Neural systems of cognitive demand avoidance

Ceyda Sayalı and David Badre[1,2]

[1]Department of Cognitive, Linguistic, and Psychological Sciences
[2]Brown Institute for Brain Sciences
Brown University

Address correspondence to:
Ceyda Sayalı
Box 1821
Brown University
Providence, RI 02912-1978
tel 401-863-1544, fax 401-863-2255
email ceyda_sayali@brown.edu

Abstract

Cognitive effort is typically aversive, evident in people's tendency to avoid cognitively demanding tasks. The 'cost of control' hypothesis suggests that engagement of cognitive control systems of the brain makes a task costly and the currency of that cost is a reduction in anticipated rewards. However, prior studies have relied on binary hard versus easy task subtractions to manipulate cognitive effort and so have not tested this hypothesis in "dose-response" fashion. In a sample of 50 participants, we parametrically manipulated the level of effort during fMRI scanning by systematically increasing cognitive control demands during a demand-selection paradigm over six levels. As expected, frontoparietal control network (FPN) activity increased, and reward network activity decreased, as control demands increased across tasks. Further, greater activity in this network during a task was correlated with avoidance of that task. However, this relationship with avoidance was not attributable to the change in FPN activity that was due to the cognitive control manipulation, lending only partial support to the cost of control hypothesis. By contrast, we unexpectedly observed that the inhibition of a task-negative brain network corresponding to the Default Mode Network (DMN) across levels of the cognitive control manipulation predicted the change in avoidance. In summary, we find partial support for the cost of control hypothesis, while highlighting the role of task-negative brain networks in modulating effort avoidance behavior.

Keywords: Cognitive effort, effort avoidance, cognitive control, effort cost, reward

Highlights

- Reward network of the brain shows decreasing activity with increasing effort requirements.
- Frontoparietal control network of the brain shows increasing activity with increasing effort requirements.
- Neither reward network nor the frontoparietal control network predicts demand avoidance behavior
- Default Mode Network predicts demand avoidance behavior in those who avoid effort.

Introduction

Cognitive effort influences our everyday decisions about whether to perform challenging mental tasks. Most people prefer less cognitively effortful tasks when given a choice, a phenomenon known as ‘demand avoidance’ (Kool et al., 2010). However, it is not yet settled what cognitive and neural mechanisms underlie this avoidance behavior, or whether the deployment of these mechanisms varies among individuals.

One account of demand avoidance behavior is the ‘cost of effort' hypothesis, according to which the brain codes cognitive effort as disutility (Botvinick, 2007). Consistent with this hypothesis, the value of reward is discounted as a function of effort requirements (Westbrook et al., 2013). In the brain, this hypothesis predicts that this cost should be computed by the same networks that typically process reward, such as the mesolimbic dopaminergic system, including the medial prefrontal cortex and the ventral striatum (VS). This prediction is supported by at least one fMRI study which observed reduced activation in VS following effortful task performance (Botvinick et al., 2009). However, the engagement of these systems may be affected by whether one is performing the effortful task or selecting whether to perform it. For example, Schouppe et al. (2014) found that VS increased, rather than decreased, its activity during the selection of the effortful task. Thus, evidence that the brain registers effort as disutility has received limited, albeit conflicting, support.

A second question concerns what makes a task cognitively effortful. An influential hypothesis proposes that tasks are effortful to the degree that they recruit cognitive control. Thus, according to the “cost of control” hypothesis, cognitive control might be recruited to the degree that its associated costs do not exceed its anticipated benefits (Shenhav et al, 2013). Consistent with this hypothesis, people avoid tasks that require more task switching (Kool et al., 2010),

greater working memory load (Westbrook et al., 2013), and greater response conflict (Schouppe et al., 2014), all of which putatively require greater cognitive control.

Given the general association of cognitive control with a dorsal fronto-parietal network (FPN) in the brain (Fedorenko et al., 2013; Niendam et al., 2012; Badre & D'Esposito, 2009; Vincent et al., 2008; Cole & Schneider, 2007; Dosenbach et al., 2007), a reasonable prediction is that engagement of FPN by a task will be associated with a tendency to avoid that task. At least one study has observed that participants who tended to avoid demanding tasks also showed increased FPN activity during effortful task performance (McGuire & Botvinick, 2010). However, this study contrasted only two effort levels and so had limited sensitivity to test whether cognitive control demands were associated, in a dose-response fashion, with changes in brain-behavior relationships within-subject. Further, this study focused on univariate activation change in FPN, but other evidence suggests that functional connectivity with the FPN may be important for cognitive effort (e.g., Ray et al., 2017).

Accordingly, in the present study, we sought to provide a more sensitive test of the brain-behavior relationships predicted by the *cost-of-control* hypothesis. Specifically, we developed a parametric version of the DST paradigm for fMRI to sensitively test how incremental changes in activity and connectivity levels in FPN, the reward network, and/or other systems related to demand avoidance. Further, we addressed how these within-subjects effects were modulated by individual differences in demand avoidance versus demand seeking behavior. Using this approach, we observed only qualified support for the cost for control hypothesis, in that though FPN activity was related to effort avoidance, this relationship was not clearly mediated by the cognitive control manipulation, as opposed to other factors. However, we discovered that effort-

related modulation of regions corresponding to the "default mode network" were associated with effort-related changes in effort avoidance.

Methods

1.1.Participants

Twenty-eight right-handed adults (aged 18-35; 14 female) with normal or corrected to-normal vision were recruited for the behavioral demand avoidance experiment (Experiment 1). Fifty-six right-handed adults (aged 18-35; 26 female) with normal or corrected to-normal vision were recruited for the fMRI experiment (Experiment 2). Two participants withdrew from the fMRI study and did not complete all experimental phases. Two fMRI participants' data were excluded due to an operator error that resulted in a failure to record behavioral data during scanning. One participant reported during debriefing that they failed to follow instructions. One participant was excluded due to head movement greater than our voxel size across all sessions. So, in total, data from six participants in Experiment 2 were excluded prior to analysis of the fMRI data. Thus, 50 participants were included in the behavioral and fMRI analyses of Experiment 2. All participants were free of neurological or psychiatric conditions, were not taking drugs affecting the central nervous system, and were screened for contraindications for MRI. Participants provided informed consent in accordance with the Research Protections Office at Brown University.

1.2.Behavioral Task

In both the Experiment 1 and 2, participants performed a parametric variant of the demand selection task (DST; Fig 1). In an initial Learning phase, the participant associated virtual card "decks", denoted by particular symbols, with a certain effort level through

experience with trials drawn from that deck. Then, in a second Test phase, participants chose which deck they wished to perform (Selection epoch) and then performed the trials drawn from that deck (Execution epoch). A distinguishing characteristic of this version of the DST relative to prior versions is that we varied the effort level parametrically over six levels based on the frequency of task switching required by a deck. This allowed us to characterize the functional form of changes in behavior or fMRI signal due to our cognitive control manipulation. We now provide detail on each of these phases of the behavioral task.

1.2.1. The Task "Decks"

Throughout all phases of the experiment, participants performed trials drawn from virtual decks (Fig 1a). Blocks consisted of sequences of 13 consecutive trials drawn from a given deck. On each trial, the participant categorized a presented digit as either odd/even (parity judgment) or greater/less than 5 (magnitude judgment). The color of a circle (green or blue) surrounding the digit cued whether to perform the parity or magnitude judgment on each trial based on a pre-instructed mapping. The participant indicated their categorization by pressing the left or the right arrow key on the keyboard. The response deadline was 1.5 sec. Trials were separated by .2 sec. Response mappings and color-to-task mappings were counterbalanced across participants. Digits 1-4 and 6-9 were used with equal frequency across both tasks and were randomized for order of presentation.

In order to manipulate cognitive effort across decks, we varied the frequency of task switching required by a particular deck (Fig 1b). The probability of switching tasks from trial to trial within a sequence drawn from each deck increased across 6 effort levels: 8%, 23%, 38%, 54%, 69%, 85%. A higher frequency of task switching is associated with a higher experience of cognitive effort (Monsell, 2003), and has been shown to be aversive (Arrington & Logan, 2004).

At the beginning of each block, a shape was presented for 1 sec to indicate which deck was being performed, and this shape was also tiled in the background throughout performance of the block. Participants were told that this shape represented the symbol on the back of the virtual card deck from which trials for that sequence were drawn. Each effort level was associated with a particular deck symbol. Participants were not told about this relationship, but could learn through experience that certain decks were more effortful than others to perform. The mappings between deck symbols and effort levels were randomized across participants.

1.2.2. Practice Phase

During an initial "Practice Phase" participants gained familiarity with the trial structure, the categorization decisions, and the color and response mappings. After being given instructions regarding the categorization and response rules (Fig 1a), they practiced two 13-trial blocks. Each task sequence included the presentation of a deck symbol, and the subsequent performance of 13 consecutive task-switching trials. In the first block of the Practice Phase, feedback was presented after the button press as either 'Correct' in pink, 'Incorrect' in yellow, or 'Please respond faster' if the participant failed to answer in 1.5s. In the second block of the Practice phase, feedback was omitted as would be the case in the actual experiment. The deck symbols that participants encountered in the Practice Phase were for practice only and were not presented during the learning or test phases to avoid any association of error feedback with a particular deck symbol.

1.2.3. Learning Phase

In the Learning Phase (Fig. 1c), participants learned the association between 6 deck symbols and an effort level. Each deck was performed 15 times in random order. In both experiments, this phase was performed in a behavioral testing room outside the magnet.

1.2.4. Test Phase

In the Test Phase (Fig 1d), two decks were presented and the participant chose which they would like to perform (Selection epoch). The participants were told to freely choose between decks prior to a 3 sec deadline. We note that in contrast to other DST variants (Gold et al., 2014), participants in this task were not told about the difficulty manipulation to avoid biasing choices based on participants' explicit beliefs about effort. Once the participant made their selection, the selected deck turned blue and both options remained on the screen until the end of the 3 sec deadline. In the event of a non-response, the same choice pair was re-presented at the end of the entire experiment until they made a decision. Each pair was selected from the set of fifteen unique (un-ordered) pair combinations of all six decks, excluding self-pairing (e.g., deck #1 paired with deck #1). Each deck was presented either on the left or the right side of the screen, counterbalanced for location across trials. The Selection epoch was followed by the execution of the selected effort level task sequence (Execution epoch). The sequence of events in this epoch was the same as during the Learning phase.

In the fMRI study, only the Test phase was performed in the MRI scanner, and participants used a MRI-compatible button box to indicate their decisions. The Execution trials were optimized as blocks. The Selection and Execution events were separated in time by a jittered time interval (mean 2 secs) so that signal related to each could be analyzed independently. The Test phase was separated into four, approximately 15 minute-long scanning runs. In each run, each pair was presented 3 times in a pseudo-randomly intermixed order, making a total of 180 decision trials across 4 blocks.

2.3. Behavioral Data Analysis

Trials with response times below 200 ms were excluded from further analysis. Execution trials on which participants missed the response deadline were also excluded from further

analysis (approximately %1 of execution trials in both phases). Response times were calculated using only correct trials.

Data were analyzed using a mixed-design analysis of variance (ANOVA) (within subject factor: Effort, between subject factor: Avoidance group). If the sphericity assumption was violated, Greenhouse-Geisser correction was used. Significant interactions were followed by simple effects analysis, the results of which are presented with False Detection Rate (FDR) correction. Alpha level of .05 was used for all analyses. Error bars in all figs stand for within-subject error.

Choice behavior was assessed by calculating the probability of selecting an effort level across all selection trials on which that effort level was presented as an option during the Test Phase. The decision difference analyses included the calculation of the choice probability and the decision time to select the easier task across all decisions with the same difference in difficulty levels between effort options in the selection epoch of the Test Phase. For example, choice probability associated with a difficulty difference of 1 would be computed by averaging the probability of choosing the easier task across all choice pairs that differed by 1 effort level (i.e., 1 vs 2, 2 vs 3, 3 vs 4, 4 vs 5 and 5 vs 6).

2.3.1. Permutation analysis of choice behavior

In order to appropriately test for changes in behavioral selection rates across effort levels while also defining group membership (i.e., demand avoiders or demand seekers) based on the same choice behavior, we conducted a permutation procedure. First, we estimated likelihood distributions of the slope change in selection rates across effort levels in two different choice-defined groups. Specifically, for each of 10,000 iterations, we conducted the following procedure. (1) We selected a random half of each participant's data and calculated an overall rate

of choosing the harder task from those data. (2) We assigned each participant to either the Demand Seeker or Demand Avoider group depending on whether their selection rate was greater or less than 50% respectively. (3) Using the other half of each participant's data, we fit a linear function to selection rates across all effort levels, and extracted the slope (β). (4) We averaged the β values for each group separately. Over 10,000 iterations for each group, this procedure yielded likelihood distributions of selection rate slopes across effort levels for Demand Avoiders and Demand Seekers.

