event-related activity, or significantly greater event-related activity than state-related activity. Two constraints were that (a) the same pattern had to hold in both W_2 and W_{12} conditions, and (b) there could not be a significant interaction between window-size (WIN: W_2 vs. W_{12}) and temporal dynamics of activation (DYN: state vs. effect). These constraints ensured that activation in the resulting regions reflected only a main effect of DYN that was not significantly moderated by WIN. This first analysis yielded two maps, one containing regions that

displayed greater event-related than state-related activity, and the other containing regions that showed the opposite pattern.

Second, to test whether activity in some brain regions displayed a significant WIN effect or WIN × DYN interaction (for instance, significantly greater state-related activity in W_2 relative to W_{12}, but no differences in event-related activity), we conducted a series of conjunction analyses that combinatorially tested for all possible patterns of differences in activity. Specifically, it was possible for (a) both sustained and transient neural activity to be greater in one window-size condition than in the other; (b) sustained but not transient activity to be greater in one condition than the other; and (c) transient but not sustained activity to be greater in one condition than the other. Since there were two directions for which any of the above could be true (i.e., $W_2 > W_{12}$ or $W_{12} > W_2$), there were six conjunctions in all. Table 1 lists the individual effects that were entered into each conjunction.

There were several constraints imposed on each test. First, in addition to significant differences between conditions for a particular effect type, a region also had to show activity that was significantly different from baseline fixation. Second, a stipulation for the four conjunctions that tested for increases specific to one DYN effect (e.g., greater state-related activity but not greater event-related activity) was that regions also had to show a significant WIN × DYN interaction. Conversely, for the two conjunctions testing for increases common to both DYN effects (state/event), we stipulated that regions not show a WIN × DYN interaction. Finally, an additional constraint for the 4 DYN-specific conjunctions was that only activity in the DYN condition (state/event) under consideration showed a significant WIN effect in the direction being tested.

To illustrate, consider the analysis testing for regions that showed greater state-related activity in W_2 than W_{12} but not greater event-related activity. For this test, the following effects were entered into the conjunction: (a) W_2 state-related activity significantly greater than W_{12} state-related activity; (b) W_2 state-related activity significantly greater than baseline; (c) W_2 event-related activity not greater than W_{12} event-related activity; and (d) significant WIN × DYN interaction. The use of these multiple constraints thus ensured conceptual consistency while making the overall analysis more conservative.

3. Results

3.1. Behavioral performance

To verify that the WIN manipulation had the intended behavioral effect, we compared choice of the good deck in the short-WIN condition (W_2) with that in the long-WIN condition (W_{12}). Mean percentage choice of the good deck for the W_2 and the W_{12} conditions were 71% and 23% across all trials, respectively, and each differed significantly from chance ($t(25) = 6.4$ and -12.5, respectively, both $Ps < 0.001$) as well as from each other ($t(25) = 10.41$, $P < 0.001$). Moreover, all 26 participants performed better in W_2 than in W_{12} (i.e., made more choices of the good deck and earned more money).

3.2. fMRI data

Table 2 lists all clusters that showed significantly greater state-related activity than event-related activity in both the W_2 and W_{12} conditions relative to the fixation baseline (Table 2, state > event; Fig. 2, green). As predicted, the majority of clusters fell within putative cognitive control regions, including bilateral PFC, medial frontal cortex, parietal cortex, and lateral cerebellum. The magnitude of the state- and event-related responses for a representative region (right inferior parietal cortex) is illustrated in Fig. 3A. The network of regions showing greater event-related than state-related activation was more widespread (Table 2; event > state), recruiting both posterior cortical regions likely associated with perceptual encoding and response execution (e.g., visual and somatomotor cortex), along with subcortical regions potentially involved in reward processing and choice selection (e.g., basal ganglia, amygdala, ventromedial frontal cortex). The diffuse nature of the activation was likely due to the non-selective nature of the event-related contrast (i.e., comparison against fixation) and the fact that transient activation is in general likely to be more widespread than sustained activity. Because our analysis procedure did not decompose more general transient components of task processing from those selective to decision making, we do not speculate further on the significance of this event-related activity. The decomposition of event-related activity is best accomplished by comparing different kinds of trials to each other based on

Table 1
Effects entered into individual conjunction analyses

Conjunction	Description	Sustained effect	Transient effect	WIN × DYN interaction
1	W_2 greater overall	$W_2 > W_{12}$	$W_2 > W_{12}$	No
2	W_2 greater sustained	$W_2 > W_{12}$	W_2 not $> W_{12}$	Yes
3	W_2 greater transient	W_2 not $> W_{12}$	$W_2 > W_{12}$	Yes
4	W_{12} greater overall	$W_{12} > W_2$	$W_{12} > W_2$	No
5	W_{12} greater sustained	$W_{12} > W_2$	W_{12} not $> W_2$	Yes
6	W_{12} greater transient	W_{12} not $> W_2$	$W_{12} > W_2$	Yes

