

Effort discounting in human nucleus accumbens

MATTHEW M. BOTVINICK, STACY HUFFSTETLER, AND JOSEPH T. MCGUIRE
Princeton University, Princeton, New Jersey

A great deal of behavioral and economic research suggests that the value attached to a reward stands in inverse relation to the amount of effort required to obtain it, a principle known as *effort discounting*. In the present article, we present the first direct evidence for a neural analogue of effort discounting. We used fMRI to measure neural responses to monetary rewards in the human nucleus accumbens (NAcc), a structure previously demonstrated to encode reference-dependent reward information. The magnitude of accumbens activation was found to vary with both reward outcome and the degree of mental effort demanded to obtain individual rewards. For a fixed level of reward, the NAcc was less strongly activated following a high-demand for effort than following a low demand. The magnitude of this effect was noted to correlate with preceding activation in the dorsal anterior cingulate cortex, a region that has been proposed to monitor information-processing demands and to mediate in the subjective experience of effort.

At the level of everyday commonsense knowledge, there is a close relationship between reward and effort. This is evident, for example, when a potential payoff is judged not to be worth the work it would require, or when a level of reward for some effort is considered to be unfair. The same direct connection between reward and effort is also found in formal theories, including behavioral and economic accounts of decision making (see Kivetz, 2003; Walton, Kennerley, Bannerman, Phillips, & Rushworth, 2006), social theories of equity (Walster, Walster, & Berscheid, 1978), and legal theories of distributive justice (Locke, 1690/1987). In these and other contexts, a common proposition is that effort carries a negative value or cost, sometimes referred to as the *disutility of effort*, and that this cost provides a reference against which earned rewards are evaluated (Figure 1). According to this basic principle, referred to in some contexts as *effort discounting*, a reward carries a higher net value if it is easily obtained than if it is obtained only through great effort (Kivetz, 2003; Phillips, Walton, & Jhou, 2007; Rudebeck, Walton, Smyth, Bannerman, & Rushworth, 2006). In effect, effort sets in place a reference point against which rewards are measured.

Effort discounting and the close relationship between reward and effort that underlies it clearly represent more than mere cultural convention. Rodents, birds, and non-human primates have been shown to weigh effort against reward in decision making (Phillips et al., 2007; Salamone, Cousins, & Bucher, 1994; Stevens, Rosati, Ross, & Hauser, 2005; Tsunematsu, 2001; Walton, Bannerman, Alterescu, & Rushworth, 2003; Walton et al., 2006; some contrary conclusions from the same literature are discussed below), and capuchin monkeys have been shown to reject rewards smaller than those received by conspecifics

for an equal expenditure of effort (Brosnan & de Waal, 2003). Psychopharmacological interventions and lesions to specific brain structures have been observed to alter the relative weighting of effort and reward information in decision making (Denk et al., 2005; Floresco & Ghods-Sharifi, 2007; Salamone, Correa, Mingote, & Weber, 2003; Salamone et al., 1994; Walton et al., 2003). Furthermore, effort discounting relates closely to another form of discounting, delay discounting, for which specific neural substrates have been identified (see, e.g., Roesch, Taylor, & Schoenbaum, 2006; Rudebeck et al., 2006). Given such findings, it seems plausible that the tight relationship between reward and effort that holds at the behavioral level may reflect the operation of basic neural mechanisms.

In the present experiment, we used fMRI to investigate the integration of reward and effort information in the human brain. Our specific objective was to test for a neural correlate of effort discounting. The experiment focused on the nucleus accumbens (NAcc), a basal ganglia structure that has been found in numerous studies to respond to reward outcomes, often in a reference-dependent fashion (Breiter, Aharon, Kahneman, Dale, & Shizgal, 2001; Delgado, Nystrom, Fissell, Noll, & Fiez, 2000; Elliott, Friston, & Dolan, 2000), and has also been heavily implicated in effort-based decision making (Salamone, Correa, Farrar, & Mingote, 2007; Salamone et al., 2003; Salamone et al., 1994; Walton et al., 2006). In alignment with a recent study of delay discounting (Roesch et al., 2006), we looked for discounting effects at the time of reward receipt, following effort expenditure. On the basis of theories of effort discounting (e.g., Kivetz, 2003; Phillips et al., 2007), we hypothesized that the NAcc response to reward would vary inversely with the level of effort demanded prior to reward delivery.

M. M. Botvinick, matthewb@princeton.edu

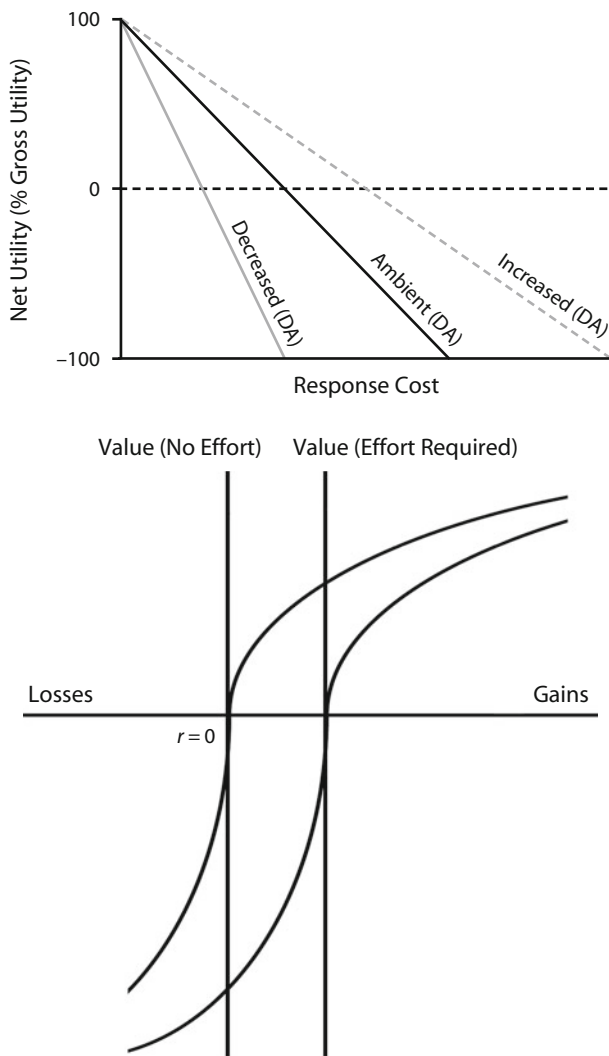


Figure 1. Two diagrammatic representations of effort discounting: one from animal behavior research and one from behavioral economics. **Top:** Based on choice behavior in animals, Phillips, Walton, and Zhou (2007) proposed that the subjective value (net utility) of a fixed reward varies inversely with the effort required to obtain it (response cost). They reviewed work suggesting that the strength of this discounting effect varies with concentrations of dopamine (DA) within the NAcc. Adapted from Phillips et al. (2007), Figure 2, with kind permission from Springer Science+Business Media. **Bottom:** According to prospect theory (Tversky & Kahneman, 1992), the relationship between a reward's magnitude and its subjective value is characterized by a curvilinear value function. Kivetz (2003) proposed that this function shifts to the right as the effort demanded to obtain the reward increases. The figure shows two value curves: one relevant to rewards requiring a fixed amount of effort (right), the other to rewards requiring no effort (left). Horizontal axis: reward magnitude, including both wins and losses, with 0 at the point labeled $r = 0$. Vertical axis: subjective value, with the horizontal axis crossing at 0. Adapted from Kivetz (2003), Figure 1, by permission. Copyright 2003 INFORMS, the Institute for Operations Research and the Management Sciences, 7240 Parkway Drive, Suite 310, Hanover, MD 21076.

Although effort discounting has been considered to result from both physical and mental effort, our focus in the present experiment was on mental effort only. One advantage of this choice is that relatively solid information

is available concerning neural response to both cognitive demand and the associated subjective sense of mental effort. In particular, these have been proposed to engage the dorsal anterior cingulate cortex (ACC; Botvinick, Braver, Barch, Carter, & Cohen, 2001; Botvinick, Cohen, & Carter, 2004; Naccache et al., 2005). Studying the effects of mental effort on reward processing thus presents the opportunity to compare neural encodings of reward (within the NAcc) with encodings of effort itself (within the ACC).

METHOD

Participants

The experiment was performed at the University of Pennsylvania. A total of 45 participants were recruited from the university's student and staff population. All denied any history of neurological or psychiatric conditions. Participants were randomly sorted into an experimental group (30 total, 22 females) and a control group (15 total, 8 females).¹ Of the 30 experimental participants, 7 females were ultimately excluded from analysis, 3 due to scanner malfunction and 4 due to excessive movement, on the basis of criteria described below. Of the 15 control participants, 2 were excluded due to excessive movement. All participants included in the final analysis were right-handed and were from 19 to 33 years of age. They provided informed consent and were paid for their participation.