Next, we estimated a null distribution using the following procedure over 10,000 iterations: (1) We randomly drew a set of participants in the same number as were assigned to choice-defined groups. (2) Using a random half of each participant's data, we fit a linear function to selection rates across all effort levels, and extracted the slope (β). (3) We averaged the β values in this group. Over 10,000 iterations, this procedure produced a permuted null distribution of avoidance rate slopes across effort levels. We fit the likelihood and null histograms with a normal distribution to assess their overlap and assess statistical significance.

2.4. *MRI procedure*

In the fMRI experiment (Experiment 2), whole-brain imaging was performed with a Siemens 3T Prisma MRI system using a 64-channel head coil. A high-resolution T1-weighted 3D multi-echo MPRAGE image was collected from each participant for anatomical visualization. Each of the four runs of the experimental task involved around between 450 and 660 functional volumes depending on the participant's response time, with a fat-saturated gradient-echo echo-planar sequence (TR = 2s, TE=28ms, flip angle = 90°, 38 interleaved axial slices, 192 mm FOV with voxel size of 3x3x3 mm). Head motion was restricted with padding, visual stimuli were rear projected and viewed with a mirror attached to the head coil.

2.5. *fMRI Analysis*

Functional images were preprocessed in SPM12 (http://www.fil.ion.ucl.ac.uk/spm). Before preprocessing, data were inspected for artifacts and variance in global signal (tsdiffana, art_global,art_movie). Functional data were corrected for differences in slice acquisition timing by resampling slices to match the first slice. Next, functional data were realigned (corrected for motion) using B-spline interpolation and referenced to the mean functional image. 1-2 sessions were excluded from 3 other participants prior to behavioral analysis due to movement during data collection in the scanner. Functional and structural images were normalized to Montreal Neurological Institute (MNI) stereotaxic space using affine regularization followed by a nonlinear transformation based on a cosine basis set, and then resampled into 2x2x2 mm voxels using trilinear interpolation. Lastly, images were spatially smoothed with an 8 mm full-width at half-maximum isotropic Gaussian kernel.

A temporal high-pass filter of 128 (.0078 Hz) was applied to our functional data in order to remove noise. Changes in MR signal were modeled under assumptions of the general linear model (GLM). Two GLMs were devised: a linear effort-level GLM and an independent effort-level GLM. Both GLMs included nuisance regressors for the six motion parameters (x,y,z,pitch,roll,yaw) and four run regressors for the 'Linear Effort-Level GLM' and one run regressor for the 'Independent Effort Level GLM'. In both GLMs, we modeled the Selection and the Execution epochs separately, though, our results focus on the Execution epoch only, as we specifically aimed to test the cost of control hypothesis. Results from the Selection epoch are the topic of a different report.

2.5.1. Linear Effort-Level GLM

The linear effort-level GLM tested which voxels in the brain parametrically increased or decreased linearly with effort level. Two event regressors were used. Execution events were modeled as a boxcar that onset with the presentation of the first trial stimulus of the sequence and ended with the participant's response to the final item. Thus, the duration of this boxcar varied as a function of response time. Second, the Selection event regressor modeled each selection choice event with a fixed boxcar of three secs. We used parametric modulators on these event regressors to test the linear effect of effort level. The Execution event regressor was modulated by an Effort Level parametric regressor corresponding to the effort level of that task sequence (1 through 6). The Selection event regressor was modulated by (a) an Effort Level parametric regressor based on the chosen effort level (1 through 6), and (b) a Difference regressor that scaled with the difference between the chosen and the unchosen effort option on that selection trial (1 through 5). Note that as implemented by SPM, each parametric regressor includes only the variance unique to it and shared with those ordered after it. Thus, for example, the Effort Level regressor includes variance explained over and above that shared with the Execution event regressor. The Difference regressor did not yield statistically reliable results and is not discussed further. The Execution and Selection event regressors, along with their parametric modulators, were modeled separately for each scanning run within the GLM.

2.5.2. Independent Effort Level GLM

The independent effort level GLM sought to characterize the signal change related to each effort level independently of each other or of any particular function (e.g., linear). This GLM included twelve event regressors, one for each effort level (1 through 6) by epoch (Execution and Selection). Events in the Execution regressors were modeled as boxcars that onset with the presentation of the first trial stimulus of the sequence and ended with the

participant's response to the final item. Events in the Selection regressors were modeled with an 3 sec boxcar at the onset of each choice pair. The Selection onset regressor was further modulated by a parametric Difference regressor that scaled with the difference between the chosen and the unchosen effort option on that selection trial. Boxcars and durations were the same as in the linear effort level model. In this GLM, four run epochs and a linear drift over the whole experiment were included as nuisance regressors.

For both GLMs, SPM-generated regressors were created by convolving onset boxcars and parametric functions with the canonical hemodynamic response (HRF) function and the temporal derivative of the HRF. Beta weights for each regressor were estimated in a first-level, subject-specific fixed-effects model. For group analysis, the subject-specific beta estimates were analyzed with subject treated as a random effect. At each voxel, a one-sample t-test against a contrast value of zero gave us our estimate of statistical reliability. For whole brain analysis, we corrected for multiple comparison using cluster correction, with a cluster forming threshold of $\alpha = .001$ and an extent threshold, k, calculated with SPM to set a family-wise error cluster level corrected threshold of $p<.05$ for each contrast and group (See Table 1). Note that the higher cluster forming threshold helps avoid violation of parametric assumptions such that the rate of false positive is appropriate (Eklund et al., 2016; Flandin & Friston, 2016).

2.5.3. ROI analysis

ROI definition is described below. For each ROI, a mean time course was extracted using the MarsBar toolbox (http://marsbar.sourceforge.net/). The GLM design was estimated against this mean time series, yielding parameter estimates (beta weights) for the entire ROI for each regressor in the design matrix.

2.5.3.1. Independent ROIs

In order to test the cost of control hypothesis, we defined a fronto-parietal control network ROI defined from a functionally neutral group ([Network 12] Yeo et al., 2011) along with a VS ROI (Badre et al., 2014), as VS activity has been consistently observed in cognitive effort literature (Botvinick et al., 2009; Schouppe et al., 2014). We also included a Default Mode Network ROI from a functionally neutral group ([Network 16] Yeo et al., 2011).

2.5.3.2. PCA ROIs

In order to have an unbiased analysis of the neural response function across effort levels, we adopted a data-driven Principal Component Analysis (PCA) approach on the whole brain (126,866 voxels) to explore the shape of the neural functions of different brain areas that respond to effort execution, and their relation to demand avoidance behavior. We ran PCA over averaged β estimates of each voxel in the whole brain as they were derived from the Independent Effort Level GLM task Execution Level 1-6 onset regressors across all participants.

We ranked all voxels as a function of PC weight and extracted the first 3000 voxels that loaded the highest (Positive PC) and the least (Negative PC) on the selected components as separate ROIs (Decot et al., 2017). These ROIs further have been analyzed to yield beta estimates for each Positive and Negative PC. The number of voxels to select (3000) was an arbitrary cut off and was not chosen after looking at the results. However, to ensure that outcomes did not depend on this choice, we also tried different numbers of voxels, such as 500 and 10,000, and did not find a qualitative change in average beta estimates ROIs yield.

2.5.3.3. Functional connectivity analysis

For each participant, functional connectivity analysis was implemented in MATLAB using the CONN toolbox (http://www.nitrc.org/projects/conn; Whitfield-Gabrieli & Nieto-Castanon, 2012). CONN uses the CompCor method (Behzadi et al., 2007) which segments white

matter (WM) and cerebrospinal fluid (CSF), as well as the realignment parameters entered as confounds in the first-level GLM analysis (Behzadi et al., 2007), and the data were band-pass filtered to 0.01 Hz to 0.1 Hz.

We conducted an ROI-to-ROI analysis to test the functional connectivity between and within the first principal component clusters identified in the PC analysis for the entire sample. PC clusters further have been separated into smaller spatially independent clusters in order to conduct within-PC connectivity analysis. Positive PC1 yielded 4, and Negative PC1 yielded 8 spatially independent clusters. Accordingly, 4 Positive PC1 clusters and 8 Negative PC1 clusters were included as ROIs. Fisher's Z-transformed correlations were computed for each ROI pair.

Next, we focused on two connectivity measures. First, we calculated the average connectivity Between Positive-Negative PC1, within Positive PC1 and within Negative PC1 clusters at each effort execution level, in order to observe the connectivity change in these measures with increasing effort execution. Next, we collapsed connectivity coefficients for Between Positive-Negative PC1, within Positive PC1 and within Negative PC1 across effort levels and compared average connectivity rates between demand groups.

2.6. Brain – behavior analysis

In order to understand which brain regions that underlie effort execution predict effortful task selection, all PCs as well as the connectivity measures were related to the selection behavior during the Selection Epoch. We analyzed this relationship in two ways: 1) correlating the change in effort selection rates to the change of brain activity during task execution across effort levels ('Effort level'), 2) correlating individual effort selection rate to the brain activity during task execution at that effort level ('Individual task selection'), unranked with respect to task switching frequency. The former analysis focuses on activity change in a brain region due to the task-

switching manipulation. The latter considers how overall activity levels, independent of a linear task-switching manipulation, correlate with effort avoidance.

For the ‘Effort level’ analysis, we computed the slope of the linear function for the selection rates, β estimates of a PC ROI across effort levels and connectivity coefficients of within and between PCs. We multiplied the slope we obtained from the selection rates with -1, in order to indicate that decreasing likelihood to select more effortful tasks are conceptualized as 'demand avoidance' in our task. Then, we correlated the demand avoidance slopes and the β estimate slopes for all demand avoiders, in order to address the role of this region in explaining individual variability in demand avoidance behavior.

For the 'Individual task selection’ analysis, we correlated the β estimate of an ROI at each effort level execution with the selection rate of the corresponding effort level to derive a β value for each participant. This resulting β value indicates the relationship between selecting an effort level given the estimated activity in this region, without presupposing a linear effort manipulation.

3. Results

3.1. Demand avoidance behavior

Experiment 1 established the basic behavioral effects of our parametric effort manipulation on performance. Overall, participants were highly accurate on the categorization task across both phases (mean error: 12% in the Learning Phase, 15% in the Test Phase), and participants missed the deadline on few trials (2.3% of Learning phase trials, *SE*=0.003, 1.4% of Test phase trials, *SE*=0.01).

Both RT and error rates showed a linear effect of effort condition, increasing with a higher proportion of task switches across both switch and repeat trials. There was a significant effect of Effort on error rates, ($F(2.76,74.53)=29.88$, $p< .001$, $\eta_p^2 =.53$) and RT ($F(2.69,72.71)=131.17$, $p< .001$, $\eta_p^2 =.83$). And, the increase in both was linear over effort levels (errors: $F(1,27)=61.17$, $p< .001$, $\eta_p^2 =.69$; RT: $F(1,27)=248.19$, $p< .001$, $\eta_p^2 =.90$). These increases were also significant for both switch and repeat trials tested separately (switch errors: $F(1,27)=29.73$, $p< .001$, $\eta_p^2 =.52$; repeat errors: $F(1,27)=42.59$, $p< .001$, $\eta_p^2 =.61$; switch RT: $F(1,27)=0.37$, $p < .001$, $\eta_p^2 =.61$; repeat RT: $F(1,27)=42.89.98$, $p< .001$, $\eta_p^2 =.61$).

During the Selection phase, 25 of the 28 participants (89%) selected the easier task more than 50% of the time overall (Fig 2A). Further, there was a significant effect of effort level on choice behavior, ($F(5,135)=7.23$, $p<.001$, $\eta_p^2 =.21$), such that higher effort levels were associated with higher avoidance rates than lower effort levels (Fig 3A). This pattern across effort levels was fit by a linear effect ($F(1,27)=24.54$, $p<.001$, $\eta_p^2 =.48$). The probability of selecting the easier task also significantly changed depending on the difference between effort levels given as options ($F(3,81.09)=6.07$, $p=.001$, $\eta_p^2 =.18$), such that larger differences increased the probability of choosing the easier task. This effect was linear across effort levels ($F(1,27)=16.83$,

p<.001, η_p^2 =.38). Decision time (Fig 3A) was mostly unaffected by effort level with only a marginal omnibus effect of effort level on decision time (F(2.98,80.56)=2.62, p=.06, η_p^2 =.09) and no effect of difficulty difference (Supplemental Results 1).