Table 2
Co-ordinates for significantly activated regions

Contrast	BA	Coordinates			Size (mm^3)
		x	y	z	
State > event (both conditions)					
Left frontopolar cortex	10	−38	51	0	2079
Right dorsolateral prefrontal cortex	9	40	36	30	351
Left dorsolateral prefrontal cortex	6/9	−40	15	45	1323
Medial frontal cortex	8/32	−2	30	45	999
Right lateral temporal cortex	21	64	−36	−12	864
Left inferior parietal cortex	40	−46	−60	45	4509
Right inferior parietal cortex	40	46	−57	45	2619
Left lateral cerebellum		−34	−72	−48	1107
Right lateral cerebellum		44	−75	−42	756
Event > state (both conditions)					
Ventromedial frontal cortex	10	−2	60	0	1323
Anterior cingulate cortex	24/32	2	3	39	972
Supplementary motor cortex	6	2	−15	72	2484
Left primary motor cortex	4	−34	−30	66	2970
Left somatomotor cortex	1/2/3/4	−56	−21	−30	10,260
Right primary motor cortex	4	50	−6	42	5778
Left superior temporal cortex	22	−50	6	−6	972
Left ventral temporal cortex	18/37	−28	−51	−18	15,795
Right tempoparietal cortex	22/40	56	−36	24	10,854
Left superior parietal cortex	7	−26	−57	48	1242
Right inferior parietal cortex	40	26	−72	36	11,097
Left occipitoparietal cortex	19	−28	−84	15	11,502
Striate visual cortex	17/18	8	−75	−3	27,081
Left extrastriate visual cortex	18/19	−34	−78	−9	31,455
Right extrastriate visual cortex	18/19	38	−72	−6	60,021
Left insula/basal ganglia/thalamus		−20	−9	6	39,879
Right insula/basal ganglia		26	−6	15	12,339
Right insula		34	3	0	19,386
Right thalamus/brainstem		10	−30	−6	45,009
Right amygdala/parahippocampal gyrus		28	0	−27	1674

outcome measures (e.g., reward earned), an analysis which is beyond the scope of the current report (but see Ref. [62] for such analyses).

Of the six focused conjunction analyses described in Table 1, only two yielded significant results. Several
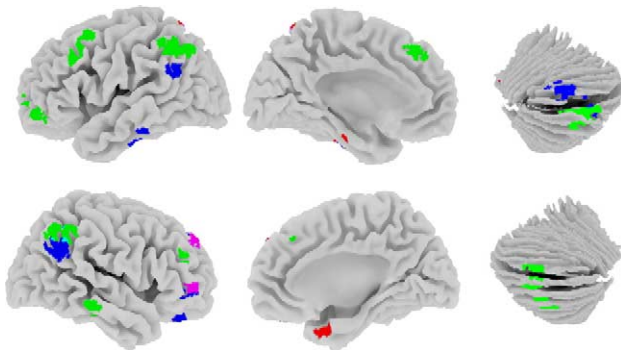

Fig. 2. Brain regions showing significant activation for analyses reported in the text. Green: regions showing greater state-related activity than event-related activity in both experimental conditions. Blue: regions showing a selective increase in state-related activity in W_{12} relative to W_2. Red: regions showing a selective increase in event-related activity in W_2 relative to W_{12}. Purple: overlap between blue and red regions.

regions showed significantly more state-related activity in W_{12} than in W_2 but no differences in event-related activity (Table 3, W_{12} state > W_2 state; Fig. 2, blue), and several showed more event-related activity in W_2 than in W_{12} but no differences in state-related activity (Table 3, W_2 event > W_{12} event; Fig. 2, red). Interestingly, a number of the regions showing relatively greater state-related activity in the W_{12} condition than in the W_2 condition were contiguous with regions showing greater state-related activity in both conditions. This pattern held true bilaterally in inferior parietal cortex and in left lateral cerebellum. Magnitudes for event- and state-related activity in right inferior parietal cortex are depicted in Fig. 3B. It is important to remember that regions showing a greater sustained response in the W_{12} than in the W_2 condition (blue) were constrained in that they also had to show an interaction between window-size (W_2 vs. W_{12}) and type of activity dynamics (state vs. event), whereas regions showing a larger state-related response in both conditions (green) were constrained not to show such an interaction. Thus, the fact that these two types of regions are contiguous in several cognitive control areas is by no means forced and likely reflects increased recruitment of additional regions in W_{12}, presumably in
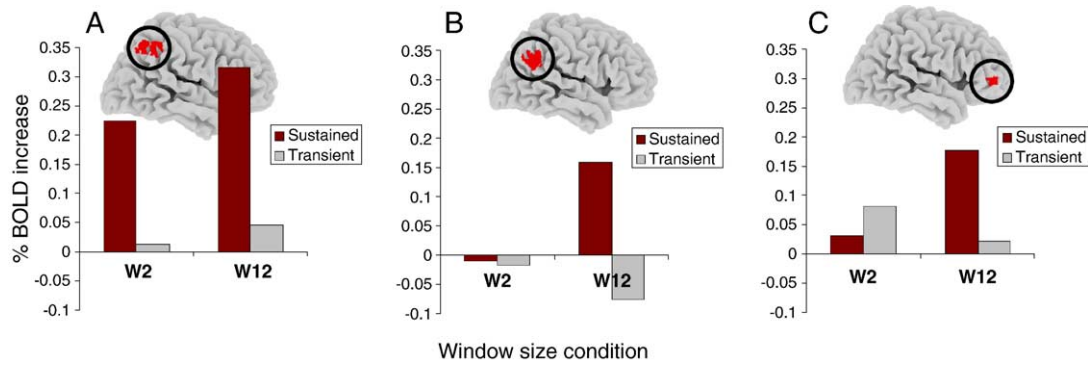
Fig. 3. Estimated magnitude of sustained and transient neural activity as a function of experimental condition for 3 selected regions. (A) Activation for a region of right inferior parietal cortex displaying greater state-related activity than event-related activity in both conditions. (B) Activation in a contiguous region displaying a selective state-related increase in W_{12} relative to W_2. (C) Right frontopolar region showing both greater state-related activity in W_{12} than in W_2 and greater event-related activity in W_2 than in W_{12}.