Behavioral Task

The experimental protocol was approved by the University of Pennsylvania School of Medicine's institutional review board. Participants in both the experimental and control groups performed in a task-switching paradigm involving probabilistic reward (see Figure 2). In each task block, participants viewed a series of 10 individually presented numerals ranging from 1 to 9 but excluding 5. All visual stimuli were displayed on a back-projection screen placed at the head of the scanner bore, which was viewed through a mirror mounted on the head coil. The numerals were presented on a black background in 72-point Arial font. The participants responded to each numeral by pressing one of two buttons on a fiber-optic response pad. The correct response depended on the color in which the numeral was presented. If the target numeral appeared in yellow, the participants were to perform a parity judgment, depressing the button beneath their right index finger to indicate even and a button beneath their right middle finger to indicate odd. If the target numeral appeared in blue, the participants used the same two buttons to perform a magnitude judgment (index finger, lower than 5; middle finger, greater than 5).

The experiment involved two types of task block. In low-demand blocks, all digits were of the same color. In high-demand blocks, the color alternated across successive digits, requiring the participant to make effortful and inefficient switches between tasks (see Monsell, 2003). In a parallel behavioral study reported elsewhere (Botvinick, 2007), participants given a free choice between these two task conditions showed a consistent bias against the high-demand option, consistent with the idea that this carried relatively high disutility.

Each block was preceded by a brief visual cue that indicated whether the upcoming block would be of the high- or low-demand type. A solid disk (yellow or blue) indicated a low-demand block and additionally indicated the relevant task. A vertically split disk, half yellow and half blue, indicated a high-demand block, and provided no information about the color of the first digit. Within each scan and over the course of the experiment as a whole, the two block types and the two classification tasks occurred with equal frequency, following a different randomized sequence for each participant.

Immediately following the last digit in each block, a message appeared, reading "Deciding your pay . . ." After a variable interval, this was replaced with a large white dollar sign, indicating that the




Figure 2. The sequence of events within low-demand (left) and high-demand (right) blocks. Each block included 10 numerals and could end with either a \$ or an X.

participant had earned \$1 for the just-completed block, or a large white X, indicating that nothing had been earned for the block.

The specifics of stimulus timing were as follows: Disk cue duration was 2,000 msec. Digit duration was 1,500 msec, with a 500-msec interstimulus interval. Message duration ranged arbitrarily from 2,000 to 8,000 msec in 2,000-msec steps. Reward cue duration was 2,000 msec. The interblock intervals ranged arbitrarily from 2,000 to 8,000 msec in 2,000-msec steps.

At the beginning of each session, the experimenter read the instructions aloud to the participant. The participants were then given 5 min of practice on the task-switching paradigm, without reward cues. During the fMRI experiment, participants performed the task (with rewards) throughout four scanning periods, each approximately 10 min in duration (see the Image Acquisition section below). Each scanning period contained 18 task blocks. Within each period, the sequence of high- and low-demand blocks and of \$ versus X cues was randomly generated, subject to the constraint that each scan could contain no more than five and no fewer than three occurrences of any specific combination of block type and cue type. This resulted in roughly equal numbers of each block/cue combination over the course of the experiment. Note that it was important to prevent significant imbalances in reward rate between high- and low-demand blocks in order to avoid corresponding differences in reward prediction. At the same time, to assure that the participants would take interest in the reward cues, it was necessary to create the impression that the range of possible cumulative rewards was fairly wide. With this in mind, we told participants prior to participation that they would earn between \$10 and \$50; and as part of the instructions, they were told that the amount of pay would depend directly on the number of \$ cues that occurred during their session. All of the participants received \$40 for their participation (slightly greater than the number of \$ cues, which ranged from 32 to 39).²

Prior to the fMRI experiment, the participants were informed that the rate of pay would be equal for high- and low-demand blocks. In addition, it was made clear that pay would not be contingent upon speed or accuracy during task performance. Nevertheless, the participants were encouraged to respond as quickly as possible without sacrificing accuracy.

The decoupling of reward probability from reaction time (RT), error rate (ER), and block type was a critical aspect of the experimental design and merits further comment. Note, first, that rewards were not intended to serve as incentives to effort. Effort was manipulated entirely through externally imposed task demands, rather than through differential incentives. Our interest was in how external demands for effort would affect reward processing, rather than in how rewards would affect the voluntary mobilization of effort on a fixed task. In order to prevent confusion on this point, we refer to the difference between block types in our experiment as relating to “demand” rather than to “effort.”

Second, note that if reward had been made dependent on RTs or ERs, this would have introduced differences in reward rates between high- and low-demand blocks, because RTs and ERs could be expected to differ between the two. This would likely have led to different expectancies for reward between the two block types. Specifically, the participants would have grounds for predicting a \$ outcome on low-demand blocks more strongly than on high-demand blocks, simply on the basis of observed frequencies. This is important because NAcc reward responses have been shown to vary depending on reward expectancy (see, e.g., Breiter et al., 2001). Thus, making reward probability dependent on RTs or ERs would have introduced a serious confound.

It was also important to inform the participants that reward probabilities would be equivalent between block types, in order to prevent the culturally based assumption that the experimenters would reward high-demand blocks with more frequent \$ outcomes than they would low-demand blocks. Loosely speaking, the idea of effort discounting suggests that rewards are valued in reference to what is felt to be deserved, not only to what is predicted. By explicitly informing the participants of the equivalent reward probabilities between block types, we sought to neutralize the prediction factor to get the participants to focus on the entitlement factor.

Note that it was possible, in principle, for the participants to perceive the reward events as being completely independent of the temporally associated task-performance blocks. However, pilot work and informal interviews suggested that this was prevented by framing rewards as “pay.” In order to understand why this framing may

have been effective, consider that pay-for-work in everyday life is rarely immediately contingent on the details of work performance. For instance, a factory worker's hourly pay is typically not immediately dependent on the worker's hourly output. In general, in our culture (indeed, including the setting of the typical psychology experiment), one is usually paid a preestablished, fixed amount for performing work within certain, usually implicit, performance bounds. This was precisely the situation in our experiment. It was considered that any effort-discounting effect observed in the fMRI data would provide supportive, if not indisputable, evidence that the participants had understood reward events to be related to block completion.

The control version of the task was designed to verify that any effort-discounting effect observed in the experimental group was specific to reward processing, and not due to nonspecific or incorrectly modeled effects of demand on NAcc activation. The control task was identical to the experimental task, with two exceptions. First, the message at the end of each block read "Calculating, please wait . . ." Second, the message was followed by an S or a K, rather than a \$ or an X. Control participants were told that the message and these letters were for the information of the experimenter and were irrelevant to the participant's task. However, the instructions were to pay attention during the letter cues, nonetheless, "to make sure you don't miss the beginning of the next block." The instructions prior to the experiment did not mention any block-by-block rewards for task performance. The participants were told they would earn \$20 for their participation, and they received this amount at the end of the experiment.³

Image Acquisition

We collected fMRI images using a Siemens Trio scanner operating at 3 Tesla with a Siemens eight-channel head coil. Each scanning session began with the acquisition of a high-resolution, T1-weighted, axial anatomical image. The anatomical scan was followed by four scans of axial gradient-echo echoplanar images with the following specifications: repetition time (TR) = 2,000 msec, time to echo = 30 msec, $\alpha = 90^\circ$, matrix = 64×64 , field of view = 24 cm, $3 \times 3 \times 3$ mm voxels. We collected 33 contiguous slices during each TR. Four functional scans were completed, each containing 306 TRs.

Image Analysis

Offline data processing was performed using the VoxBo software package (www.voxbo.org). After image reconstruction, the anatomical images were corrected for field inhomogeneities and transformed to standardized space, defined by a Montreal Neurological Institute (MNI) template. The 4-D functional time-series data were interpolated to correct for slice-acquisition timing, spatially realigned to the first TR of the first scan of the study (to correct for participant movement), spatially normalized, and spatially smoothed using a 4-mm fixed width at half-maximum kernel. Within each scan, the signal at each voxel was mean normalized and linearly detrended. Data for participants showing greater than 2 mm of movement in any plane during a single functional scan were excluded from further analysis.