Experiment 2 applied the same behavioral task as in Experiment 1, but the Test phase was scanned using fMRI. As with Experiment 1, performance on the categorization task was strong overall (mean error: 12%) with relatively few omissions (<1.5% of trials). Also as with Experiment 1, both RT and error rates showed a linear effect of effort, indicating that our effort manipulation was successful in increasing task difficulty (errors: F(1,49)=49.11, p< .001, η_p^2 =.50; RT: F(1,49)=190.48, p< .001, η_p^2 =.80). And, these effects were evident on both Switch and Repeat trials (switch errors: F(1,49)=47.67, p< .001, η_p^2 =.49 ; repeat errors: F(1,49)=20.37, p< .001, η_p^2 =.29; switch RT: F(1,49)=0.36, p = .55, η_p^2 =.01 ; repeat RT: F(1,49)=635.24, p< .001, η_p^2 =.93)

Unlike in Experiment 1, Experiment 2 participants did not consistently avoid the effortful task as a group during the Selection phase. Rather, nearly half of participants (N=24) selected the easier task more than 50% of the time (Fig 2B). This suggests that, as a group, participants were not avoiding the effortful change more than one expects by chance.

Importantly, however, the diagnostic feature of demand avoidance in our paradigm is the linear change in selection rates across effort levels, as observed in Experiment 1. Therefore, if the subgroup with overall selection rates above 50% (Demand Avoiders) are engaged in effort avoidance, we should also see a linear change in their selection rates across effort levels as was observed in Experiment 1. Likewise, those with selection rates below 50% (Demand Seekers) might show the opposite pattern, with a linearly increasing tendency to choose the harder task. However, simply testing the slope across effort levels within each group might be biased by the

way the groups are defined, as the same choice data would be used for both computations. Thus, to test this prediction, we performed a permutation procedure (see section 2.3.1 for details) that allowed us to generate likelihood distributions for the slope of change in selection rates given an overall tendency (>50%) to avoid versus seek the harder task, as well as a null distribution. This procedure computed the selection rate slope and group membership on independent halves of the data over multiple (10,000) random draws from the data set.

Fig 4 plots the overlap of the likelihood distributions for the Demand Avoider and Demand Seeker groups with the permuted null distribution. First, the mean slope was negative for Demand Avoiders and positive for the Demand Seekers, consistent with a respective tendency to avoid versus select a task as function of its effort. Second, these distributions overlapped minimally with the permuted null distribution (Demand avoider: p = .008, Demand seeker: p = .003). Thus, the likelihood of randomly assigning individuals to two groups and obtaining the demand avoidance slopes we observe using a random half of the data is less than 1%.

For Demand avoiders, there was a significant effect of effort level on choice behavior, $F(5,125)=8.24$, $p= .00$. Similar to the pilot behavioral study, they were significantly less likely to choose an effort level with increasing effort requirement; there was a significant linear effect, $F(1,25)=31.46$, $p= .00$. There was a significant effect of effort level on decision time $F(5,125)=3.05$, $p= .01$. The decision time to select an effort level linearly increased with increasing effort levels, $F(1,25)=6.34$, $p= .19$. For Demand Seekers, there was a significant effect of effort level on choice behavior, $F(3.22,73.98)=4.67$, $p= .004$. They were significantly more likely to choose an effort level with increasing effort requirement, there was a significant

linear effect, $F(1,23)=34.20$, $p= .00$. The decision time to choose an effort level was similar ($M=$ 1.00 s., $SD= 0.25$) across all effort levels, $F(2.88,66.25)=2.55$, $p= .07$.

In summary, the behavioral data provide evidence for two groups of participants that differed in their effort choices: Demand Avoiders (Fig 3B), show a decreasing tendency to choose a task as more cognitive effort is required, and Demand Seekers (Fig 3C) exhibited the opposite pattern (see Supplementary Results 1 for additional analysis). We note that based on a behavioral experiment reported in the Supplementary materials (see Supplementary Results 2), this higher tendency to demand seek in the fMRI group versus the behavioral group may be a self-selection bias due to a systematic difference in Need for Cognition among those who volunteer for fMRI. Nevertheless, these behaviorally defined groupings were maintained throughout the subsequent fMRI analyses in order to control for individual differences.

3.2. The functional form of univariate brain activity over effort levels

Figure 5 plots corrected whole brain activation from the main univariate contrasts. During task performance (Execution > Baseline), both groups activated a fronto-parietal network commonly seen during cognitive control tasks and showed deactivated (Baseline > Execution) regions of inferior parietal lobule and precuneus that overlap the "default-mode network" (DMN) as well as regions of posterior cingulate cortex and medial prefrontal cortex commonly associated with reward. There was no significant difference between Demand Avoiders and Demand Seekers in these basic task-positive and task-negative activation patterns.

The frontoparietal control network showed a positive correlation with effort level (Parametric+) and conversely DMN and VS negatively correlated with parametric increases

(Parametric-) in effortful task execution. There were also no groups differences between for the Parametric+ or the Parametric- contrasts (Fig 5C-D; Table 2).

In addition to whole-brain analyses, we investigated the pattern of activation in independent ROIs, in order to examine their underlying neural function across the execution of increasing effort levels. We defined ROIs in the FPN (Network12; Yeo et al, 2011), DMN (Network16; Yeo et al., 2011), and/or reward network (VS; Badre et al., 2014), as described in 2.5.3. Then, using the "Independent Effort Level GLM" (see 2.5.2), we estimated the activation in the network for each effort level separately. These estimates are plotted in Fig 6.

As expected, the FPN ROI showed a linear trend across effort levels (Fig 6B). There was a significant effect of effort on FPN β estimates, $F(5,240)=5.19$, $p< .001$, $\eta_p^2 =.10$. The β estimates linearly increased with increasing effort requirements of the task, $F(1,48)=20.85$, $p< .001$, $\eta_p^2 =.30$. There was no effect of avoidance group on β estimates, $F(1,48)=0.57$, $p= .45$, $\eta_p^2 =.01$.

The DMN showed a negative linear trend in activation across effort levels (Fig 6A). There was a significant effect of effort on DMN β estimates, $F(5,240)=4.18$, $p< .001$, $\eta_p^2 =.08$. The β estimates linearly decreased with increasing effort requirements of the task, $F(1,48)=9.77$, $p= .003$, $\eta_p^2 =.17$. There was no effect of avoidance group on β estimates, $F(1,48)=1.06$, $p= .31$, $\eta_p^2 =.02$.

An ROI in VS also showed a negative linear trend (Fig 6C). There was a significant effect of effort on VS β estimates, $F(3.26,156.52)=5.94$, $p= .001$, $\eta_p^2 =.11$. The β estimates linearly decreased with increasing effort requirements of the task, $F(1,48)=9.40$, $p= .004$, $\eta_p^2 =.16$. There was no effect of avoidance group on β estimates, $F(1,48)=0.32$, $p= .58$, $\eta_p^2 =.01$. We note that in all cases a linear function was the curve of best fit relative to alternatives

(summarized in Supplementary Table S1), with the exception of VS that, though showing a reliable linear function, was better fit by a quadratic due to an apparent asymptote at level 4.

Finally, we conducted a whole brain PCA analysis to ensure that other networks or clusters of voxels do not hold non-linear functions with increasing effort execution that might influence avoidance behavior, but that would not be detected by the linear parametric regressor used in the GLM or by our ROIs. Thus, we adopted a PCA approach to the whole brain analysis. This analysis is used to identify a small set of variables that explain the maximum amount of variance in a large data set. In our case, we used the PCA to identify voxels in the whole brain that share a similar underlying activation functional form across the execution of effort levels.

From 126,866 voxels included in the analysis, the PCA located 3 PCs that explained 90% of the variance of the data. The first PC explained 57% of the variance. The percent variance of the data explained by each subsequent PCs reached asymptote by 4-5 PC, and 5 PCs explained 100% of the total variance (Fig 7). We extracted 3000 voxels that loaded the most positively (Positive PCs) and most negatively (Negative PCs) on each component and extracted their βs in order to observe the functional form of each PC.

Positive PC1 (Fig 8A, Table 3) overlapped the frontoparietal network, and included the left SMA, left superior parietal lobule, left middle frontal gyrus, left inferior frontal gyrus. As we observed in Parametric+ whole-brain analysis, there was a significant effect of effort on FPN cluster β estimates as defined by Positive PC1, $F(3.84,184.18)=19.08$, $p< .001$, $\eta_p^2 =.28$. The β estimates from these voxels linearly increased with increasing effort requirements of the task, $F(1,48)=69.24$, $p< .001$, $\eta_p^2 =.59$, converging with the preceding a priori analyses and indicating that the most voxel-wise variance in effort execution can be explained by a linear trend in the brain regions that contribute to controlled behavior.

Negative PC1 (Fig 8B, Table 4) overlapped with the DMN, and included the medial prefrontal cortex, posterior cingulate cortex, left and right angular gyrus, left and right temporal gyrus, and left and right inferior frontal gryus. The β estimates linearly decreased with increasing effort requirements of the task, $F(1,48)=27.45$, $p< .001$, $\eta_p^2 =.36$.

Positive PC2 (Fig 8A, Table 3) included anterior cingulate cortex, right orbitofrontal cortex, caudal segments of the right parietal cortex, left precuneus, and cuneus. There was a significant effect of effort on Positive PC2 β estimates, $F(3.51,168.37)=4.05$, $p= .005$, $\eta_p^2 =.08$. The β estimates had a quadratic trend across increasing effort requirements of the task, $F(1,48)=19.63$, $p<.001$, $\eta_p^2 =.29$. No cortical or subcortical regions loaded negatively on PC2.

Regions that loaded most positively on PC3 (Fig 8A, Table 3) included the dorsal portions of the occipital lobe. Consistent with its functional definition, there was not a significant effect of effort on Positive PC3 β estimates, $F(3.78,181.22)=2.34$, $p= .06$, $\eta_p^2 =.05$. Despite the visual nature of the task, this component overall showed a task-negative trend across effort levels.

Negative PC3 (Fig 8B, Table 4) included right inferior frontal gyrus, medial frontal cortex, left and right angular gyrus, brainstem and left insula, regions that are mainly adjacent to the regions observed in Positive PC2. There was a significant effect of effort on Negative PC3 β estimates, $F(3.71,178.25)=5.17$, $p= .001$, $\eta_p^2 =.10$. Similar to that of Positive PC2, the β estimates had a quadratic trend across increasing effort requirements of the task, $F(1,48)=29.34$, $p<.001$, $\eta_p^2 =.38$.

Overall, PCA results indicate that a linear function explained more than 50% of the voxel-wise variance in the shape of activation function across effort levels with Positive PC1 showing a positive trend and Negative PC1 showing a negative trend.

3.2. Connectivity change across effort levels

As FPN recruitment linearly increases with increasing effort execution, DMN shows decreased activation. Prior studies have reported that while regions within each of these networks correlated positively with each other, they are negatively correlated between these networks depending on task-requirements and/or individual differences (Gao & Lin, 2012; Elton & Gao, 2014).

In order to test whether connectivity within or between networks differed across effort levels, we calculated the mean connectivity separately for all ROI pairs that paired Between Positive-Negative PC1, within Positive PC1 and within Negative PC1 clusters at each level of effort execution. There was no effect of effort level on within Positive PC1 connectivity, $F(5,245)=1.23$, $p=.29$, $\eta_p^2=.03$, within Negative PC1 connectivity, $F(4.02,197.11)=.26$, $p=.90$, $\eta_p^2=.005$, or between Positive PC1 and Negative PC1 connectivity, $F(5,245)=.47$, $p=.78$, $\eta_p^2=.009$, indicating that connectivity between Positive PC1 and Negative PC1 network did not change across effort levels.

3.3.Brain – Behavior Analysis

Having established the functional form of activity change across effort levels, we tested whether brain regions tracking effort execution predicted effort selections. ROIs were defined from Positive and Negative PC1 to enter the brain-behavior analysis (though results are similar if other definitions of these networks are used, see Supplementary Results 3). Two correlation techniques were adopted to examine the type of brain-behavior relationship: 1) relationship between brain activity and effort selection rate at each effort level, 2) relationship between change in brain activity and change in effort selection rates across effort levels.

Demand avoiders preferred decks that yielded less activity in the Positive PC1 ROI (M_β=-0.20., SD_β=.30) during execution, t(25) = -3.42, p = .002). Conversely, demand seekers preferred decks that yielded greater activity in this ROI (M_β=0.19., SD_β=.44), t(23) = 2.13, p = .04. However, change in Positive PC1 ROI during task execution across effort levels (i.e., due to the task switching manipulation) did not reliably predict change in effort selection rates in either group (demand avoiders: r(26) = -.18, p=.39; demand seekers: r(24) = .36, p=.08). Thus, though we observe a relationship between activation in the fronto-parietal network and avoidance behavior, we did not find evidence that Positive PC1 activation and task selection was related to the parametric effort manipulation (Fig 9A).