response to increased task demands. We return to this point below.

It also is intriguing to note that the only two conjunction analyses that provided significant results did so in opposite directions. In the W_{12} condition, several regions showed greater state-related activity, but no regions showed greater event-related activity; in the W_2 condition, several regions showed greater event-related activity but not state-related activity. Were these results due to chance, we would expect an approximately equal distribution of regions showing greater state- and event-related activity in both the W_2 and W_{12} conditions. Moreover, because the two observed effects were potentially compatible (i.e., a region could in principle show both relatively greater state-related activity in W_{12} and relatively greater event-related activity in W_2), we conducted an overlap analysis to determine whether any regions were sensitive to both effects. Four regions survived this test (Table 3, W_{12} state > W_2 state AND W_2 event > W_{12} event; Fig. 2, purple), two of which were located in right prefrontal cortex, one in the left superior parietal lobe and one in rostromedial cerebellum. The magnitudes of the event-related and state-related activity are displayed for the region in right frontopolar PFC (Fig. 3C). A crossover interaction is clearly evident; an identical pattern of results holds in the other three regions (not depicted).

Table 3
Co-ordinates for significantly activated regions

Contrast	BA	Coordinates			Size (mm³)
		x	y	z	
W_{12} state > W_2 state					
Right frontopolar cortex	10/46	40	42	3	891
Right dorsolateral prefrontal cortex	8/9	14	45	42	621
Right ventral prefrontal cortex	47/11	16	33	−18	405
Left extrastriate cortex	19/20	−56	−33	−27	1296
Left tempoparietal junction	39	−56	−60	24	1620
Right tempoparietal junction	39	50	−60	33	4671
Right inferior parietal cortex	40	50	−60	51	216
Left superior parietal cortex	7	−14	−69	60	270
Medial rostral cerebellum		−20	−27	−36	864
Lateral cerebellum		−50	−63	−33	999
W_2 event > W_{12} event					
Right frontopolar cortex	10/46	40	45	6	459
Right dorsolateral prefrontal cortex	8/9	14	48	42	621
Left lateral temporal cortex	21	−62	−6	−9	405
Left superior parietal cortex	7	−16	−69	60	513
Medial rostral cerebellum		−20	−30	−30	486
Uncus	28	22	6	−27	918
W_{12} > W_2 state + W_2 > W_{12} event					
Right frontopolar cortex	10/46	40	42	6	459
Right dorsolateral prefrontal cortex	8/9	14	48	42	486
Left superior parietal cortex	7	−14	−69	60	270
Medial rostral cerebellum		−16	−30	−33	216

This type of crossover pattern between more event-related activity in W_2 and more state-related activity in W_2 is strongly suggestive of a change in processing mode, whereby different temporal dynamics govern activation in these regions depending on task condition. If the cross-over pattern is indicative of a trade-off or shift between sustained and transient processing modes, then the pattern should be observable at the level of individual differences as well. Consequently, a prediction suggested by a sustained-transient processing mode shift across conditions is that individuals who show a greater difference between W_{12} and W_2 state-related activity should show a correspondingly greater difference between W_2 and W_{12} event-related activity. We tested this prediction by computing the Pearson correlation coefficient between the differences in state-related and event-related activity estimates (i.e., W_{12} state $-$ W_2 state correlated with W_2 event $-$ W_{12} event) within each of the four brain regions that showed the crossover pattern. A strong positive correlation was found in all four regions (rs = 0.57 $-$ 0.78, all Ps < 0.003). That is, the greater a participant's sustained response in the W_{12} condition relative to the W_2 condition, the greater his or her event-related response in the W_2 relative to the W_{12} condition.

We then examined how specific this correlational effect was to the regions identified with the group-level cross-over interaction pattern. To test specificity, we performed an identical correlational analysis on 10 other regions, including 3 regions that displayed greater state-related activity than event-related activity in both conditions (Fig. 2, green) and 7 randomly selected regions that did not appear in any of our analyses. Surprisingly, similar effect sizes were obtained in all cases (all rs > 0.62, all Ps < 0.001). Thus the trade-off between sustained and transient neural activity appears to be pervasive at the level of individual participants and is observable even in regions that at the group level were constrained not to show an interaction between task condition and effect type. This result is not easily interpretable in the present context, but it is potentially suggestive of a more general relationship between sustained and transient neural activity during task performance.