Subsequent analysis focused on five a priori regions of interest (ROIs). A bilateral NAcc ROI (mean voxel count = 35) was defined according to the guidelines proposed by the Massachusetts General Hospital Center for Morphometric Analysis (www.cma.mgh.harvard.edu/). Four further ROIs were defined for comparison with NAcc (see the Results section for further comments on the selection of these ROIs). A medial prefrontal cortex ROI (44 voxels) was estimated from Table 1 and Figure 1 of Knutson, Fong, Adams, Varner, and Hommer (2001). This area was reported in that study to show sensitivity to cues indicating reward outcomes.

The resulting ROI was centered at Talairach coordinates 3, 40, -12; the Talairach coordinates were mapped to MNI space using the tal2mni.m function developed at Cambridge University (see imaging.mrc-cbu.cam.ac.uk/imaging/MniTalairach). A bilateral ACC ROI (206 voxels) was based on a meta-analysis of studies implicating the ACC in cognitively demanding tasks (Ridderinkhof, Ullsperger, Crone, & Nieuwenhuis, 2004). The region chosen was

based on Figure 1B of Ridderinkhof et al. (2004) and was shaped so as to encompass the region of highest density within the cluster of ACC locations depicted in that figure (centered approximately at Talairach coordinates 0, 14, 41). The resulting ROI corresponded closely to a region identified in a study by Bush et al. (2002) as being involved in reward evaluation. Orbitofrontal and amygdala ROIs were drawn based on inspection of anatomical images. Coverage of the OFC was adequate in 22 of 23 experimental participants. In initial analyses, separate ROIs were drawn for medial and lateral divisions of the OFC (mean voxel count = 190 and 821, respectively), following previous work suggesting differential function across this divide (e.g., O'Doherty, Kringelbach, Rolls, Hornak, & Andrews, 2001). However, data from these subregions were pooled after no significant difference was found in their response to reward cues (\$ vs. X). A left dorsolateral prefrontal cortex (DLPFC) ROI (82 voxels) centered at Talairach coordinates -48, 21, 21 was drawn from Luks, Simpson, Feiwell, and Miller (2002), which reported activation within this region in association with task switching, the source of cognitive demand in our study.

Within each ROI, spatially averaged signal time courses were passed through a notch filter, removing frequencies above 0.25 Hz and below 0.0032 Hz. The resulting time courses for each ROI were then analyzed using the general linear model (GLM), as implemented in VoxBo. The model included covariates for 10 task events: solid-disk cue, split-disk cue, digit within high-demand block, digit within low-demand block, message following high-demand block, message following low-demand block, \$ following high-demand block, \$ following low-demand block, X following high-demand block, and X following low-demand block. Each of these covariates was convolved with a canonical hemodynamic response function. The intertrial interval provided a baseline. The model also included covariates for intercept, scan effects, and any individual mean signal spikes greater than 3.5 standard deviations from the experimental mean.

Parameter estimates (beta values) were extracted for each ROI for each event type. The resulting values were entered into further analyses, as enumerated in the Results section. ROI analyses of reward-cue responses assumed the form of a repeated measures ANOVA with factors for block type (high vs. low demand) and reward cue (\$ vs. X). Analyses comparing reward-cue activity between the experimental and control groups took the form of a repeated measures ANCOVA with factors for block and cue type (\$ and S vs. X and K) and participant group. These analyses also included, as a covariate, the difference between the parameter estimates for ROI activation during high-demand and low-demand block performance.⁴ Analyses comparing reward-cue responses between ROIs involved repeated measures ANOVAs with factors for block type, cue type, and ROI.

In order to further establish regional specificity of findings from earlier analysis, a whole-brain exploratory analysis was also conducted. Data were preprocessed and smoothed as described above. To each voxel's time course we applied the same general linear model that was applied to average time courses in the foregoing ROI analyses. Each participant's four parameter estimates for reward cues were then taken to group-level analyses. Single-subject general linear models were estimated using VoxBo software; group-level ANOVAs were carried out using AFNI (Cox, 1996). Specific contrasts and associated significance thresholds are described in the Results section.

As discussed further in the Results section, a third set of GLM analyses was conducted in order to examine correlations between ACC activation during numeral classification (within individual blocks) and the NAcc response to the subsequent reward cues. The model included individual covariates for each task block and for each reward cue, as well as one covariate for all disk cues and one covariate for all message events. Parameter estimates for the reward-cue response within the NAcc ROI were entered into a linear regression with regressors for cue type, block type, and ACC parameter estimate, as well as for the NAcc parameter estimate for numeral-task performance in the block preceding each reward cue. The NAcc

parameter estimate was included in order to partial out gradual fluctuations in NAcc activation.

RESULTS

Validation of Effort Manipulation

Both behavioral and fMRI data were consistent with the expectation that task blocks involving continual task switching would be more difficult than blocks requiring no switching. Manual response RTs were slower during high-demand blocks ($M = 853$ msec) than during low-demand blocks ($M = 645$ msec) [$t(22) = -15.6, p < .0001$]. Manual response percent correct (PC) was lower during high-demand blocks ($M = 93.7$) than during low-demand blocks ($M = 96.6$) [$t(22) = 3.31, p = .003$]. This effect held for both parity (for low-demand blocks, mean RT = 664 msec, mean PC = 96.5; for high-demand blocks, mean RT = 866 msec, mean PC = 92.7) and magnitude (for low-demand blocks, mean RT = 614 msec, mean PC = 96.7; for high-demand blocks, mean RT = 837 msec, mean PC = 94.7). In a questionnaire completed at the end of the experiment, the participants rated high-demand blocks as being more effortful than low-demand blocks ($M = 7.11$ vs. 3.78 , respectively, on a scale of 1 to 10; Wilcoxon signed-rank test, $z = -4.15, p < .0001$). The questionnaire also asked participants to indicate what they believed would have been "fair pay" for the performance of individual high-demand and low-demand blocks. This question was included in the questionnaire for the last 9 participants only, who consistently assigned a higher dollar amount to high-demand blocks than to low-demand blocks [$M = \$1.89$ vs. $\$0.89$, respectively; $t(8) = 6.93, p = .0001$], consistent with the idea that the effort required by high-demand blocks was associated with significant disutility, and consistent as well with the choice data obtained in the parallel behavioral study (see the Method section).

A further indication that the demand manipulation was successful was that task performance induced greater activation in high-demand blocks than in low-demand blocks within the dorsal ACC [$t(22) = 7.00, p < .0001$], a region widely considered to monitor task difficulty and demand for cognitive control (Botvinick et al., 2001; Botvinick et al., 2004; Botvinick, Nystrom, Fissell, Carter, & Cohen, 1999; Davis et al., 2005; Kerns et al., 2004; Ridderinkhof et al., 2004) and also proposed to mediate in the subjective experience of mental effort (Naccache et al., 2005).

In order for the control group to serve its intended role, it was important to confirm that participants in that group exerted themselves no less than did participants in the experimental group. Consistent with this, RTs and ERs revealed no significant differences between the control and experimental groups. Indeed, both RTs and ERs were numerically, if not statistically, lower in the control group [for high-demand blocks, mean RT = 783 msec, mean PC = 95.5; for low-demand blocks, mean RT = 591 msec, mean PC = 98.2; main effect of group: RT, $F(1,33) = 3.59, p = .07$; ER, $F(1,33) = 0.83, p = .37$].⁵ Ironically, this pattern raises the unanticipated possibility that participants in the control group may have exerted slightly greater effort overall than did participants in the

experimental group. In light of this, it is critical to note that no interaction was found between participant group and block type, consistent with the conclusion that the difference in demand was comparable between groups [RT, $F(1,33) = 0.83, p = .37$; ER, $F(1,33) = 0.03, p = .87$]. Also consistent with this conclusion were the findings that, within the control group, ACC activation during task blocks was greater for high-demand blocks than for low-demand blocks [$t(12) = 4.28, p = .001$] and that the size of this effect did not differ significantly from that of the experimental group [$t(34) = -0.88, p = .39$].

Effort Discounting in NAcc

In order to analyze responses to reward cues, activation values within NAcc were entered into a two-way repeated measures ANOVA with factors for cue (\$ vs. X) and preceding block type (high vs. low demand; see the Method section). Consistent with the findings of previous studies, the NAcc activated more strongly in response to the \$ cue than to the X cue [main effect of cue, $F(1,22) = 10.97, p = .003$].⁶ We find it critical that NAcc activation also showed a main effect of block type, according to which activation in response to reward cues was higher following low-demand blocks than following high-demand blocks [$F(1,22) = 10.32, p = .004$; see Figure 3A]. This difference was evident for \$ cues [$t(22) = 2.22, p = .037$] and for X cues [$t(22) = 3.09, p = .005$]. Indeed, in the high-demand condition, the X cue was associated with a relative decrease in NAcc activation, consistent with the idea that under sufficient effort demands, a missed reward can, in fact, register as a loss (see the bottom panel of Figure 1; Kivetz, 2003). The context dependence of NAcc responses to cues marking zero reward also fits with that of previous studies, in which NAcc responses to such cues were affected by reward expectancies (Breiter et al., 2001; Knutson, Taylor, Kaufman, Peterson, & Glover, 2005).