The Negative PC1 ROI, by contrast, did show this relationship (Fig 9B). Negative PC1 β estimates did not individually predict task selection behavior independent of our effort manipulation in either group (demand avoiders: t(25) = -0.27, p = .79, demand seekers: t(23) = -2.04, p = .053). However, Pearson correlation between Negative PC1 β estimate slopes and demand avoidance behavior across demand avoider participants was significant (r(26) = .49, *p* =.01). In other words, those individuals who had shallower Negative PC1 slopes (i.e., more positive, less deactivated) across effort levels during task execution showed greater demand avoidance. This relationship was absent for demand seekers (r(24) = .23, *p* =.29). The demand group*Negative PC1 slope interaction was statistically reliable (*b*= -0.11, *t*(46) = -7.74, *p* < .001).

Additionally, we tested for the relationship between connectivity within and between Positive PC1 and Negative PC1 and demand avoidance, in order to see if individual differences in connectivity predict effort selections. Neither connectivity within nor between these networks correlated with behavioral avoidance rates (see Table S2). Thus, together with our analysis on

connectivity change across effort levels, we found no evidence that changes in connectivity were related to either the experience or behavior related to cognitive effort.

Effortful task-execution naturally comes with performance costs such as increased error rates and response times. Thus, the higher likelihood of time on task or errors might lead to demand avoidance rather than engagement of the cognitive control system. In order to differentiate the role of performance measures from that of Positive PC1 and Negative PC1 in demand avoidance, we next regressed out the role of performance from brain-behavior relations. A stepwise multiple regression was conducted to evaluate whether Negative PC1 slope, RT slope and ER slope were necessary to predict selection slope in demand avoiders. At step 1, the Negative PC1 slope entered into the regression equation and was significantly related to selection slope (Negative PC1 slope: $b = -.59$, $t(24) = -2.78$, $p = .01$, $F(1,24) = 7.73$, $p = .01$). The multiple correlation coefficient was .49, indicating approximately 24% of the variance of the selection slope could be accounted for by Negative PC1 slope. At step 2, ER and RT slopes were regressed out of the selection slope, and Negative PC1 slope entered into the regression on the residuals of ER and RT slopes. Negative PC1 slope still reliably predicted the selection slope even after the effects of ER and RT were removed (Negative PC1 slope: $b = -.50$, $t(24) = -2.37$, $p = .03$, $F(1,24) = 5.62$, $p = .03$, $R^2=.16$), indicating that change in Negative PC1 activity during effort execution predicted change in effort selections over and beyond what could be attributed to the change in performance.

As reported above, the change in Positive PC1 activation as a linear function of effort levels did not correlate with avoidance rates, but the level of activation in this set of voxels did correlate generally with avoidance rates. Thus, a similar analysis was conducted to probe whether the overall relationship between activation in the Positive PC1 and individual selection

rates was explained by performance measures. After removing variance due to the behavioral measures, Positive PC1 activation no longer predicted effort selections for either group (demand avoiders: $t(25) = 0.74$, $p = .47$; demand seekers: $t(23) = -0.15$, $p = .88$). However, RT and/or ER continued to predict effort selections following the reverse analysis, when the effect of Positive PC1 was removed from selection rates (demand avoiders: RT: $t(25) = -2.97$, $p = .01$, ER: $t(25) = -2.09$, $p = .05$; demand seekers: RT: $t(23) = 1.40$, $p = .18$, ER: $t(23) = 2.86$, $p = .01$).

Next, we tested for the effects of task-switching probability on the selection rates when the effects of Positive PC1 activation was controlled in order to see if our linear effort manipulation predicted demand avoidance over and beyond Positive PC1 activity. Task-switching probability continued to predict effort selections even when the effect of Positive PC1 was removed (demand avoiders: $t(25) = -4.37$, $p < .001$; demand seekers: $t(23) = 4.37$, $p < .001$), indicating that selection behavior was explained better with our linear task manipulation than Positive PC1 activity.

The preceding analyses suggest that the lack of relationship between change in Positive PC1 and demand avoidance could be due to non-linearities in task performance that drive Positive PC1 activity and which do not scale linearly with our task-manipulation. In order to test for the effects of non-linear task-performance on Positive PC1 activity, we next analyzed the effect of performance (RT and Errors) on Positive PC1 activity when the effect of task-switching probability was controlled. Both RT and ER predicted Positive PC1 activity individually (RT: demand avoiders: $t(25) = 4.98$, $p < .001$; demand seekers: $t(23) = 5.37$, $p < .001$; ER: demand avoiders: $t(25) = 2.59$, $p = .02$; demand seekers: $t(23) = 3.47$, $p = .002$). When the effect of task-switching probability was removed, RT no longer predicted Positive PC1 activity (demand avoiders: $t(25) = 0.96$, $p = .34$; demand seekers: $t(23) = 1.15$, $p = .26$). However, ER continued to

predict Positive PC1 activity for demand avoiders even when the effect of task-switching probability was removed (demand avoiders: $t(25) = -2.24$, $p = .03$; demand seekers: $t(23) = 1.18$, $p = .25$).

Thus, for demand avoiders, ER drove Positive PC1 recruitment over and beyond what could be attributed to the linear effort manipulation. These results indicate that Positive PC1 activity tracks a combination of error and control related signals. Given that the selection rates are explained better by our linear task manipulation than Positive PC1 activity, additional error-related signal in Positive PC1 might explain why individual task activation in Positive PC1 correlated with demand avoidance but not the change across effort levels. Additional analyses that tested the effects of task-switching probability on selection rates showed that when the effect of task performance was removed, task-switching probability, no longer predicted selection rates either (see Supplementary Results 4).

Overall, we have observed that change in Negative PC1 predicts change in effort selection as a function of effort levels in demand avoiders over and beyond what can be explained by performance measures ER and RT. Positive PC1 activation did not predict effort selections over and beyond what is explained by performance measures in either group. While both performance measures can predict Positive PC1 activity as well as effort selections, the effect of performance on both Positive PC1 activity and effort selections is mostly captured by our linear effort manipulation on task-switching probability.

Discussion

The present research aimed to rigorously test the 'cost of control' hypothesis of effort by using a parametric version of DST and observing brain activity in two separate epochs of effort-based-decision-making. The 'Cost of control' hypothesis makes two core proposals: 1) effort registers as disutility by the brain, 2) the cost derives from cognitive control demands. Thus, we predicted that increasing effort would decrease brain activity in the reward network and increase control-related brain activity in FPN. And, further, FPN-activity due to effort execution would predict effort selection. We found only partial support for these predictions, and rather located unexpected evidence that engagement of the DMN (or failure to disengage) due to effort during a task influences effort-based decisions. We consider each of these findings in turn.

Consistent with the idea that effort might register as disutility, we observed that reward network, including VS, VPFC and PCC, linearly reduced activity as a function of increasing effort execution, providing support for the 'cost of effort' hypothesis and consistent with prior reports (Schouppe et al., 2014; Botvinick et al., 2009). We also observed a saturating decrease in VS activation as effort levels increased, consistent with observations by Botvinick et al. (2009). However, VS activity during task execution did not predict effort selections, indicating that effort costs, as tracked by VS during task-execution, do not directly influence demand avoidance behavior in DST paradigms in the absence of monetary reward. Nevertheless, we only scanned during the "Test phase" when the association of effort with the different task conditions had already been established. It is reasonable to hypothesize that the costs computed by VS may be differentially important during the "Learning phase" when effort associations are being acquired.

To test the hypothesis that effort costs arise from engagement of the cognitive control system, we manipulated effort by varying cognitive control demands in a task-switching

paradigm. Consistent with the literature that implicates FPN activity in tasks demanding cognitive control, we observed that most of the variance in our data could be explained by a linearly increasing activation function in this network with increasing task-switching probability.

However, the linear change in FPN activity due to the cognitive control manipulation was not predictive of a change in demand avoidance behavior. Instead, single-level task activity in FPN predicted effort selection rates, independent of a linear effort manipulation. So, for example, a demand avoiding participant selected the 5th effort level the least, if that effort level yielded the highest FPN recruitment. This would be the case even if the 6th effort level required the most task-switching. The reason for this paradoxical relationship might be because the linear functional form of FPN activation is confounded by other factors, such as those related to task performance. We found that performance-related factors like error likelihood predicted FPN activity even after the effects of our experimental cognitive control manipulation were controlled.

Thus, we observe only partial support for the cost of control hypothesis, at least as it might stem from engagement of the FPN network. Performance indices like error-likelihood and RT likely correlate with cognitive control demands in a task, but they could relate to other factors, as well. Thus to avoid circularity in relating cognitive control to FPN activation and demand avoidance, the task switching manipulation provided an independent definition of cognitive control. And using this definition of cognitive control demands, we did not find evidence that changes in activation in FPN attributable to cognitive control were related to demand avoidance.

We note that our observations are not necessarily inconsistent with prior reports. A previous study that also utilized a task-switching paradigm to manipulate cognitive effort

(McGuire and Botvinick, 2010) showed that a change in FPN activity predicted demand avoidance. However, this study used a binary comparison, where a change in brain activity and task selections between low and high effort conditions was not corrected by the change in intermediate task levels. By contrast, our parametric manipulation provided greater sensitivity to separate changes in FPN activation due to the cognitive control manipulation versus other factors related to performance.

A related open question concerns whether demands on cognitive control are themselves effortful, or rather, cognitive control tends to consume more time and it is the time on task that the brain taxes as costly. Kool et al. (2010) showed that participants avoided switching between tasks even when staying on task meant longer task durations. However, we observed that the effect of task-switching probability on effort selections was cancelled when the effects of task-performance were controlled, indicating that our experimental manipulation of cognitive control did not account for demand avoidance in the absence of time-based and error-likelihood differences between conditions. Together with the FPN findings, these results imply that the cost associated with effort cannot be solely explained by recruiting cognitive control during effort execution. Instead error-avoidance, as well as the opportunity-cost of time could constitute the sources of effort costs. Future studies should aim at equating error-rates and time-on-task across effort levels in order to empirically differentiate the effects of task performance from cognitive control recruitment on effort decisions.

Whereas we found partial support for the cost of control hypothesis with regard to engagement of the FPN, a novel discovery in the present study was that the DMN is a robust correlate of effort avoidance. Specifically, demand avoider participants who showed a diminished negative change in DMN activity across cognitive control defined effort levels

showed the highest demand avoidance rate. This relationship between DMN and effort selections persisted even after performance measures such as RT and ER were controlled.

The reason for this relationship between DMN activity and effort avoidance may be an important avenue for future research. While FPN recruitment has been shown to underlie cognitive control, DMN has been associated with a wide range of internal mental functions, including self-referential thoughts, mind-wandering, and episodic future planning (Buckner et al., 2008; Weismann et al., 2006). Inhibition of DMN has been argued to support task-related attention (Spreng, 2012), while the inability to inhibit DMN has been related to dysfunctional cognitive processes such as rumination (Nejad et al., 2013) and pathologies such as depression (Anticevic et al., 2012; Lemogne et al., 2012). Accordingly, in our study, those participants who showed higher effort avoidance rates could have had increased difficulty suppressing DMN activity or relied more on the processes it entails, which in turn might have registered as a cost. Future studies should seek to replicate this discovery and to determine what factor drives this change in DMN activity across effort levels.

It is possible that this DMN activity is related to reward and/or value processing that predicts effort selections. A common observation in the literature is that vmPFC positively tracks subjective value and predicts subjects' preferences between choices (Rushworth et al., 2011). Given this functional association, the 'cost of effort' hypothesis would predict that individuals will avoid tasks that yield the least activity in the reward network. However, in our study, we observed that individuals who showed the least reduction in this network showed the greatest demand avoidance. In addition to linear reductions in VPFC and PCC activity with increasing effort, independent ROI definitions of VS showed that VS reduces its activity across increasing effort, however, DMN but not VS predicted demand avoidance behavior (see Supplementary

Results 3). These results together suggest that while reward network tracks effort in congruence with the 'cost of effort' hypothesis, this cost does not predict effort selections.

We also probed the relationship of functional connectivity to effort avoidance. Stronger functional or structural connectivity within-FPN has been associated with working memory performance and general intelligence (Cole et al., 2012; Gordon et al., 2012; Nagel et al., 2011), executive capacity (Gordon, Lee, et al., 2011), cognitive dysfunctionalities (Rosenberg et al., 2016) and procrastination behavior (Wu et al., 2016). However, we did not find evidence that connectivity within the FPN network or between FPN and DMN networks predicted effort selections.