4. Discussion

As predicted, greater state-related activation common across both the short and long decision-making conditions was observed almost exclusively in brain regions associated with working memory and cognitive control, namely lateral PFC, medial frontal cortex, and parietal cortex. This finding supports our analysis that the choice task requires temporally extended integration of information, a capacity subserved by neural mechanisms that actively can maintain internal representations over sustained intervals [8,11,54]. In contrast, event-related activation was extremely diffuse and included most of the visual and motor cortex, in addition to both working memory and reward processing regions. Thus,

whereas the transient activation locked to specific trials likely reflected perceptual and motor planning processes as well as the neural response to reward outcomes, sustained activity appeared to engage primarily working memory circuits associated with maintenance of information on-line. In the present context, such activation could reflect either the simultaneous representation of multiple trial outcomes (as participants attempted to extract a pattern) or the sustained representation of active strategies (such as the intention to continue choosing from a particular deck).

Interestingly, only event-related but not state-related activation was found in reward- or punishment-related regions such as the ventromedial frontal cortex, amygdala, or ventral striatum, providing support for the hypothesis that there are dissociable contributions of affective-motivational and cognitive neural systems to decision-making behavior. Previous studies have suggested that affective-motivational brain regions such as the striatum and ventromedial frontal cortex are involved in representing the anticipation or response to individual rewards (e.g., Refs. [18,49]). However, there is no reason to presume that the same regions also must be involved in the prediction, integration, and representation of those outcomes at a higher level. It may be more parsimonious to suppose that temporal integration functions rely on the same neurocognitive mechanisms regardless of whether the contents they operate over happen to be past actions and rewards or some other type of information more commonly associated with working memory (numbers, words, etc.). Our results appear to support this notion.

As expected, there were striking differences in the choice responses of the participants in the two conditions of the task. These differences were reflected by greater focal state-related activity in the W_{12} condition than in the W_2 condition, and greater focal event-related activation in the W_2 than in the W_{12} condition, with no other patterns present. As noted above, the state-related increases in the W_{12} condition were mostly located in regions contiguous with areas that showed greater state-related than event-related activity in both the W_2 and W_{12} conditions, suggesting that increased task demands in the W_{12} conditions result in increased recruitment of these regions. Also consistent with this interpretation, several regions exhibited non-selective state-related activation in one hemisphere but W_{12}-specific state-related activation in the contralateral homologue. This pattern was present in both lateral temporal cortex and frontopolar cortex and is consistent with research suggesting that the contralateral homologue of a task-relevant region often comes on-line as the task-related load increases [44,59]. Thus, taken together, the results suggest an interpretation in which the increasing demands on temporal integration from the W_2 to the W_{12} condition produce a corresponding increase in sustained activity in parietal and prefrontal regions. This increased activity could reflect the extra processing resources devoted to carrying out temporal integration operations.

As appealing as this account may be, it runs into difficulty explaining the differences in choice behavior between the two conditions. Specifically, participants exhibited extremely poor self-control in the W_{12} condition, showing a strong preference (77% of selections) for the bad deck. This is problematic in that it is not immediately clear why participants recruit greater cognitive control when performance is so poor. One alternative possibility that can account for the behavioral performance pattern is that participants may have employed a win–stay lose–shift (WSLS) strategy throughout the experiment, staying with the same deck following an increase in reward and switching decks following a decrease in reward [41]. Such a strategy is consistent with the poorer performance in the W_{12} condition because the immediate gain associated with the bad deck is more noticeable than are the delayed gains associated with the good deck. A WSLS strategy, however, is not consistent with the observed pattern of brain activation. First, such a strategy is thought to be fairly automatic, reactive to specific outcomes, and require little sustained cognitive effort. It therefore does not explain the observed sustained activity in control-related brain regions. Second, a WSLS also provides no obvious explanation for the increase in sustained activity in these regions in the W_{12} condition relative to the W_2 condition.

Another possibility that might account for both the behavioral and neuroimaging results is that greater sustained activation in the W_{12} condition may correspond not to increased temporal integration, but to increased reliance on an acquired strategy that must be maintained throughout the task. On this account, the critical variable relating behavioral performance to neural activity is not success per se but absolute difference from chance. In other words, although participants performed poorly in the W_{12} condition (23%), the degree to which they selected from the bad deck suggests they had developed a strong and consistent preference for it (after all, no preference = 50%) and were not simply guessing. In contrast, although participants developed a preference for the good deck in the W_2 condition, the magnitude of this preference above the no preference point (71%–50% = 21%) was somewhat smaller than the degree of preference below that of no preference in the W_{12} condition (50%–23% = 27%). This difference in magnitudes was marginally significant, $t(25) = 1.94$, $P = 0.06$. The fact that greater state-related activation in the W_{12} condition also was found in a region of right frontopolar prefrontal cortex further supports this view. Previous research suggests a role for this area in complex planning and subgoal processing [1,7,34], a characterization consistent with the representation of a strategy for repeated selection over time. Finally, in a previous report based on the same data but using an individual differences approach [62], we found that greater event-related activity in right dorsolateral PFC predicted poorer performance in the W_2 condition but better performance in the W_{12} condition, possibly reflecting the fact that

some participants were less likely to rely on strategy and more likely to guess on the W_2 trials.