Although the difference between the responses to \$ and X cues was larger on average following high-demand blocks than following low-demand blocks, the factors for reward cue and block type did not show a significant interaction [$F(1,22) = 1.54, p = .24$]. Note that such an interaction would follow from some theories of effort discounting, but not from some others. Consider, for example, the two accounts diagrammed in Figure 1. The account illustrated in the upper panel of the figure does not predict an interaction, since net utility is a linear combination of gross reward and effort cost. In contrast, the theory depicted in the lower panel does predict an interaction.⁷ This is due to the shape of the utility function shown in the figure, which builds in an assumption of loss aversion (by assuming a stronger curvature above the x -axis than below it). Of course, both of these predictions depend on the assumption that there is a linear relationship between net utility and NAcc BOLD, and, in this respect, the absence of an interaction in our data must be interpreted with caution. The material point is that the absence of a significant interaction in our data does not contradict any strong prediction arising from the idea of effort discounting.

The absence of a significant interaction between demand and reward cue in the NAcc does, however, raise one

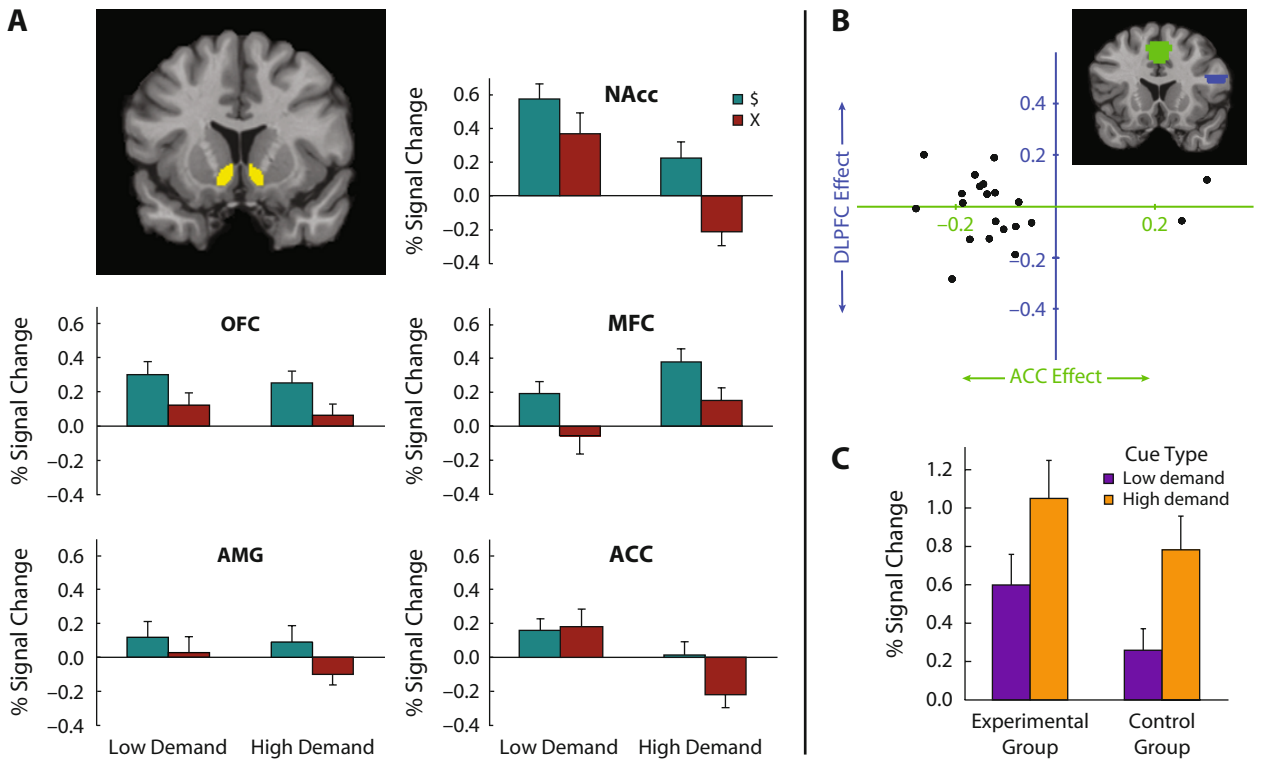


Figure 3. (A) Top left: Coronal section showing the location of the NAcc ROI. The chart labeled “NAcc” shows the mean parameter estimates for NAcc activation in response to \$ and X cues occurring at the completion of high-demand and low-demand blocks, after subtraction of corresponding parameter estimates from the control condition. Bars indicate standard errors. The other charts show these estimates for the OFC (orbitofrontal cortex), MFC (medial prefrontal cortex), AMG (amygdala), and ACC (anterior cingulate cortex). (B) Top: Coronal section showing the location of the ACC and DLPFC ROIs. Bottom: Regression coefficients for each participant in the experimental group, relating reward-cue responses in the NAcc to task-related activity in the ACC and DLPFC. The position along the horizontal axis reflects the relationship between ACC activation during task performance and the NAcc response to the subsequent reward cue. The position along the vertical axis reflects the relationship between DLPFC activation during task performance and the NAcc response to the subsequent reward cue. A more consistent relationship was observed between ACC and NAcc than between DLPFC and NAcc. (C) Mean NAcc responses to the visual cues occurring at the beginning of high- and low-demand blocks, for the experimental group and the control group. Bars indicate standard errors.

important interpretive issue. Specifically, it leaves open the possibility that the apparent discounting effect in the NAcc might instead reflect baseline differences between the high- and low-demand conditions—that is, a simple carryover (or rebound from) activation in the NAcc arising during block performance. Fortunately, this interpretation is ruled out by data from the control group, as discussed next.

Comparison With the Control Group

Results from the control group indicated that the effect obtained in the experimental group was specific to reward processing and not due to nonspecific or incorrectly modeled effects of effort on NAcc activation. A two-way ANOVA on NAcc activation in this control condition (see the Method section) revealed no significant main effect of either cue type [S vs. K; $F(1,12) = 2.72, p = .13$] or block type [$F(1,12) < 0.001, p = .99$]. Pooling across cue types, the mean and standard error for NAcc activity (percent change) were 0.12 and 0.09, respectively, in low-demand blocks and 0.012 and 0.10, respectively, in high-demand blocks.

In order to directly compare the profile of NAcc activation in the experimental and control groups, a mixed-effects

ANCOVA was conducted including within-participants factors for cue and block type and a between-participants factor for participant group (experimental vs. control; see the Method section). This revealed a significant interaction between block type and participant group [$F(1,34) = 16.09, p = .0003$], supporting the conclusion that the impact of demand on the response to the ensuing reward cue differed between the experimental and control groups.

Note that this difference between groups was not explained by any difference in NAcc activity leading up to reward-delivery events. In both groups, NAcc activity during block performance was slightly higher during low-demand than during high-demand blocks. This effect was at trend level in the experimental group [$t(22) = 1.87, p = .074$] and was statistically significant in the control group [$t(22) = 2.77, p = .011$]. However, the size of this difference between block types did not differ significantly between groups [$t(34) = 1.06, p = .30$]. Moreover, the magnitude of the difference for individual participants was included as a covariate in the ANCOVA comparing the experimental and control groups (see the Method section). Thus, the difference in the discounting effect be-

tween groups is unlikely to reflect a carryover from earlier task-induced differences in NAcc activity.

Comparison With Other Regions

In order to investigate the anatomical specificity of the activation pattern observed in the NAcc, we evaluated activation in four other ROIs: two associated in previous research with processing of reward (but not demands for effort), and two associated with both reward and effort processing.

The orbitofrontal cortex (OFC) has been extensively implicated in the representation of reward information (e.g., Rolls, 2004; Schultz, Tremblay, & Hollerman, 2000). However, Rudebeck et al. (2006) provided evidence that the OFC does not play a critical role in the representation of effort costs. Consistent with this, we found that OFC activity was significantly affected by reward information [main effect of reward cue, $F(1,21) = 8.85, p = .007$], but not by the level of effort demanded [main effect of block type, $F(1,21) = 2.12, p = .160$] (see Figure 3A).⁸ The main effect of demand observed in the NAcc differed significantly from the pattern observed in the OFC [$F(1,21) = 5.29, p = .032$].