Finally, individual differences are an important variable in effort avoidance that has not received much attention in the literature. In our fMRI task, we observed high individual variability in effort avoidance, such that roughly half the participants were better characterized as demand seekers than demand avoiders. This rate of variability has not been reported previously in the literature and was not the case in our own behavior-only task (Experiment 1). The conflicting results might be due to three reasons: 1) self-selection bias in fMRI experiments, 2) context effects in fMRI settings, 3) timing effects for subject pool participants. The behavioral participants in our behavior-only study mostly consisted of undergraduate participants who volunteered for course credit. However, our fMRI participants consisted of a more variable sample who volunteered for fMRI scanning at different months of the year. A recent study showed that PET scan volunteers significantly scored higher in sensation-seeking trait compared to behavior-only volunteers (Oswald et al., 2013). Sensation-seeking is also a trait that positively correlated with Need for Cognition (NfC), a self-report inventory that tests effort avoidance and that negatively correlates with Effort Discounting scores (Schuller, 1999). We confirmed in a

follow up study that fMRI volunteers scored significantly higher on NfC compared to subject pool participants (see Supplementary Results 2). Thus, observed differences in effort avoidance behavior might be due to a self-selection bias in fMRI samples. Future studies of cognitive effort should closely attend to this factor.

Regardless of the source of the individual differences, we were able to test how the factors affecting effort-based decisions differed across these groups. Although there were no differences in brain activity and behavioral performance between demand groups during task execution, demand avoiders avoided those tasks that yielded greater FPN recruitment, time-on-task and error-rates, while demand seekers chose them. This discrepancy in the way separate demand groups utilized the same information indicates that FPN, and task performance could influence effort-based value-based decisions generally, even if not necessarily in terms of a cost signal. On the other hand, reduced DMN inhibition across effort levels influenced effort avoidance only in demand avoiding participants, indicating that change in DMN activity across effort levels entered into effort-based decisions exclusively as a cost signal.

In conclusion, we adopted a parametric version of DST in order to test the 'cost of control' hypothesis and explore the neural mechanisms that underlie effort execution. We have observed that the reward network reduces activity in response to executing more effortful tasks, in congruence with the 'cost of effort' hypothesis. FPN and DMN both predict effort avoidance behavior, although the effect of FPN is mediated by performance measures such as time-on-task and error-likelihood, indicating that behavioral task performance predicts effort-based decisions better than FPN activity. As behavioral task performance can represent the required cognitive control to perform well at a task, increased time-on-task and error-likelihood could constitute costs associated with the opportunity-cost of time and error-avoidance. Additionally, high

individual variability in effort avoidance behavior in our task showed that the direction greater FPN activity and behavioral task performance influenced effort-based decisions depended on the demand group, indicating that such factors do not exclusively influence effort-based decisions as a cost. On the other hand, it was shown that reduced DMN inhibition was an exclusive cost signal that differentiated demand avoiders from demand seekers, promising to be a new avenue for future effort-based research.

Acknowledgements: We are grateful to members of the Badre lab, the Frank lab, and the Shenhav lab for their helpful interactions over this work. In particular, we thank Amitai Shehav and Michael Frank for their insightful feedback. We also thank Britanny Ciullo and Nada Hamzah for their assistance with data collection. The present work was supported by R01 awards from the National Institute of Neurological Disease and Stroke (NS065046) and the National Institute of Mental Health (MH099078), a MURI award from the Office of Naval Research (N00014-16-1-2832), and a fellowship from the James S. McDonnell Foundation. CS was also supported through a T32 institutional training grant from the NIH (3T32NS062443-08S1).

References

Anticevic, A., Cole, M. W., Murray, J. D., Corlett, P. R., Wang, X. J., & Krystal, J. H. (2012). The role of default network deactivation in cognition and disease. Trends in cognitive sciences, 16(12), 584-592.

Arrington, C. M., & Logan, G. D. (2004). The cost of a voluntary task switch. Psychological science, 15(9), 610-615.

Badre, D., & D'Esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical?. *Nature reviews. Neuroscience*, *10*(9), 659.

Badre, D., Lebrecht, S., Pagliaccio, D., Long, N. M., & Scimeca, J. M. (2014). Ventral striatum and the evaluation of memory retrieval strategies. Journal of cognitive neuroscience, 26(9), 1928-1948.

Behzadi, Y., Restom, K., Liau, J., & Liu, T. T. (2007). A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. Neuroimage, 37(1), 90-101.

Botvinick, M. M. (2007). Conflict monitoring and decision making: reconciling two perspectives on anterior cingulate function. Cognitive, Affective, & Behavioral Neuroscience, 7(4), 356-366.

Botvinick, M. M., Huffstetler, S., & McGuire, J. T. (2009). Effort discounting in human nucleus accumbens. Cognitive, Affective, & Behavioral Neuroscience,9(1), 16-27. PubMed PMID: 19246324.

Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network. Annals of the New York Academy of Sciences, 1124(1), 1-38.

Cole, M. W., & Schneider, W. (2007). The cognitive control network: integrated cortical regions with dissociable functions. *Neuroimage*, *37*(1), 343-360.

Cole, M. W., Yarkoni, T., Repovš, G., Anticevic, A., & Braver, T. S. (2012). Global connectivity of prefrontal cortex predicts cognitive control and intelligence. Journal of Neuroscience, 32(26), 8988-8999.

Decot, H. K., Namboodiri, V. M., Gao, W., McHenry, J. A., Jennings, J. H., Lee, S. H., ... & Deisseroth, K. (2017). Coordination of brain-wide activity dynamics by dopaminergic neurons. *Neuropsychopharmacology*, *42*(3), 615-627.

Dosenbach, N. U., Fair, D. A., Miezin, F. M., Cohen, A. L., Wenger, K. K., Dosenbach, R. A., ... & Schlaggar, B. L. (2007). Distinct brain networks for adaptive and stable task control in humans. *Proceedings of the National Academy of Sciences*, *104*(26), 11073-11078.

Eklund, A., Nichols, T. E., & Knutsson, H. (2016). Cluster failure: why fMRI inferences for spatial extent have inflated false-positive rates. Proceedings of the National Academy of Sciences, 201602413.

Elton, A., & Gao, W. (2014). Divergent task-dependent functional connectivity of executive control and salience networks. Cortex, 51, 56-66.

Flandin, G., & Friston, K. J. (2016). Analysis of family-wise error rates in statistical parametric mapping using random field theory. arXiv preprint arXiv:1606.08199.

Fedorenko, E., Duncan, J., & Kanwisher, N. (2013). Broad domain generality in focal regions of frontal and parietal cortex. *Proceedings of the National Academy of Sciences*, *110*(41), 16616-16621.

Gao, W., & Lin, W. (2012). Frontal parietal control network regulates the anti-correlated default and dorsal attention networks. Human brain mapping, 33(1), 192-202.

Gordon, E. M., Lee, P. S., Maisog, J. M., Foss-Feig, J., Billington, M. E., VanMeter, J., & Vaidya, C. J. (2011). Strength of default mode resting-state connectivity relates to white matter integrity in children. Developmental science, 14(4), 738-751.

Gordon, E. M., Stollstorff, M., Devaney, J. M., Bean, S., & Vaidya, C. J. (2012). Effect of dopamine transporter genotype on intrinsic functional connectivity depends on cognitive state. Cerebral Cortex, 22(9), 2182-2196.

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision making and the avoidance of cognitive demand. Journal of Experimental Psychology: General, 139(4), 665. PubMed PMID: 20853993.

Lemogne, C., Delaveau, P., Freton, M., Guionnet, S., & Fossati, P. (2012). Medial prefrontal cortex and the self in major depression. Journal of affective disorders, 136(1), e1-e11.

McGuire, J. T., & Botvinick, M. M. (2010). Prefrontal cortex, cognitive control, and the registration of decision costs. Proceedings of the National Academy of Sciences, 107(17), 7922-7926. PubMed PMID: 20385798.

Monsell, S. (2003). Task switching. Trends in cognitive sciences, 7(3), 134-140.

Nagel, I. E., Preuschhof, C., Li, S. C., Nyberg, L., Bäckman, L., Lindenberger, U., & Heekeren, H. R. (2011). Load modulation of BOLD response and connectivity predicts working memory performance in younger and older adults. Journal of Cognitive Neuroscience, 23(8), 2030-2045.

Nejad, A. B., Fossati, P., & Lemogne, C. (2013). Self-referential processing, rumination, and cortical midline structures in major depression. Frontiers in human neuroscience, 7, 666.

Niendam, T. A., Laird, A. R., Ray, K. L., Dean, Y. M., Glahn, D. C., & Carter, C. S. (2012). Meta-analytic evidence for a superordinate cognitive control network subserving diverse executive functions. *Cognitive, Affective, & Behavioral Neuroscience*, *12*(2), 241-268.

Oswald, L. M., Wand, G. S., Zhu, S., & Selby, V. (2013). Volunteerism and self-selection bias in human positron emission tomography neuroimaging research. Brain imaging and behavior, 7(2), 163-176.

Ray, K. L., Lesh, T. A., Howell, A. M., Salo, T. P., Ragland, J. D., MacDonald, A. W., ... & Carter, C. S. (2017). Functional network changes and cognitive control in schizophrenia. *NeuroImage: Clinical*, *15*, 161-170.

Rushworth, M. F., Noonan, M. P., Boorman, E. D., Walton, M. E., & Behrens, T. E. (2011). Frontal cortex and reward-guided learning and decision-making. *Neuron*, *70*(6), 1054-1069.

Rosenberg, M. D., Finn, E. S., Scheinost, D., Papademetris, X., Shen, X., Constable, R. T., & Chun, M. M. (2016). A neuromarker of sustained attention from whole-brain functional connectivity. Nature neuroscience, 19(1), 165-171.

Schouppe, N., Demanet, J., Boehler, C. N., Ridderinkhof, K. R., & Notebaert, W. (2014). The role of the striatum in effort-based decision-making in the absence of reward. The Journal of Neuroscience, 34(6), 2148-2154. PubMed PMID: 24501355

Schuller, I. S. (1999). Procrastination, need for cognition and sensation seeking. Studia psychologica, 41(1), 73-85.

Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: an integrative theory of anterior cingulate cortex function. Neuron,79(2), 217-240. PubMed PMID: 23889930.

Spreng, R. N. (2012). The fallacy of a “task-negative” network. Frontiers in psychology, 3, 145.

Westbrook, A., Kester, D., & Braver, T. S. (2013). What is the subjective cost of cognitive effort? Load, trait, and aging effects revealed by economic preference. PLoS One, 8(7), e68210.

Whitfield-Gabrieli, S., & Nieto-Castanon, A. (2012). Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks. Brain connectivity, 2(3), 125-141.

Wu, Y., Li, L., Yuan, B., & Tian, X. (2016). Individual differences in resting-state functional connectivity predict procrastination. Personality and Individual Differences, 95, 62-67.

Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., ... & Fischl, B. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. Journal of neurophysiology, 106(3), 1125-1165.