It is important to note that the increasing difficulty of the task as W increases is due to a number of factors that cannot be dissociated. First, as the number of trials over which reward is computed increases, participants are required to keep track of their behavior over an increasing number of trials, imposing greater working memory demands; second, the incremental change in reward decreases, making trial-to-trial change more difficult to discern; and third, because the amount of noise added to the reward function is constant across different values of W, the ratio between reward and noise level decreases as W increases. The precise influence of these factors cannot be dissociated within this paradigm without altering several parameters besides W. However, our goal in this initial study was not to parameterize the relationship between performance and the various quantities in the reward equation but to investigate the role of cognitive control in decision making. The important point for our purposes is that increasing W systematically increases the difficulty of the task by imposing increasing demands for cognitive control, allowing comparisons to be made across conditions with differential control demands.

Given the present results and the difficulty in decomposing various effects of the window-size manipulation, it is not possible to adjudicate between the temporal integration account and the strategy–representation account. Future studies will be needed to more carefully and systematically manipulate various experimental factors related to the reward–payoff function. For example, a less-extreme manipulation of the window-size parameter might allow for a more powerful way to test for the presence of temporal integration effects. In particular, it would be ideal to use a window-size value at which choice of the good deck was significantly lower than in the W_2 condition, but still significantly greater than chance (e.g., 60% good deck selections). In such a condition, the temporal integration account predicts a smaller but still significant increase in sustained activity in comparison to the W_2 condition. In contrast, a strategy-maintenance account should predict less sustained activity than in the W_2 condition, because the deck preference would be less strongly established (i.e., closer to 50%). The ability to incrementally manipulate window size, and hence deck preference, with this paradigm suggests that such a study is feasible.

The overarching goals of the current study were both conceptual and methodological. Conceptually, our goal was to demonstrate that temporally extended decision making relies upon neural mechanisms associated with cognitive control. Methodologically, our goal was to demonstrate the utility of hybrid state-item designs in the decision-making domain. These goals were synergistic since our primary hypothesis was that there would be a qualitative distinction in activity dynamics between brain regions subserving temporally extended control processes and those associated with the anticipatory and response components of individual

trials. The results confirm this hypothesis and further suggest that experimental manipulations of temporal integration demands produce selective effects on sustained and transient activation.

It is not possible to state with certainty that the state-related effects we obtained are attributable to a task-related demand for integration of temporally extended information or sustained maintenance of strategic information, but these are both parsimonious interpretations that are empirically testable. Regardless of the ultimate explanation, there is little question that partitioning neural activity into transient and sustained components in the present study provided information that would have been unavailable using either an event-related or a blocked design. Whereas the former design overlooks maintenance-related activity and the latter design provides non-specific temporal information, hybrid designs can uncover a comprehensive spectrum of temporal dynamics. Our understanding of decision making and choice behavior, which involves a range of processes operating on different time scales, should benefit from the application of similar techniques in future neuroimaging studies.

Acknowledgment

Appendix A

We present here the technical details of the reward payoff function used in our experimental paradigm. Formally, the payoff equation giving the reward, r, earned on trial j can be expressed as:

$$r_j = B + L\frac{\sum\limits_{i=0}^{W-1} c_{j-i}}{W} - Ic_j + e$$

The four terms of this equation refer to the baseline reward value ($B = 4$), the long-term choice parameter ($L = 4$), the immediate choice parameter ($I = 2$), and a random noise component ($e = +1$ on 20% of trials, 0 on 60% of trials, and -1 on 20% of trials). The long-term choice component is computed by summing deck choices, c, where c is a binary variable: a value of 1 denotes a choice from the good deck, and a value of 0 denotes a choice from the bad deck. Deck choices are summed over the previous W trials, where W is the window-size variable (set to 2 in one condition and 12 in the other), up to and including the current trial j. The division by W gives a proportional value between 0 and 1 that is then scaled by the long-term parameter L. Note that an increase in the number of selections made from the good deck over the last W trials results in an increase in r. In contrast, the immediate choice

component is dependent only on the current choice c_j and is negative. Thus, a selection from the good deck on the current trial decreases the current reward. Note also that the scaling factor I for the immediate choice effect is only half the size of the long-term choice effect L (i.e., 4 vs. 2). This means that long-term choices will have a greater impact on accumulated reward than immediate choices.