An analogous pattern was observed in the medial prefrontal cortex (MFC), within a rostral region previously reported to respond to reward outcomes (Knutson et al., 2001; Rogers et al., 2004), but not previously associated with the processing of effort demands (see Figure 3A). As was the case for the OFC, this region showed a main effect of reward cue [$F(1,22) = 4.84, p = .039$], but not of block type [$F(1,22) = 0.36, p = .55$]. The main effect of demand observed in the NAcc differed significantly from the pattern observed in the MFC [$F(1,22) = 14.37, p = .001$].

The basolateral nucleus of the amygdala has been implicated in both reward (see Baxter & Murray, 2002) and effort (Floresco & Ghods-Sharifi, 2007) processing. In our experiment, although the main effect of reward cue did not reach significance in the experimental group [$F(1,22) = 2.45, p = .13$], the amygdala response to \$ versus X did differ significantly from that for control [$F(1,33) = 4.68, p = .038$]. However, the amygdala showed no significant effect of block type [$F(1,22) = 0.76, p = .39$], and the pattern here differed significantly from that seen in the NAcc [$F(1,22) = 5.75, p = .025$].

A different profile was observed in the dorsal ACC, another region associated with both reward (e.g., Bush et al., 2002; Shidara & Richmond, 2002) and effort (e.g., Botvinick, 2007; Botvinick et al., 2004; Walton et al., 2003) processing. The ACC showed a main effect of block type [$F(1,22) = 12.15, p = .002$], but no significant main effect of reward cue [$F(1,22) = 1.19, p = .29$].⁹ The response of the NAcc to reward cues differed significantly from that of the ACC [$F(1,33) = 9.80, p = .005$]. The ACC did show a trend toward interaction between reward and block type [$F(1,22) = 3.01, p = .10$]. However, the interaction pattern did not differ significantly from that observed in the controls [$F(1,33) = 0.04, p = .83$].

The overall pattern of findings is summarized in Table 1. In short, three regions—the OFC, the MFC, and

Table 1
Summary of Results From Region of Interest (ROI) Analyses

ROI	Effect		
	Reward Cue	Block Type	Interaction
Nucleus accumbens	+	+	–
Orbitofrontal cortex	+	–	–
Medial prefrontal cortex	+	–	–
Amygdala	(+)	–	–
Dorsal anterior cingulate	–	+	–

Note—Plus signs indicate that the relevant effects were observed in the experimental group ($p < .05$). The parenthesized plus sign indicates that the relevant effect did not reach significance in the experimental group, but did do so when the experimental group was compared with the control group.

the amygdala—displayed some sensitivity to reward, and one—the ACC—appeared to encode effort costs. However, among the regions studied, only the NAcc showed both effects. The effort-discounting effect observed in the NAcc thus appears to be regionally specific.

To further test the specificity of the pattern obtained in the NAcc, we conducted an exploratory analysis, as detailed in the Method section. Our intention in running this analysis was to localize joint main effects of cue and block type, examining whether these appeared in regions outside the NAcc. However, consistent with the assumptions about statistical power that prompted an initial ROI-based analysis, no region (including the NAcc) passed conventional Bonferroni-corrected significance thresholds, neither for this conjunction of main effects, nor for the interaction between cue and block type. As a follow-up measure, we identified the lowest uncorrected F -value threshold at which at least 1 voxel in both the left and the right NAcc survived for both main effects. This threshold turned out to be $F(1,22) = 9.67$, corresponding to an uncorrected p value of .005 for each main effect. We then surveyed the rest of the brain for voxels that survived this same threshold. A total of 118 such voxels were found: 92 in the occipital cortex (right calcarine gyrus, 13; middle gyrus, 26 right, 32 left; inferior gyrus, 15 right, 4 left; left superior gyrus, 2), 19 in the parietal cortex (superior lobule, 4 left, 8 right; left inferior lobule, 2; right supramarginal gyrus, 1; right angular gyrus, 4), 4 in the frontal cortex (left middle gyrus, 1; left inferior gyrus, 3), 2 in the left inferior temporal gyrus, and 1 in the left cerebellum. Although it is tempting to interpret the occurrence of suprathreshold voxels in some of these regions, caution is dictated by the nature of the analysis, which involves a low (and nondirectional) statistical threshold. What the analysis does clearly indicate is that the dual main effect obtained in the NAcc ROI analysis does not reflect an anatomically generalized, nonspecific pattern. Whether the effect is unique to the NAcc or shared by other specific regions is a question that demands further experimentation.

ACC As a Source of Information on Effort Demands

As discussed earlier, effort discounting can be understood as a form of reference-dependent reward processing. Specifically, the cost of effort can be viewed as setting

a reference point against which rewards are measured (see Figure 1). Thus, in considering the neural correlates of effort discounting, one important question is this: How is the reference point set?

One neural structure that seems likely to be involved is the dorsal ACC. As previously noted, the ACC has been proposed to monitor task difficulty, possibly indexed by conflicts in information processing (Botvinick et al., 2004; Ridderinkhof et al., 2004). This perspective is consistent with our finding, in the present experiment, of greater ACC activation during task performance in high-demand blocks than in low-demand blocks. Other evidence suggests that the ACC may also play a role in effort-based decision making (Botvinick, 2007; Rushworth, Walton, Kennerley, & Bannerman, 2004; Walton et al., 2003). Furthermore, on an anatomical level, the ACC sends direct connections to the NAcc (Croxson et al., 2005; Kunishio & Haber, 1994). All of these considerations make it plausible that the ACC serves as a source of effort demand information to the NAcc.

If this is accurate, then, in the context of our experiment, one would expect activation in the ACC during the numeral classification task (within individual blocks) to correlate inversely with the NAcc response to subsequent reward cues. To test this prediction, we conducted a set of regression analyses based on estimates of block-by-block activation in the ACC and the NAcc (see the Method section). A regression based on the entire group of experimental participants confirmed a highly significant inverse relationship between ACC activation during task performance in individual blocks and the NAcc response to the subsequent reward cue [standardized regression coefficient (β), $-.20, p < .0001$].¹⁰ Notably, this relationship persisted even when an indicator variable for block type was included as a regressor ($\beta = -.16, p < .0001$) and when data for low- and high-demand blocks were analyzed separately (low demand, $\beta = -.10, p = .02$; high demand, $\beta = -.22, p < .0001$). In other words, even after controlling for external demands for effort, activation in the ACC during effort expenditure predicted subsequent reductions in NAcc reward responses. This finding was found to be consistent across participants, with participant-specific regressions yielding negative regression coefficients for ACC activation in 21 of 23 cases [$M = -0.14, t(22) = -3.77, p = .001$; see Figure 3B].

In order to investigate the anatomical specificity of this relationship between the ACC and the NAcc, we performed the same regression analysis on activation in the left DLPFC (see the Method section), another region widely believed to play a role in cognitively demanding situations (Luks et al., 2002; Miller & Cohen, 2001), but not generally associated with a monitoring role (Botvinick et al., 1999). Although this region was indeed more active during task performance in high-demand blocks than in low-demand blocks [$t(22) = 9.97, p < .0001$], its level of activation on individual blocks did not show any significant correlation with the NAcc response to the ensuing reward cue, either in an omnibus analysis ($\beta = -.01, p = .67$) or in individual participant analyses [mean $\beta = -.004, t(22) = -0.14, p = .89$; see Figure 3B].

Effort Discounting or Reward Prediction?

As noted in the Method section, an important consideration affecting our experimental design was that NAcc responses to reward outcomes may be modulated by the degree to which such outcomes are predicted (Breiter et al., 2001; Knutson et al., 2005). A particular concern was that our participants might spontaneously assume that \$ outcomes would occur more frequently after high- than low-demand blocks. If the participants had adopted this assumption, this would have created an alternative explanation for our results—namely, that NAcc responses were reduced after high-demand blocks because the participants literally predicted a higher average level of reward in that context, as compared with NAcc responses after low-demand blocks. Note the critical difference between this and the effort-discounting account, which explains the observed pattern not as a result of differential reward prediction, but as reflecting a process that values rewards relative to their associated effort demands.

In order to discourage differential reward prediction, we explicitly informed the participants that the frequency of \$ outcomes would not differ between block types. The provision of these instructions militates, to some extent, against a reward prediction account of our findings. Nevertheless, it is admittedly conceivable that the participants elected to ignore the information we provided, as well as their own experience in the task, in order to believe that reward rates would differ between block types. In view of this possibility, it is worth considering two further pieces of evidence.