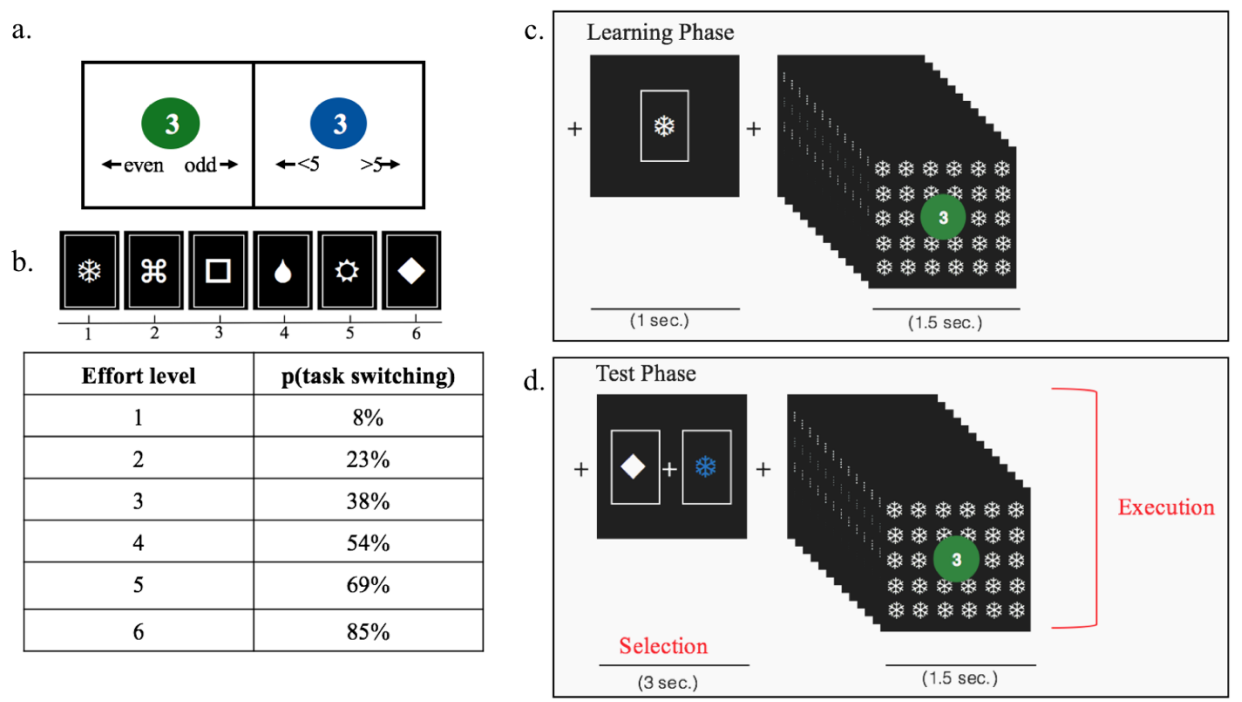

Effort level	p(task switching)
1	8%
2	23%
3	38%
4	54%
5	69%
6	85%

Figure 1 Schematic of the parametric demand selection task. (a) Example trials of the categorization task with a digit printed on a colored circle. The color cued which task to perform. The response mappings shown below the digits are for the reader's reference and did not actually appear on the screen. (b) An example of six deck cues with their associated effort level pairing, based on the probability of task switching in 13 trials from that deck. (c) Schematic of block events during the Learning Phase. First, a symbol icon is presented for 1 sec, after which the participant completes 13 categorization trials with the associated cue tiled on the background. (d) Schematic of block events during the Test Phase. In the "Selection" epoch, the participant chooses between two symbol icons that are associated with two different effort levels, with a deadline of 3 sec. In the "Execution" epoch, the participant executes the effort level associated with the selected option, while the selected cue is tiled on the background.

	Execution>Baseline	**Baseline>Execution**	**Parametric+**	**Parametric-**
Demand Avoiders	191	51	56	220
Demand Seekers	26	41	101	181

Table 1 Cluster extent (k) thresholds as computed by SPM for each contrast and demand groups using a cluster forming threshold, $\alpha = .001$).

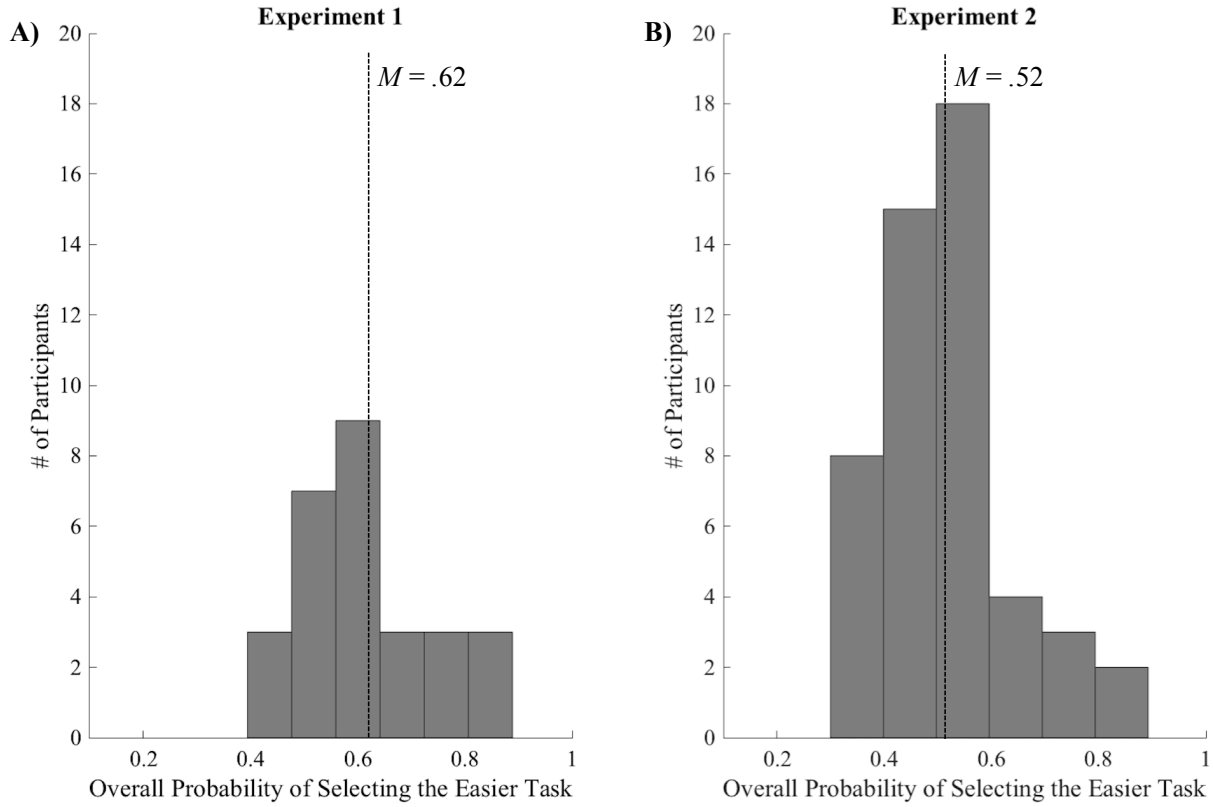

Figure 2 Distribution of participants for the overall probability of selecting the easier task across all decision trials in A) Experiment 1 and B) Experiment 2. In Experiment 1, 25 of the 28 participants (89%) selected the easier task more than 50% of the time overall. In Experiment 2, nearly half of participants (N=24) selected the easier task more than 50% of the time. The mean overall probability of selecting the easier task for each experiment is indicated with a vertical dashed line.

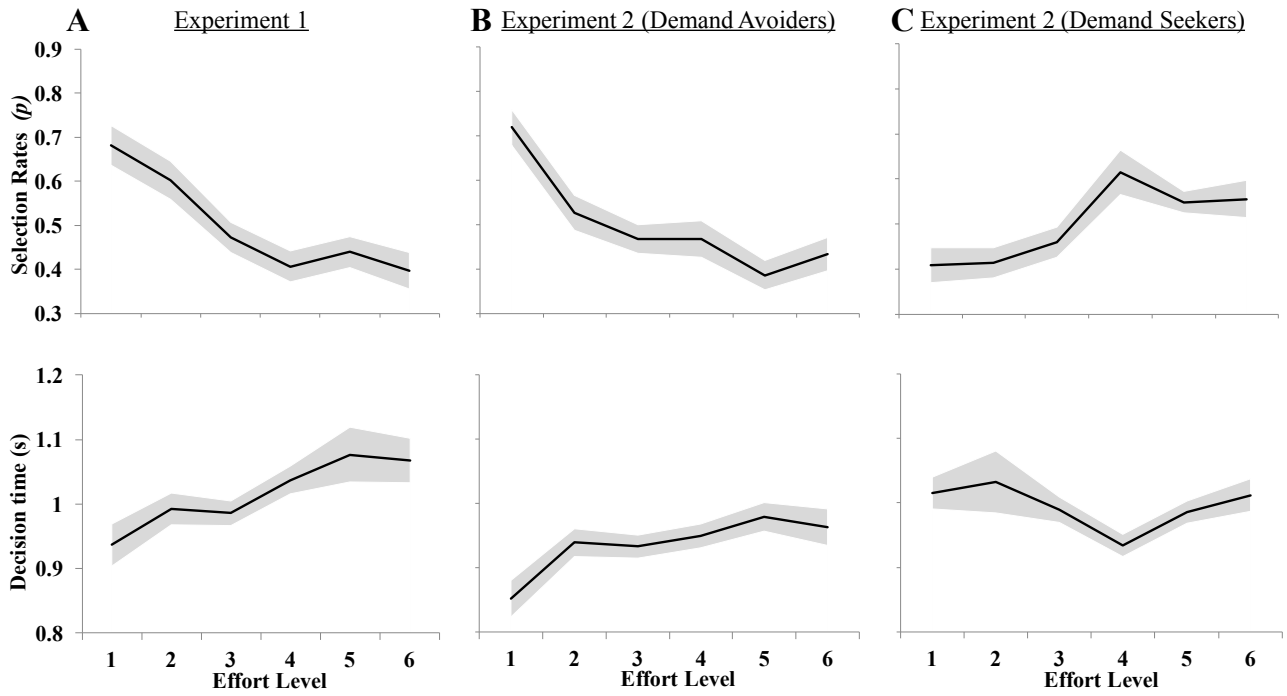

Figure 3 Choice behavior for the entire Experiment 1 sample (A), demand avoiders in Experiment 2 (B), and demand seekers in Experiment 2 (C). The top panels plot the selection rates across effort levels in terms of the probability of choosing that task when it was given as an option. In Experiment 1, the entire sample showed a decreasing tendency to choose higher effort levels. In Experiment 2, demand avoiders showed a decreasing tendency to choose a task as more cognitive effort is required, and demand seekers exhibited the opposite pattern. The bottom panels plot the decision times across effort levels. In Experiment 1, decision time was mostly unaffected by effort level with only a marginal omnibus effect of effort level on decision time. In Experiment 2, the decision time to select an effort level linearly increased with increasing effort levels for demand avoiders and the decision time to choose an effort level was similar across effort levels for demand seekers. All positive and negative trends were significant linear effects at $p < .05$. The shaded area plots standard error of the mean.

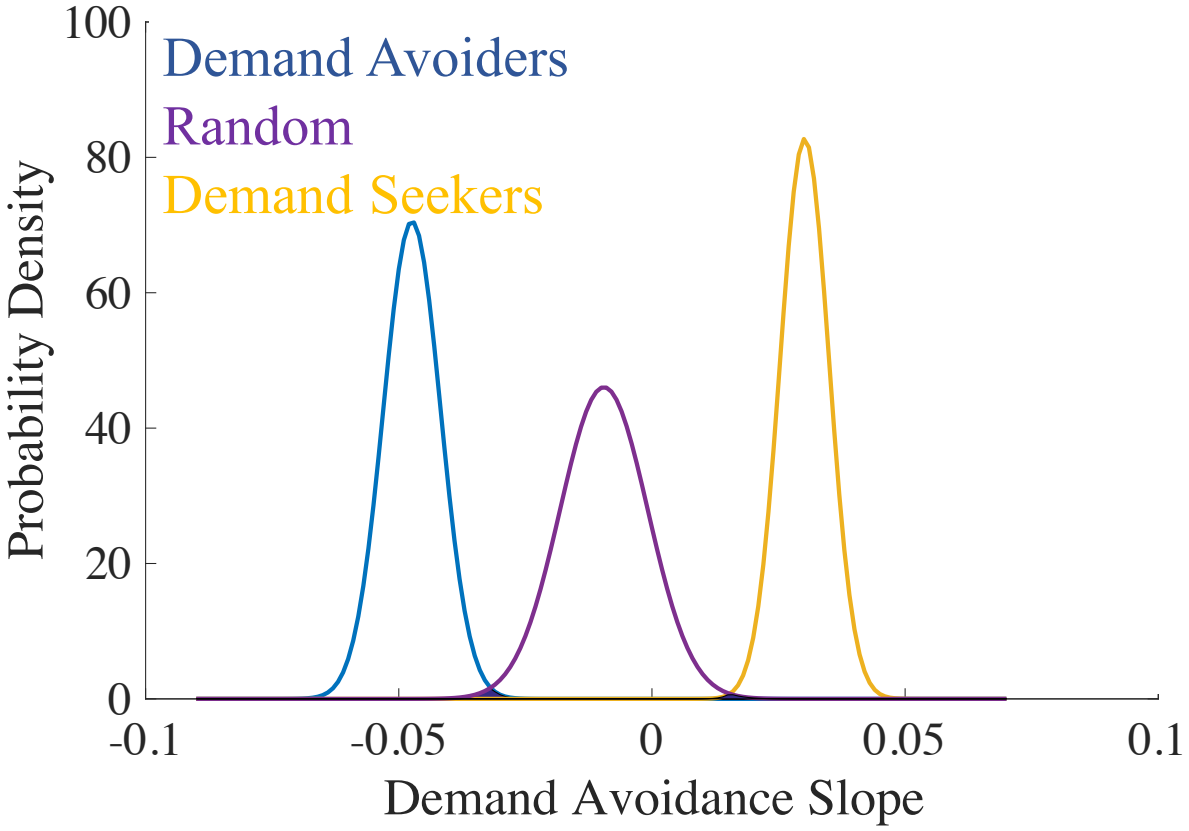

Figure 4 The overlap of the likelihood distribution for the Demand Avoider (blue) and Demand Seeker (yellow) groups with the permuted null distribution (purple). The mean demand avoidance slope for demand avoiders is negative, and the demand avoidance slope for demand seekers is positive, consistent with a respective tendency to avoid versus select a task as function of its effort. These distributions overlapped minimally with the permuted null distribution (Demand avoider: p = .008, Demand seeker: p = .003).