Fig. 4 provides a visualization of the reward function. In this figure, the two parallel lines correspond to the payoffs from the bad (top line) and good (bottom line) decks. Notice that for any given point along the x-axis, the payoff from the bad deck is always higher. However, selections from these two decks 'move' a participant in different directions along the x-axis. Each selection from the good deck moves the participant to the right, because the percentage of good-deck selections made over the last W trials (i.e., the x-axis) increases. In contrast, each selection from the bad deck moves the participant to the left along the x-axis. This dynamic corresponds to the contradictory effects of the good and bad decks: the good deck provides a smaller reward on the current trial but increases reward in the long term (since it moves a participant to the right, i.e., up the payoff line); conversely, the bad deck typically yields more on the current trial but results in a smaller payoff on future trials. Another way to express this dynamic is that repeatedly selecting the local maximum (the bad deck) leads directly to the global minimum, whereas repeatedly selecting the local minimum (the good deck) leads to the global maximum. This tension between immediate and delayed reward is what gives rise to the need for cognitive control: in order to perform optimally, participants must inhibit the impulse to pursue immediate rewards in order to increase their gains in the long term.
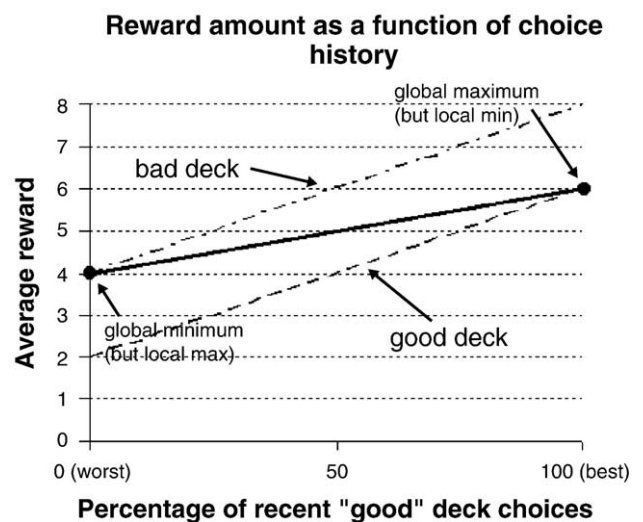

Fig. 4. Amount of reward as a function of deck selection over the previous few trials. The global and local maxima are in direct opposition; hence, reward maximization requires participants to endure a local reduction for several trials before achieving the global maximum.

The equation parameters can be related to the graph in Fig. 4 as follows. The y-intercept for the higher line indicates the baseline value ($B = 4$ points) awarded on each trial regardless of choice. The x-axis refers to the percentage of good-deck choices made within the last W trials; this percentage, expressed as a ratio and multiplied by the long-term component ($L = 4$), is equal to the contribution of the long-term (i.e., delayed) effect of choice to the overall amount earned (range: 0 to 4 points). The difference between the two parallel lines represents the immediate effect of deck choice ($I = 2$ points). Thus, selecting from the good deck on any given trial decreases the number of expected points to be earned on that trial by a value of I.

Finally, the parameter W can be conceptualized as the number of steps it takes to go from one end of the reward function to the other. For example, suppose $W = 2$, and a participant has just earned 6 points by selecting the bad deck (top line). If the participant were to then select the bad deck again on the following trial, he or she would move all the way down the rest of the line to the bottom (since it only takes 2 steps to get from end to end). Consequently, the participant would earn 4 points. However, if W had been set to a large value, say 12, a second selection from the bad deck would result in only a very small decrease in the reward amount, since the participant would only move 1/12th of the way down the line.

References

[1] S.C. Baker, R.D. Rogers, A.M. Owen, C.D. Frith, R.J. Dolan, R.S. Frackowiak, T.W. Robbins, Neural systems engaged by planning: a PET study of the Tower of London task, Neuropsychologia 34 (1996) 515–526.

[2] D.J. Barraclough, M.L. Conroy, D. Lee, Prefrontal cortex and decision making in a mixed-strategy game, Nat. Neurosci. 7 (2004) 404–410.

[3] A. Bechara, H. Damasio, A.R. Damasio, G.P. Lee, Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making, J. Neurosci. 19 (1999) 5473–5481.

[4] A. Bechara, H. Damasio, A.R. Damasio, Emotion, decision making and the orbitofrontal cortex, Cereb. Cortex 10 (2000) 295–307.

[5] A. Bechara, H. Damasio, A.R. Damasio, Role of the amygdala in decision-making, Ann. N. Y. Acad. Sci. 985 (2003) 356–369.

[6] K.C. Berridge, Pleasures of the brain, Brain Cogn. 52 (2003) 106–128.

[7] T.S. Braver, S.R. Bongiolatti, The role of frontopolar cortex in subgoal processing during working memory, NeuroImage 15 (2002) 523–536.

[8] T.S. Braver, J.D. Cohen, L.E. Nystrom, J. Jonides, E.E. Smith, D.C. Noll, A parametric study of prefrontal cortex involvement in human working memory, NeuroImage 5 (1997) 49–62.

[9] T.S. Braver, J.R. Reynolds, D.I. Donaldson, Neural mechanisms of transient and sustained cognitive control during task switching, Neuron 39 (2003) 713–726.

[10] R.L. Buckner, P.A. Bandettini, K.M. O'Craven, R.L. Savoy, S.E. Petersen, M.E. Raichle, B.R. Rosen, Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic resonance imaging, Proc. Natl. Acad. Sci. U. S. A. 93 (1996) 14878–14883.

[11] R. Cabeza, L. Nyberg, Imaging cognition II: an empirical review of 275 PET and fMRI studies, J. Cogn. Neurosci. 12 (2000) 1–47.

[12] K. Cheng, J. Pena, M.A. Porter, J.D. Irwin, Self-control in honeybees, Psychon. Bull. Rev. 9 (2002) 259–263.