The first is our finding of a correlation between the magnitude of the NAcc effort discounting effect and ACC activity during task performance. In order to explain this aspect of the data in terms of reward prediction, one would need to assume not only that the participants predicted more frequent rewards on high demand blocks, but also that the strength of this prediction was dynamically adjusted on the basis of the specific level of difficulty encountered during performance of each individual block, as indexed by ACC activation. It seems unlikely that the participants would have made such fine-grained reward predictions.

The second piece of evidence involves an aspect of our data that has not been previously discussed. Recall that each block in our experiment began with the presentation of a visual cue that indicated whether the block would involve high or low demand (see Figure 2). Note that for a participant who believed that \$ outcomes were more likely on high-demand blocks, these initial cues should convey differential reward predictions. Given previous evidence indicating that the NAcc responds to events predictive of reward (see, e.g., Knutson & Cooper, 2005; Nicola, Yun, Wakabayashi, & Fields, 2004), one might therefore expect the cue leading high-demand blocks to induce greater NAcc activation than that leading low-demand blocks. One problem with this prediction is that such a difference in cue-related activity could arise from sources other than reward prediction, such as arousal, motor preparation, or simple temporal overlap between cue- and task-related activation in the NAcc. Fortunately, a cleaner prediction can be framed: If the participants expected more frequent

\$ cues on high-demand blocks, the NAcc should respond more strongly to the high-demand cue than to the low-demand cue; and because this difference would relate specifically to reward prediction, it should *not* be observed in the control group, where no rewards were delivered.

The pattern actually observed is presented in Figure 3C. The NAcc did respond more strongly to the high-demand cue. However, contrary to a reward-prediction account, this effect appeared with comparable magnitude in both the experimental group and the control group. The difference was statistically significant in both the experimental group [$t(22) = 4.615, p = .0001$] and the control group [$t(12) = 4.0474, p = .0016$]. An ANOVA with factors for cue type and group yielded a nonsignificant trend toward a main effect of group [$F(68,1) = 2.67, p = .12$], but showed no sign of an interaction between cue and group [$F(68,1) = 0.04, p = .85$].

This pattern of NAcc responses is inconsistent with a reward prediction account, but it is open to several other explanations. For example, as suggested above, it could be connected with elevated arousal in anticipation of high-demand blocks. It is also inviting to consider that the NAcc response to the task cues might relate to the effort-discounting effect seen later in each trial. As we have discussed, effort discounting can be viewed as an instance of reference-dependent reward processing. With this in mind, the cue-related NAcc response might be interpreted as reflecting a reference-setting event, with a stronger NAcc response reflecting the establishment of a higher reference point for reward. If this interpretation is valid, one should expect to see a positive correlation across experimental participants between the size of the cue-type effect in the NAcc and the size of the effort-discounting effect. Such a correlation was indeed observed ($r = .57; p = .005$).¹¹

If the cue-related NAcc response does indeed represent a reference-setting event, it is interesting that this event occurs in the control condition, in which no outcomes are framed as rewards. This could indicate that the reference-setting event associated with anticipated effort is in some sense obligatory. Whatever its interpretation, this aspect of our findings seems to invite further experimental investigation.

DISCUSSION

Taken together, our experimental results suggest that NAcc responses to cued rewards are decremented because of the effort required to earn those rewards, a neural correlate of effort discounting. The data are also consistent with the idea that information about the demand for effort comes to the NAcc, at least in part and in some circumstances, from the dorsal ACC.

These conclusions are concordant with existing theories of both NAcc and ACC function. As noted earlier, the NAcc is widely understood to be involved in the processing of rewards (Breiter et al., 2001; McClure, York, & Montague, 2004; O'Doherty, 2004). One influential interpretation of existing findings is that the NAcc encodes errors in reward prediction (McClure et al., 2004;

O'Doherty, Dayan, Friston, Critchley, & Dolan, 2003; but see Nicola et al., 2004). A critical feature of this account, also found in some contemporary behavioral theories of reward processing (Tversky & Kahneman, 1992), is that reward is represented not in absolute terms, but is instead relative to a standard or reference point. This reference dependence accords well with the present findings, which suggest that the NAcc encodes a reward's magnitude relative to its cost in effort.

Our results are also consistent with assertions that the NAcc plays a role in effort-based decision making (Salamone et al., 2003; Walton et al., 2006). This perspective focuses on cost-benefit analyses computed in advance of effort expenditure, whereas our focus—following recent work on delay discounting (Roesch et al., 2006)—has been on the moment when reward outcomes are revealed: following effort. Nevertheless, the ability to perform cost-benefit analyses depends on knowledge gleaned from previous experience with reward (McClure et al., 2004; O'Doherty et al., 2003) and effort (Botvinick, 2007) outcomes. Thus, there may be a functional connection between the pattern of NAcc behavior reported here and the proposed role of this structure in behavioral decision making. Of note, in preliminary work, Croxson, Walton, and Rushworth (in press) have reported NAcc effort-discounting effects analogous to those we have reported, but occurring when reward cues were presented prior to effort expenditure. Exploring the relationship between Croxson et al.'s findings and the ones reported here presents an important objective for further research.

The present results accord with previous work regarding the ACC, as well. As noted above, one perspective on ACC function links it to the monitoring of task difficulty or the demand for cognitive control (Botvinick et al., 2001; Botvinick et al., 2004; Botvinick et al., 1999; Brown & Braver, 2005; Davis et al., 2005; Kerns et al., 2004; Ridderinkhof et al., 2004). The ACC has also been proposed to mediate in the subjective experience of mental effort (Naccache et al., 2005). Both of these perspectives are clearly consistent with our proposal that the ACC serves as a source of effort demand information in the evaluation of rewards. According to another (not unrelated) theoretical perspective, the ACC appears to play a key role in effort-based decision making (Botvinick, 2007; Floresco & Ghods-Sharifi, 2007; Rushworth et al., 2004; Walton et al., 2003). Walton et al. (2006) have proposed, in particular, that the ACC affects decision making through its influence on the NAcc. The resonance between this proposal and the work we have reported here is readily evident. Nevertheless, it is once again important to acknowledge the distinction between cost-benefit analyses computed ahead of action selection and the representation of rewards at the time of their conferral. It will be interesting to consider how the role of the ACC in modulating online reward processing, as studied in our experiment, may relate to its putative role in cost-benefit analysis and decision making.

The present findings dovetail well with existing neuroscientific data, and they also fit with behavior-based

accounts of effort discounting. In research in both humans and animals, effort discounting has generally been inferred from a preference for rewards that can be easily obtained over rewards demanding greater effort (Denk et al., 2005; Floresco & Ghods-Sharifi, 2007; Kivetz, 2003; Phillips et al., 2007; Rudebeck et al., 2006; Rushworth et al., 2004; Salamone et al., 2007; Salamone et al., 2003; Schweimer & Hauber, 2006; Schweimer, Saft, & Hauber, 2005; Solomon, 1948; Walton et al., 2003; Walton, Bannerman, & Rushworth, 2002; Walton et al., 2006; Zipf, 1949). The present findings complement this behavioral pattern, providing evidence of a direct neural correlate for what has so far been primarily a theoretical construct motivated by behavioral observations.

Having made this point, it is also critical to acknowledge a small but important set of animal behavior studies that have been interpreted as running counter to the notion of effort discounting. For example, several studies (Clement, Feltus, Kaiser, & Zentall, 2000; Friedrich & Zentall, 2004; Kacelnik & Marsh, 2002) have shown that birds prefer rewards usually obtained through a high degree of effort over more easily obtained rewards when free access is given to both (however, see Vasconcelos, Uruioli, & Lionello-DeNolf, 2007). Understanding the relationship between such findings and those we have reported here presents an important challenge. However, it should be noted that this challenge is not isolated to our own work: It is not yet clear how the reported preference for effort-associated rewards fits with the general finding, from a much larger body of research, that animals tend to favor low-effort routes to reward. It is also worth noting that Clement and Zentall (2002) have demonstrated analogous effects when delay is substituted for effort, meaning that the interpretive challenge applies not only to effort discounting, but also to the even more firmly established construct of delay discounting.

The work we have reported here raises a number of other issues for further investigation. For example, would the same pattern of NAcc activation hold in the case of physical, rather than mental, effort, and would the same relationship be observed between NAcc and ACC activation? Would the pattern change with alterations in NAcc dopamine concentrations, which have been shown to affect breakpoints in effort-based decision making (see Figure 1)? The experimental paradigm we employed here could also be used to investigate individual differences relevant to effort and reward, including differences in intrinsic motivation, industriousness, and need for cognition. Another dimension of interest from the point of view of individual differences is personal wealth, which has been shown to affect reward processing at both behavioral and neural levels (Tobler, Fletcher, Bullmore, & Schultz, 2007). It would be interesting to know whether this factor also affects effort discounting, as intuition suggests it might. Finally, it could be productive to investigate whether the relationship between effort and reward, as reflected in NAcc activation, is altered in clinical conditions characterized by changes in motivation, such as major de-

pression and bipolar disorder, as some parallel research (e.g., Salamone, Correa, Mingote, Weber, & Farrar, 2006) has already suggested.