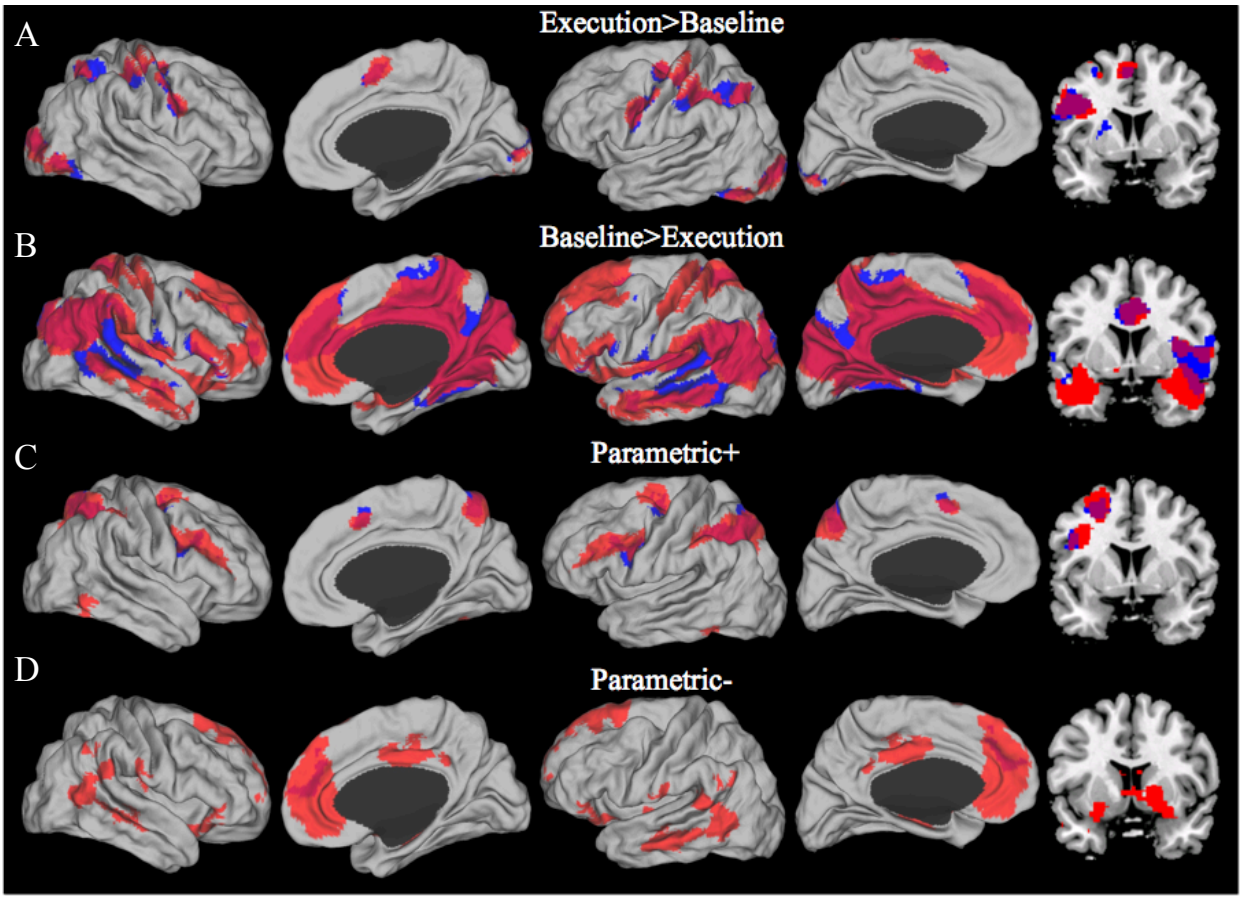

Figure 5 Four whole brain contrasts are plotted on a canonical surface with a coronal slice (Y=8) on the far right to show activation in the basal ganglia. Red: demand avoiders; blue: demand seekers. (A) The contrast of all blocks (across effort levels) from the execution phase versus baseline primarily shows the FPN and other typical task active areas. (B) The reverse contrast of baseline > execution phase blocks shows typical task negative activations including the DMN. (C) The positive parametric effort level effect shows regions that demonstrated a linear increase with effort level, including regions of the FPN. (D) The negative parametric effort level effect shows regions that were negatively correlated with effort level including regions of the DMN and ventral striatum. Images are thresholded at $p < .05$ FWE cluster corrected.

	Brain Region	Cluster size	MNI coordinates X	Y	Z	Voxel t-value
Execution > Fix						
Demand Avoidants						
	right occipital cortex	2320	26	-94	2	15.78
	left occipital cortex	1502	-26	-94	2	11.74
	left supramarginal gyrus	3830	-44	-36	42	7.19
	right superior parietal lobule	292	30	-58	44	5.33
Demand Seekers						
	left occipital cortex	2998	-16	-96	0	13.21
	right occipital cortex	3502	-28	-52	46	10.57
	right superior parietal lobule	406	34	-56	54	8.6
	left putamen	373	-22	4	10	5.6
	supplementary motor cortex	232	-8	10	52	5.14
Fix > Execution*						
Demand Avoidants						
	cuneus	6380	16	-66	22	14.82
	middle cingulate gyrus	2827	0	-18	42	13.67
	anterior cingulate gyrus	1856	-10	40	14	11.97
	left angular gyrus	472	-44	-74	34	11.4
	right middle occipital gyrus	150	50	-70	20	10.76
	left superior frontal gyrus	316	-20	48	34	10.09
Demand Seekers						
	posterior cingulate cortex	9135	-2	-32	46	17.83
	anterior cingulate gyrus	474	2	44	16	12.83
	right posterior insula	731	44	-8	-4	12.47
	left supramarginal gyrus	142	-60	-50	26	12.41
	right middle temporal gyrus	179	66	-20	-14	12.03
	right angular gyrus	381	50	-68	28	11.47
	left middle temporal gyrus	145	-58	-52	-2	11.33
	left middle occipital gyrus	148	-42	-78	6	10.42
Parametric+						
Demand Avoidants						
	left superior parietal lobule	3343	-24	-64	46	8.96
	left Inferior frontal gyrus	2146	-46	32	18	7.73
	left supplementary motor cortex	253	-8	16	44	5.91
	right superior parietal lobule	697	20	-62	54	5.62
	left inferior temporal gyrus	288	-48	-56	-14	5.3
Demand Seekers	left supramarginal gyrus	2215	-40	-42	44	7.24
	left supplementary motor cortex	550	-6	14	46	5.42
	left middle Frontal Gyrus	483	-42	26	26	5.1
Parametric-						
Demand Avoidants						
	right middle temporal gyrus	2776	52	-60	6	7.04
	medial superior frontal gyrus	10580	10	48	16	7
	right hippocampus	246	32	-28	-10	6.51
	middle cingulate gyrus	1952	4	-12	38	6.16
	left middle temporal gyrus	905	-64	-20	-14	5.8
	left precentral gyrus	308	-46	-10	54	5.24
	left posterior insula	404	-38	-16	-2	4.86
Demand Seekers						
	left superior frontal gyrus	456	-12	42	26	5.13

Table 2 Whole brain analysis for both demand groups. (* indicates that higher clustering threshold was used in order to break down big clusters for reporting purposes)

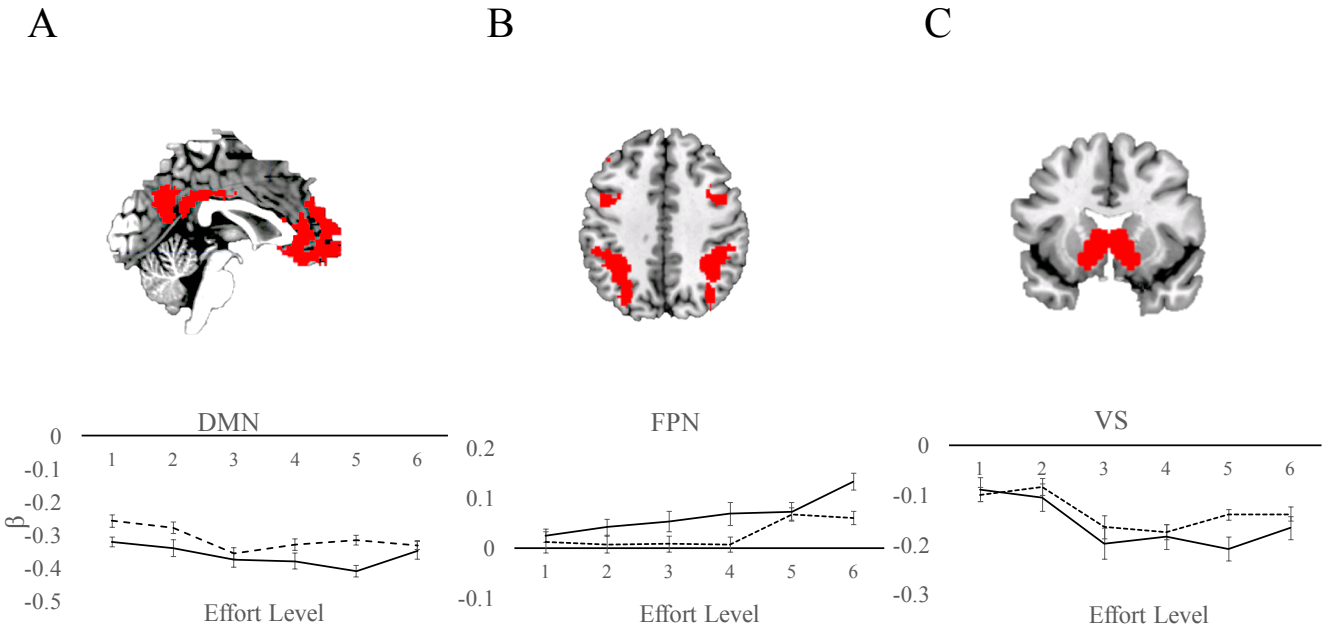

Figure 6 Estimated beta weights at each effort level for both demand groups from independent ROIs defined in (A) the default mode network (DMN), (B) the frontoparietal network (FPN), and (C) the ventral striatum (VS). Demand avoiders are indicated in solid lines. Demand seekers are indicated in dashed lines. There was no effect of avoidance group on β estimates for any of the ROIs. In all cases a linear function was the curve of best fit relative to alternatives (summarized in Supplementary Table S1), with the exception of VS. All linear trends are significant at $p < .05$ for both groups.

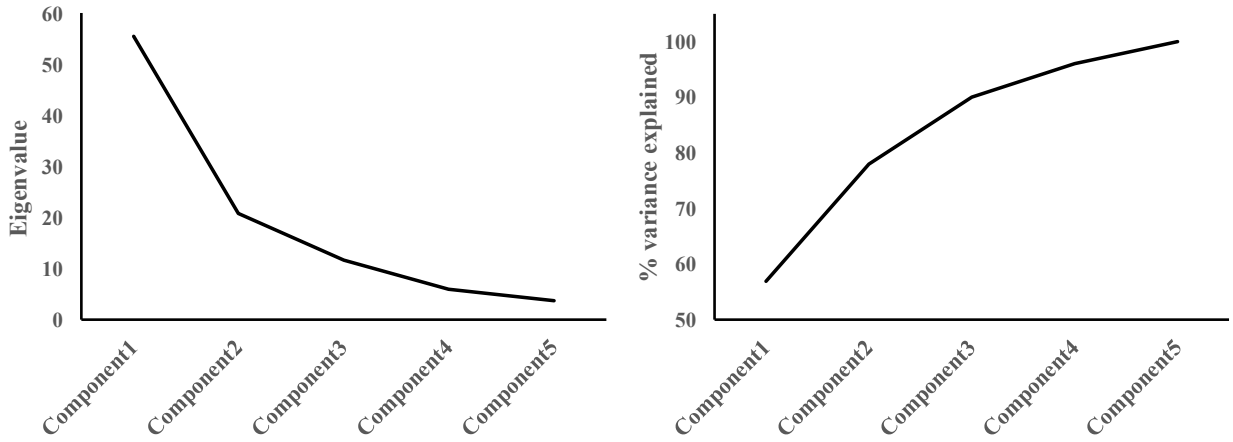

Figure 7 Summary of the top components identified by the PCA of the activation pattern over effort levels across all voxels in the brain. Left panel: Eigenvalue by each PC component. Right Panel: Variance explained by each PC component. Further analysis focused on the top 3 components as these explained 90% of the variance.