[13] J.D. Cohen, B. MacWhinney, M. Flatt, J. Provost, PsyScope: an interactive graphic system for designing and controlling experiments in the psychology laboratory using Macintosh computers, Behav. Res. Methods Instrum. Comput. 25 (1993) 257–271.

[14] A.R. Damasio, Descartes' Error: Emotion, Reason, and the Human Brain, Grosset/Putnam, New York, NY, 1994.

[15] D.J. de Quervain, U. Fischbacher, V. Treyer, M. Schellhammer, U. Schnyder, A. Buck, E. Fehr, The neural basis of altruistic punishment, Science 305 (2004) 1254–1258.

[16] D.I. Donaldson, S.E. Petersen, J.M. Ollinger, R.L. Buckner, Dissociating state and item components of recognition memory using fMRI, NeuroImage 13 (2001) 129–142.

[17] D.M. Egelman, C. Person, P.R. Montague, A computational role for dopamine delivery in human decision-making, J. Cogn. Neurosci. 10 (1998) 623–630.

[18] R. Elliott, J.L. Newman, O.A. Longe, J.F. Deakin, Differential response patterns in the striatum and orbitofrontal cortex to financial reward in humans: a parametric functional magnetic resonance imaging study, J. Neurosci. 23 (2003) 303–307.

[19] M. Ernst, E.E. Nelson, E.B. McClure, C.S. Monk, S. Munson, N. Eshel, E. Zarahn, E. Leibenluft, A. Zametkin, K. Towbin, J. Blair, D. Charney, D.S. Pine, Choice selection and reward anticipation: an fMRI study, Neuropsychologia 42 (2004) 1585–1597.

[20] S.D. Forman, J.D. Cohen, M. Fitzgerald, W.F. Eddy, M.A. Mintun, D.C. Noll, Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a cluster-size threshold, Magn. Reson. Med. 33 (1995) 636–647.

[21] K.J. Friston, A.P. Holmes, K.J. Worsley, J.-P. Poline, C.D. Frith, R.S.J. Frackowiak, Statistical parametric mapping in functional imaging: a general linear approach, Hum. Brain Mapp. 2 (1995) 189–210.

[22] K.J. Friston, S. Williams, R. Howard, R.S.J. Frackowiak, R. Turner, Movement-related effects in fMRI time-series, Magn. Reson. Med. 35 (1996) 346–355.

[23] J.R. Gray, A bias toward short-term thinking in threat-related negative emotional states, Pers. Soc. Psychol. Bull. 25 (1999) 65–75.

[24] J.R. Gray, C.F. Chabris, T.S. Braver, Neural mechanisms of general fluid intelligence, Nat. Neurosci. 6 (2003) 316–322.

[25] L. Green, J. Myerson, Alternative frameworks for the analysis of self control, Behav. Philos. 21 (1993) 37–47.

[26] L. Green, J. Myerson, D. Lichtman, S. Rosen, A. Fry, Temporal discounting in choice between delayed rewards: the role of age and income, Psychol. Aging 11 (1996) 79–84.

[27] M. Haruno, T. Kuroda, K. Doya, K. Toyama, M. Kimura, K. Samejima, H. Imamizu, M. Kawato, A neural correlate of reward-based behavioral learning in caudate nucleus: a functional magnetic resonance imaging study of a stochastic decision task, J. Neurosci. 24 (2004) 1660–1665.

[28] R.J. Herrnstein, G.F. Loewenstein, D. Prelec, W. Vaughan, Utility maximization and melioration: internalities in individual choice, J. Behav. Decis. Mak. 6 (1993) 149–185.

[29] J.M. Hinson, T.L. Jameson, P. Whitney, Impulsive decision making and working memory, J. Exp. Psychol. Learn. Mem. Cogn. 29 (2003) 298–306.

[30] J.D. Jentsch, J.R. Taylor, Impulsivity resulting from frontostriatal dysfunction in drug abuse: implications for the control of behavior by reward-related stimuli, Psychopharmacology (Berlin) 146 (1999) 373–390.

[31] B. Knutson, R. Peterson, Neurally reconstructing expected utility, Games Econ. Behav. (in press).

[32] B. Knutson, C.M. Adams, G.W. Fong, D. Hommer, Anticipation of increasing monetary reward selectively recruits nucleus accumbens, J. Neurosci. 21 (2001) RC159.

[33] B. Knutson, G.W. Fong, C.M. Adams, J.L. Varner, D. Hommer, Dissociation of reward anticipation and outcome with event-related fMRI, NeuroReport 12 (2001) 3683–3687.

[34] E. Koechlin, G. Basso, P. Pietrini, S. Panzer, J. Grafman, The role of the anterior prefrontal cortex in human cognition, Nature 399 (1999) 148–151.

[35] M.L. Kringelbach, J. O'Doherty, E.T. Rolls, C. Andrews, Activation of the human orbitofrontal cortex to a liquid food stimulus is correlated with its subjective pleasantness, Cereb. Cortex 13 (2003) 1064–1071.