Before closing, it is important to acknowledge some weaknesses of the present study that could be addressed in later experiments. First, although it was not statistically significant, a trend toward an overall difference in RTs was observed between the control group and the experimental group. As discussed above, such a difference has limited impact on the interpretation of the data, since the difference in RT between high- and low-demand blocks was comparable between groups. Nonetheless, it would be desirable to see the central neuroimaging result replicated in the context of more precisely matched behavioral conditions. Second, the results of our whole-brain analyses indicated insufficient power to investigate the effort-discounting effect in a fully exploratory fashion. Methods for accessing this effect with greater statistical power would offer a clear advantage in further study of the relationship between effort and reward processing.

AUTHOR NOTE

This study was funded by Grant P50 MH062196 from the National Institute of Mental Health to the first author. The data for this study were collected when all three authors were at the University of Pennsylvania; data analysis was conducted after all three authors moved to Princeton University. Correspondence concerning this article should be addressed to M. M. Botvinick, 3-S-13 Green Hall, Princeton University, Princeton, NJ 08540 (e-mail: matthewb@princeton.edu).

REFERENCES

- BAXTER, M. G., & MURRAY, E. A. (2002). The amygdala and reward. *Nature Reviews Neuroscience*, *3*, 563-573.
- BOTVINICK, M. M. (2007). Conflict monitoring and decision making: Reconciling two perspectives on anterior cingulate function. *Cognitive, Affective, & Behavioral Neuroscience*, *7*, 356-366.
- BOTVINICK, M. M., BRAVER, T. S., BARCH, D. M., CARTER, C. S., & COHEN, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*, 624-652.
- BOTVINICK, M. M., COHEN, J. D., & CARTER, C. S. (2004). Conflict monitoring and anterior cingulate cortex: An update. *Trends in Cognitive Sciences*, *8*, 539-546.
- BOTVINICK, M. [M.], NYSTROM, L. E., FISSELL, K., CARTER, C. S., & COHEN, J. D. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature*, *402*, 179-181.
- BREITER, H. C., AHARON, I., KAHNEMAN, D., DALE, A., & SHIZGAL, P. (2001). Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron*, *30*, 619-639.
- BROSNAN, S. F., & DE WAAL, F. B. M. (2003). Monkeys reject unequal pay. *Nature*, *425*, 297-299.
- BROWN, J. W., & BRAVER, T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science*, *307*, 1118-1121.
- BUSH, G., VOGT, B. A., HOLMES, J., DALE, A. M., GREVE, D., JENIKE, M. A., & ROSEN, B. R. (2002). Dorsal anterior cingulate cortex: A role in reward-based decision making. *Proceedings of the National Academy of Sciences*, *99*, 523-528.
- CLEMENT, T. S., FELTUS, J. R., KAISER, D. H., & ZENTALL, T. R. (2000). "Work ethic" in pigeons: Reward value is directly related to the effort or time required to obtain the reward. *Psychonomic Bulletin & Review*, *7*, 100-106.
- CLEMENT, T. S., & ZENTALL, T. R. (2002). Second-order contrast based on the expectation of effort and reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes*, *28*, 64-74.
- COX, R. W. (1996). AFNI: Software for analysis and visualization of

- functional magnetic resonance neuroimages. *Computers & Biomedical Research*, **29**, 162-173.
- CROXSON, P. L., JOHANSEN-BERG, H., BEHRENS, T. E. J., ROBSON, M. D., PINSK, M. A., GROSS, C. G., ET AL. (2005). Quantitative investigation of connections of the prefrontal cortex in the human and macaque using probabilistic diffusion tractography. *Journal of Neuroscience*, **25**, 8854-8866.
- CROXSON, P. L., WALTON, M. E., & RUSHWORTH, M. F. S. (in press). The basis of effort and delay anticipation in the human brain. *NeuroImage*.
- DAVIS, K. D., TAYLOR, K. S., HUTCHISON, W. D., DOSTROVSKY, J. O., MCANDREWS, M. P., RICHTER, E. O., & LOZANO, A. M. (2005). Human anterior cingulate cortex neurons encode cognitive and emotional demands. *Journal of Neuroscience*, **25**, 8402-8406.
- DELGADO, M. R., NYSTROM, L. E., FISSELL, C., NOLL, D. C., & FIEZ, J. A. (2000). Tracking the hemodynamic responses to reward and punishment in the striatum. *Journal of Neurophysiology*, **84**, 3072-3077.
- DENK, F., WALTON, M. E., JENNINGS, K. A., SHARP, T., RUSHWORTH, M. F. S., & BANNERMAN, D. M. (2005). Differential involvement of serotonin and dopamine systems in cost-benefit decisions about delay or effort. *Psychopharmacology*, **179**, 587-596.
- ELLIOTT, R., FRISTON, K. J., & DOLAN, R. J. (2000). Dissociable neural responses in human reward systems. *Journal of Neuroscience*, **20**, 6159-6165.
- FLORESCO, S. B., & GHODS-SHARIFI, S. (2007). Amygdala-prefrontal cortical circuitry regulates effort-based decision making. *Cerebral Cortex*, **17**, 251-260.
- FRIEDRICH, A. M., & ZENTALL, T. R. (2004). Pigeons shift their preference toward locations of food that take more effort to obtain. *Behavioural Processes*, **67**, 405-415.
- KACELNIK, A., & MARSH, B. (2002). Cost can increase preference in starlings. *Animal Behaviour*, **63**, 245-250.
- KERNS, J. G., COHEN, J. D., MACDONALD, A. W., III, CHO, R. Y., STENGER, A., & CARTER, C. S. (2004). Anterior cingulate conflict monitoring and adjustments in control. *Science*, **303**, 1023-1026.
- KIVETZ, R. (2003). The effects of effort and intrinsic motivation on risky choice. *Marketing Science*, **22**, 477-502.
- KNUTSON, B., & COOPER, J. C. (2005). Functional magnetic resonance imaging of reward prediction. *Current Opinion in Neurology*, **18**, 411-417.
- KNUTSON, B., FONG, G. W., ADAMS, C. M., VARNER, J. L., & HOMMER, D. (2001). Dissociation of reward anticipation and outcome with event-related fMRI. *NeuroReport*, **12**, 3683-3687.
- KNUTSON, B., TAYLOR, J., KAUFMAN, M., PETERSON, R., & GLOVER, G. (2005). Distributed neural representation of expected value. *Journal of Neuroscience*, **25**, 4806-4812.
- KUNISHIO, K., & HABER, S. N. (1994). Primate cingulo-striatal projection: Limbic striatal versus sensorimotor striatal input. *Journal of Comparative Neurology*, **350**, 337-356.
- LOCKE, J. (1987). *The second treatise of government*. London: Collier MacMillan. (Original work published 1690)
- LUKS, T. L., SIMPSON, G. V., FEIWELL, R. J., & MILLER, W. L. (2002). Evidence for anterior cingulate cortex involvement in monitoring preparatory attentional set. *NeuroImage*, **17**, 792-802.
- MCCLURE, S. M., YORK, M. K., & MONTAGUE, P. R. (2004). The neural substrates of reward processing in humans: The modern role of fMRI. *Neuroscientist*, **10**, 260-268.
- MILLER, E. K., & COHEN, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, **24**, 167-202.
- MONSELL, S. (2003). Task switching. *Trends in Cognitive Sciences*, **7**, 134-140.
- NACCACHE, L., DEHAENE, S., COHEN, L., HABERT, M.-O., GUICHART-GOMEZ, E., GALANAUDE, D., & WILLER, J.-C. (2005). Effortless control: Executive attention and conscious feeling of mental effort are dissociable. *Neuropsychologia*, **43**, 1318-1328.
- NICOLA, S. M., YUN, I. A., WAKABAYASHI, K. T., & FIELDS, H. L. (2004). Firing of nucleus accumbens neurons during the consummatory phase of a discriminative stimulus task depends on previous reward predictive cues. *Journal of Neurophysiology*, **91**, 1866-1882.
- O'DOHERTY, J. P. (2004). Reward representations and reward-related learning in the human brain: Insights from neuroimaging. *Current Opinion in Neurobiology*, **14**, 769-776.
- O'DOHERTY, J. P., DAYAN, P., FRISTON, K., CRITCHLEY, H., & DOLAN, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron*, **38**, 329-337.
- O'DOHERTY, J. P., KRINGELBACH, M. L., ROLLS, E. T., HORNAK, J., & ANDREWS, C. (2001). Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature Neuroscience*, **4**, 95-102.
- PHILLIPS, P. E. M., WALTON, M. E., & JHOU, T. C. (2007). Calculating utility: Preclinical evidence for cost-benefit analysis by mesolimbic dopamine. *Psychopharmacology*, **191**, 483-495.
- RIDDERINKHOF, K. R., ULLSPERGER, M., CRONE, E. A., & NIEUWENHUIS, S. (2004). The role of the medial frontal cortex in cognitive control. *Science*, **306**, 443-447.
- ROESCH, M. R., TAYLOR, A. R., & SCHOENBAUM, G. (2006). Encoding of time-discounted rewards in orbitofrontal cortex is independent of value representation. *Neuron*, **51**, 509-520.
- ROGERS, R. D., RAMNANI, N., MACKAY, C., WILSON, J. L., JEZZARD, P., CARTER, C. S., & SMITH, S. M. (2004). Distinct portions of anterior cingulate cortex and medial prefrontal cortex are activated by reward processing in separable phases of decision-making cognition. *Biological Psychiatry*, **55**, 594-602.
- ROLLS, E. T. (2004). The functions of the orbitofrontal cortex. *Brain & Cognition*, **55**, 11-29.
- RUDEBECK, P. H., WALTON, M. E., SMYTH, A. N., BANNERMAN, D. M., & RUSHWORTH, M. F. S. (2006). Separate neural pathways process different decision costs. *Nature Neuroscience*, **9**, 1161-1168.
- RUSHWORTH, M. F. S., WALTON, M. E., KENNERLEY, S. W., & BANNERMAN, D. M. (2004). Action sets and decisions in the medial frontal cortex. *Trends in Cognitive Sciences*, **8**, 410-417.
- SALAMONE, J. D., CORREA, M., FARRAR, A., & MINGOTE, S. M. (2007). Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. *Psychopharmacology*, **191**, 461-482.
- SALAMONE, J. D., CORREA, M., MINGOTE, S. M., & WEBER, S. M. (2003). Nucleus accumbens dopamine and the regulation of effort in food-seeking behavior: Implications for studies of natural motivation, psychiatry, and drug abuse. *Journal of Pharmacology & Experimental Therapeutics*, **305**, 1-8.
- SALAMONE, J. D., CORREA, M., MINGOTE, S. M., WEBER, S. M., & FARRAR, A. M. (2006). Nucleus accumbens dopamine and the forebrain circuitry involved in behavioral activation and effort-related decision making: Implications for understanding anergia and psychomotor slowing in depression. *Current Psychiatry Reviews*, **2**, 267-280.
- SALAMONE, J. D., COUSINS, M. S., & BUCHER, S. (1994). Anhedonia or anergia? Effects of haloperidol and nucleus accumbens dopamine depletion on instrumental response selection in a T-maze cost/benefit procedure. *Behavioural Brain Research*, **65**, 221-229.
- SCHULTZ, W., TREMBLAY, L., & HOLLERMAN, J. R. (2000). Reward processing in primate orbitofrontal cortex and basal ganglia. *Cerebral Cortex*, **10**, 272-283.
- SCHWEIMER, J., & HAUBER, W. (2006). Dopamine D1 receptors in the anterior cingulate cortex regulate effort-based decision making. *Learning & Memory*, **13**, 777-782.
- SCHWEIMER, J., SAFT, S., & HAUBER, W. (2005). Involvement of catecholamine neurotransmission in the rat anterior cingulate in effort-related decision making. *Behavioral Neuroscience*, **119**, 1687-1692.
- SHIDARA, M., & RICHMOND, B. J. (2002). Anterior cingulate: Single neuronal signals related to degree of reward expectancy. *Science*, **296**, 1709-1711.
- SOLOMON, R. L. (1948). The influence of work on behavior. *Psychological Bulletin*, **45**, 1-40.
- STEVENS, J. R., ROSATI, A. G., ROSS, K. R., & HAUSER, M. D. (2005). Will travel for food: Spatial discounting in two new world monkeys. *Current Biology*, **15**, 1855-1860.
- TOBLER, P. N., FLETCHER, P. C., BULLMORE, E. T., & SCHULTZ, W. (2007). Learning-related human brain activations reflecting individual finances. *Neuron*, **54**, 167-175.
- TSUNEMATSU, S. (2001). Effort- and time-cost effects on demand curves for food by pigeons under short session closed economies. *Behavioural Processes*, **53**, 47-56.
- TVERSKY, A., & KAHNEMAN, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk & Uncertainty*, **5**, 297-323.