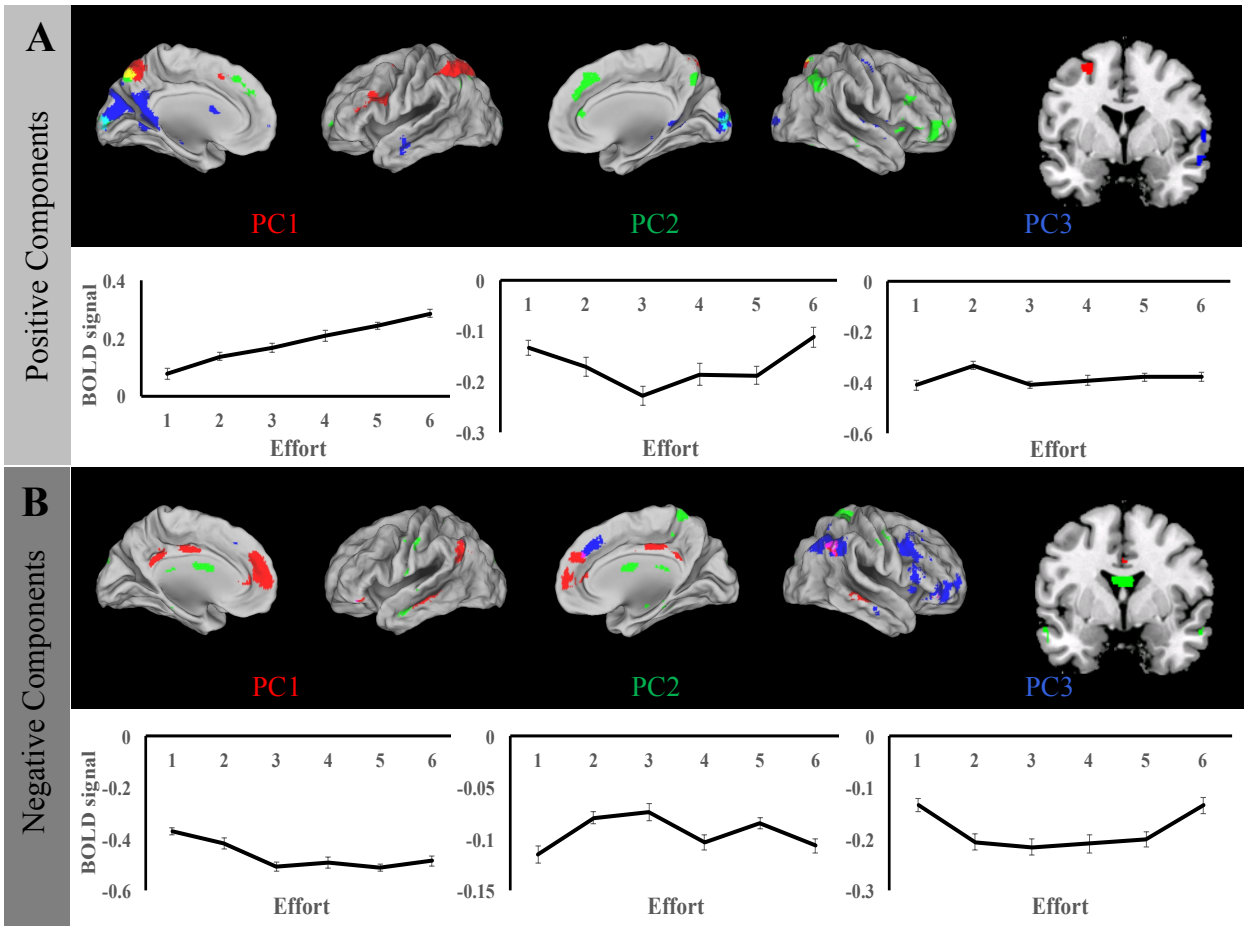

Figure 8 Visualization of PCA components and their neural functions across effort levels. A) Results from an analysis of the first 3000 voxels that load positively on components 1-3. Component1 showed a significant linear trend; Component2 showed a significant quadratic trend; Component3 beta estimates were flat across effort levels. B) Results from an analysis of the first 3000 voxels that load negatively on components 1-3. Component1 showed a significant negative linear trend; Component2 was not fit by any function. Component3 beta estimates showed a quadratic trend across effort levels. On the brain: Red: Component1; Green: Component2; Blue: Component3.

	Brain Region	Cluster size	MNI coordinates X	Y	Z
Positive Components					
Component1					
	left SMA	117	-1	17	47
	left superior parietal lobule	2382	-23	-58	50
	left middle frontal gyrus	75	-27	2	57
	left inferior frontal gyrus	426	-44	17	26
Component2					
	left superior frontal gyrus	16	-24	58	1
	left angular gyrus	158	-40	-61	52
	left orbital gyrus	10	-41	44	-9
	left temporal pole	20	-46	16	30
	left precuneus	288	1	-68	46
	left cuneus	297	2	-93	3
	right anterior cingulate gyrus	1044	2	38	28
	right middle frontal gyrus	411	36	52	-2
	right angular gyrus	391	42	-58	47
	right occipital fusiform gyrus	16	44	-63	-18
	right frontal operculum	198	46	21	-3
	right middle frontal gyrus	112	46	24	26
	right middle temporal gyrus	28	63	-23	-14
Component3					
	left caudate	22	-14	13	19
	left calcarine cortex	2515	-1	-70	13
	left hippocampus	8	-26	-22	-16
	left thalamus	8	-2	-21	11
	left precuneus	10	-4	-70	55
	left thalamus	33	-4	9	14
	left middle frontal gyrus	166	-57	-10	-17
	right SMA	9	-6	61	0
	right hippocampus	23	25	-26	-13
	right occipital fusiform gyrus	19	26	-70	-15
	right precentral gyrus	35	42	-14	57
	right superior temporal gyrus	33	59	0	-14
	right central operculum	36	61	0	3
	right superior temporal gyrus	11	66	-16	7
	right cerebellum	46	8	-65	-7

Table 3 Cluster coordinates for the Positive PCs.

	Brain Region	Cluster size	MNI coordinates X	Y	Z
Negative Components					
Component1					
	medial prefrontal cortex	1409	0	50	15
	left inferior frontal gyrus	19	-48	29	-9
	left angular gyrus	307	-51	-60	31
	left middle temporal gyrus	179	-63	-28	-7
	left posterior cingulate gyrus	725	1	-29	38
	right inferior frontal gyrus	11	47	41	-7
	right angular gyrus	99	55	-51	35
	right middle temporal gyrus	244	61	-26	-9
Component2					
	Component not reported due to sensitivity to artifacts				
Component3					
	left insula	19	-42	15	-3
	left angular gyrus	131	-43	-54	53
	brainstem	99	1	-20	-5
	right middle frontal gyrus	24	27	28	48
	right superior frontal gyrus	447	2	33	36
	right inferior frontal gyrus	1483	42	31	14
	right angular gyrus	759	45	-56	42
	right middle frontal gyrus	27	61	-14	-21

Table 4 Cluster coordinates for the Negative PCs.

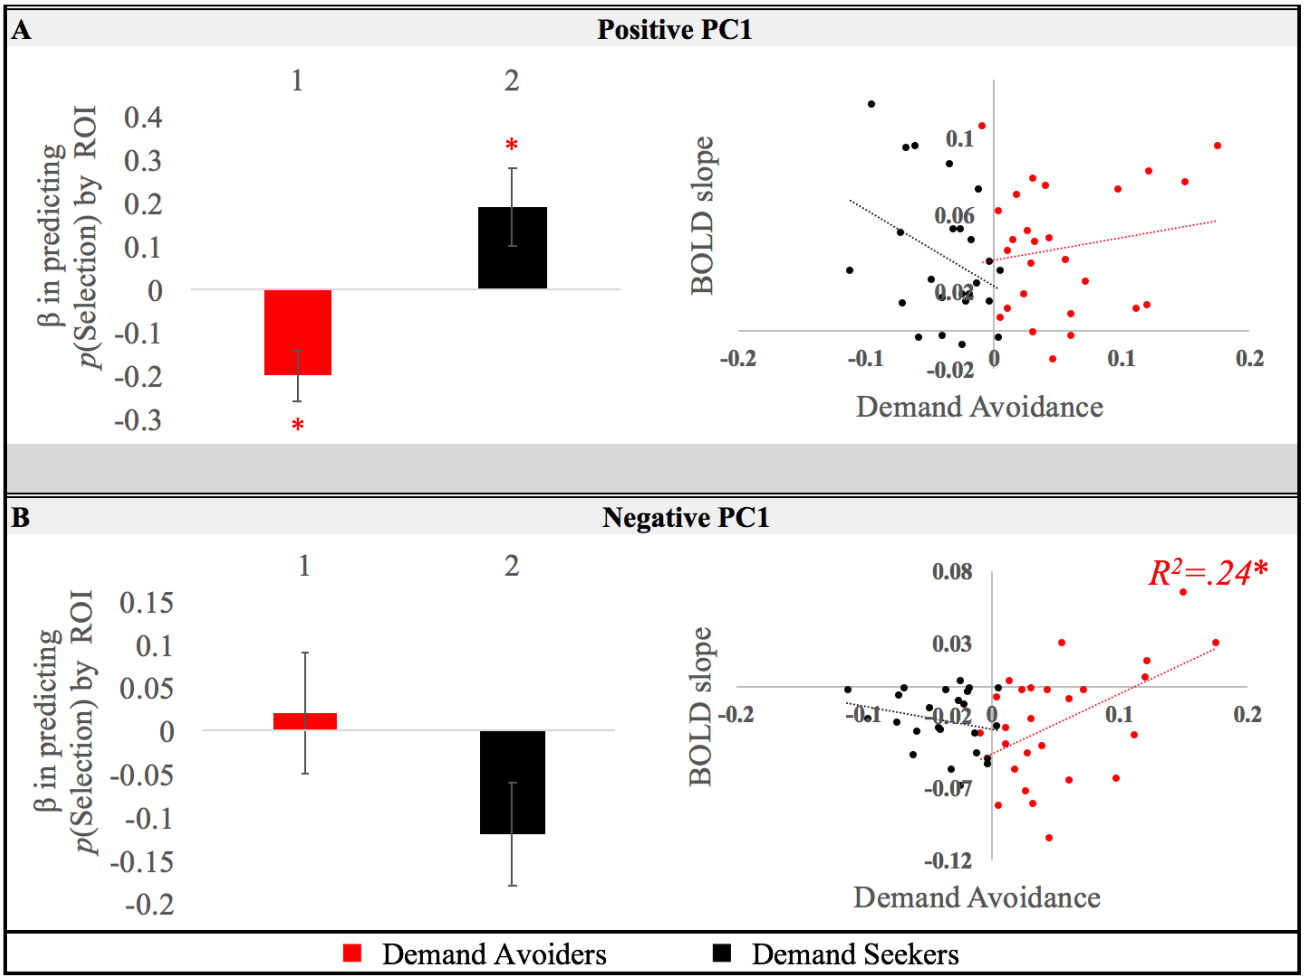

Figure 9 Principal Component cluster correlations with demand avoidance behavior are plotted for (A) Positive and (B) Negative PC1. Results for demand avoiders are plotted in red. Results for demand seekers are plotted in black. Left panel: Individual effort selection rate was correlated to the brain activity during task execution at that effort level unranked with respect to task switching frequency for each demand group separately. Overall β in predicting *p*(Selection) is plotted and was calculated by averaging correlation coefficients for each PCA ROI across all individuals in a demand group. Activity in the Positive PC1 network was correlated with effort selections, and in opposite directions depending on the demand group. Right panel: The 'Effort level' analysis tested whether the change in effort selection rates was correlated to the change of

brain activity in each PCA ROI during task execution across effort levels. Negative PC1 showed a positive relationship between the slope of change in this network across effort levels and effort selections. This was only the case in the demand avoidance group. Error bars plot standard error of the mean. * $p < .05$.