[36] E. Kudadjie-Gyamfi, H. Rachlin, Temporal patterning in choice among delayed outcomes, Org. Behav. Human Decis. Process. 65 (1996) 61–67.

[37] J.E. LeDoux, Emotion circuits in the brain, Annu. Rev. Neurosci. 23 (2000) 155–184.

[38] A.W. Logue, Research on self-control: an integrating framework, Behav. Brain Sci. 11 (1988) 665–709.

[39] M.P. McAvoy, J.M. Ollinger, R.L. Buckner, Cluster size thresholds for assessment of significant activation in fMRI, NeuroImage 13 (2001) S198.

[40] P.R. Montague, G.S. Berns, Neural economics and the biological substrates of valuation, Neuron 36 (2002) 265–284.

[41] M. Nowak, K. Sigmund, A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game, Nature 364 (1993) 56–58.

[42] J. O'Doherty, M.L. Kringelbach, E.T. Rolls, J. Hornak, C. Andrews, Abstract reward and punishment representations in the human orbitofrontal cortex, Nat. Neurosci. 4 (2001) 95–102.

[43] J.P. O'Doherty, R. Deichmann, H.D. Critchley, R.J. Dolan, Neural responses during anticipation of a primary taste reward, Neuron 33 (2002) 815–826.

[44] A.M. Passarotti, M.T. Banich, R.K. Sood, J.M. Wang, A generalized role of interhemispheric interaction under attentionally demanding conditions: evidence from the auditory and tactile modality, Neuropsychologia 40 (2002) 1082–1096.

[45] H. Rachlin, L. Green, Commitment, choice and self-control, J. Exp. Anal. Behav. 17 (1972) 15–22.

[46] J.R. Reynolds, D.I. Donaldson, A.D. Wagner, T.S. Braver, Item- and task-level processes in the left inferior prefrontal cortex: positive and negative correlates of encoding, NeuroImage 21 (2004) 1472–1483.

[47] J. Rilling, D. Gutman, T. Zeh, G. Pagnoni, G. Berns, C. Kilts, A neural basis for social cooperation, Neuron 35 (2002) 395–405.

[48] R.D. Rogers, A.M. Owen, H.C. Middleton, E.J. Williams, J.D. Pickard, B.J. Sahakian, T.W. Robbins, Choosing between small, likely rewards and large, unlikely rewards activates inferior and orbital prefrontal cortex, J. Neurosci. 19 (1999) 9029–9038.

[49] E.T. Rolls, The orbitofrontal cortex and reward, Cereb. Cortex 10 (2000) 284–294.

[50] B.R. Rosen, R.L. Buckner, A.M. Dale, Event-related functional MRI: past, present, and future, Proc. Natl. Acad. Sci. U. S. A. 95 (1998) 773–780.

[51] A.G. Sanfey, J.K. Rilling, J.A. Aronson, L.E. Nystrom, J.D. Cohen, The neural basis of economic decision-making in the ultimatum game, Science 300 (2003) 1755–1758.

[52] G. Schoenbaum, B. Setlow, M.P. Saddoris, M. Gallagher, Encoding predicted outcome and acquired value in orbitofrontal cortex during cue sampling depends upon input from basolateral amygdala, Neuron 39 (2003) 855–867.

[53] W. Schultz, Getting formal with dopamine and reward, Neuron 36 (2002) 241–263.

[54] E.E. Smith, J. Jonides, Storage and executive processes in the frontal lobes, Science 283 (1999) 1657–1661.

[55] A.Z. Snyder, Difference image versus ratio image error function forms in PET–PET realignment, in: D. Bailer, T. Jones (Eds.), Quantification of Brain Function Using PET, Academic Press, San Diego, 1996.

[56] J. Talairach, P. Tournoux, Co-planar Stereotaxic Atlas of the Human Brain, Thieme, New York, 1988.

[57] K.M. Visscher, F.M. Miezin, J.E. Kelly, R.L. Buckner, D.I. Donaldson, M.P. McAvoy, V.M. Bhalodia, S.E. Petersen, Mixed block/event-related designs separate transient and sustained activity in fMRI, NeuroImage 19 (2003) 1694–1708.

[58] N.D. Volkow, J.S. Fowler, G.J. Wang, J.M. Swanson, Dopamine in drug abuse and addiction: results from imaging studies and treatment implications, Mol. Psychiatry 9 (2004) 557–569.

[59] D.H. Weissman, M.T. Banich, Global-local interference modulated by communication between the hemispheres, J. Exp. Psychol. Gen. 128 (1999) 283–308.

[60] R.P. Woods, S.R. Cherry, J.C. Mazziotta, Rapid automated algorithm for aligning and reslicing PET images, J. Comput. Assist. Tomogr. 16 (1992) 620–633.

[61] R.P. Woods, S.T. Grafton, C.J. Holmes, S.R. Cherry, J.C. Mazziotta, Automated image registration: I. general methods and intrasubject, intramodality validation, J. Comput. Assist. Tomogr. 22 (1998) 139–152.

[62] T. Yarkoni, T.S. Braver, J.R. Gray, and L. Green, Prefrontal brain activity predicts temporally extended decision-making behavior, J. Exp. Anal. Behav. (Submitted for publication).