- VASCONCELOS, M., URCUIOLI, J. P., & LIONELLO-DE NOLF, K. M. (2007). Failure to replicate the "work ethic" effect in pigeons. *Journal of Experimental Analysis of Behavior*, **87**, 383-399.
- WALSTER, E. H., WALSTER, G. W., & BERSCHIED, E. (1978). *Equity: Theory and research*. Boston: Allyn & Bacon.
- WALTON, M. E., BANNERMAN, D. M., ALTERESCU, K., & RUSHWORTH, M. F. S. (2003). Functional specialization within medial frontal cortex of the anterior cingulate for evaluating effort-related decisions. *Journal of Neuroscience*, **23**, 6475-6479.
- WALTON, M. E., BANNERMAN, D. M., & RUSHWORTH, M. F. S. (2002). The role of rat medial frontal cortex in effort-based decision making. *Journal of Neuroscience*, **22**, 10996-11003.
- WALTON, M. E., KENNERLEY, S. W., BANNERMAN, D. M., PHILLIPS, P. E. M., & RUSHWORTH, M. F. S. (2006). Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Networks*, **19**, 1302-1314.
- ZIPE, G. K. (1949). *Human behavior and the principle of least effort*. Cambridge, MA: Addison-Wesley Press.

NOTES

1. As detailed in the Analysis section, the core analyses for the experimental group tested for main effects within this group, whereas the control group was intended mainly to allow testing for difference in the effect of block type between groups. On this basis, we reasoned that a larger sample was needed for the experimental than for the control group.
2. Note that this procedure had the result that approximately 72% of the participants received greater pay than would have been the case under a random schedule of rewards. The mild deception involved in the experiment was reviewed and approved by the University of Pennsylvania's institutional review board.
3. The control group was paid at the standard rate for fMRI studies at the University of Pennsylvania. Although an alternative approach might have been to pay the control participants the mean amount from the range that the experimental participants were led to expect (\$30), this seemed unlikely to yield different results. It is worth reiterating that neither payment nor rewards in our study were intended as incentives to

effort, and neither behavioral nor fMRI data indicated that they operated as such.

4. This covariate was included as a conservative measure. Note that a central objective of these ANCOVAs was to rule out the possibility that effort-discounting effects might reflect a simple carryover of activation from the preceding task blocks. Such a carryover effect would presumably vary in size with the difference in ROI activation between high- and low-demand blocks. Thus, including this difference as a covariate increased the test's sensitivity, maximizing our ability to detect a carryover effect.

5. Behavioral data for 1 control participant were unavailable.

6. Averaging across demand conditions, the \$ cue activated the NAcc more than did the neutral cues used in the control group [$t(34) = -2.85$, $p = .007$]; the X cue did not [$t(34) = -0.54$, $p = .59$].

7. This prediction holds in the case where all gross reward magnitudes are greater than or equal to 0, as was the case in our experiment.

8. Note that correction for multiple comparisons was not made in the analyses in this section. Our objective was to establish the distinctiveness of the pattern of activation in the NAcc. Thus, it was most conservative to adopt a low threshold for rejecting the null hypothesis in other regions.

9. For clarity, we emphasize that this contrast was based on ACC activation at the time of reward receipt. This differs from the contrast reported earlier, under the Validation of Effort Manipulation section, in which we focused on ACC activation during block performance.

10. Consistent with our finding of effort discounting for both reward outcomes, and as predicted by effort-discounting theory (see Figure 1, bottom panel), the inverse correlation was statistically significant for both \$ cues ($p = .004$) and X cues ($p = .004$).

11. This correlation across participants might have resulted from simple individual differences in the amplitude of the BOLD response, unrelated to underlying neural processes. In order to rule out this explanation, we repeated the analysis, computing the correlation after partialling out the difference, for each participant, between the \$ and X responses in the NAcc. This did nothing to diminish the strength of the correlation ($r = .65$; $p = .001$).

(Manuscript received December 27, 2007;
revision accepted for publication July 15, 2008.)