

Washington University in St. Louis

Washington University Open Scholarship

Arts & Sciences Electronic Theses and
Dissertations

Arts & Sciences

Winter 12-15-2016

Neural Dynamics Tracking Subjective Cognitive Effort

John Andrew Westbrook
Washington University in St. Louis

Follow this and additional works at: https://openscholarship.wustl.edu/art_sci_etds

Recommended Citation

Westbrook, John Andrew, "Neural Dynamics Tracking Subjective Cognitive Effort" (2016). *Arts & Sciences Electronic Theses and Dissertations*. 1010.
https://openscholarship.wustl.edu/art_sci_etds/1010

This Dissertation is brought to you for free and open access by the Arts & Sciences at Washington University Open Scholarship. It has been accepted for inclusion in Arts & Sciences Electronic Theses and Dissertations by an authorized administrator of Washington University Open Scholarship. For more information, please contact digital@wumail.wustl.edu.

WASHINGTON UNIVERSITY IN ST. LOUIS
Division of Psychological & Brain Sciences

Dissertation Examination Committee:

Todd Braver, Chair

Cynthia Cryder

Ian Dobbins

Joel Myerson

Jeff Zacks

Neural Dynamics Tracking Subjective Cognitive Effort

by

Andrew Westbrook

A dissertation presented to
The Graduate School
of Washington University in
partial fulfillment of the
requirements for the degree
of Doctor of Philosophy

December 2016
St. Louis, Missouri

© 2016, John Andrew Westbrook

Table of Contents

List of Figures	iv
List of Tables	vi
Acknowledgments	vii
Abstract	viii
Chapter 1: Introduction.....	1
1.1 Brain Regions Tracking Effort.....	2
1.1.1 Limitations of Cognitive Effort Literature.....	8
1.2 Cognitive Effort Discounting (COGED) Paradigm	10
1.3 Overview of Study & Predictions	13
1.3.1 Predictions about regions tracking objective load	14
1.3.2 Predictions about networks tracking subjective effort.....	15
1.3.2 Predictions about subjective effort tracking in the dACC, dlPFC, and VS	18
Chapter 2: Methods.....	20
2.1 Procedure Overview	20
2.2 General Task Descriptions	21
2.2.1 N-back.....	21
2.1.2 COGED.....	22
2.2 Participants	23
2.3 Imaging Procedure	25
2.4 Scanning Parameters	26
2.5 Image Processing.....	27
Chapter 3: Behavioral Results	29
3.1 N-back Performance.....	29
3.2 Decision-Making Behavior	34
3.4 Summary	38
Chapter 4: Brain Regions Tracking Cognitive Effort.....	40
4.1 Brain Regions Varying by Load.....	41
4.1.1 Load Modulation by Network.....	43

4.2 Do Brain Networks Vary by Load and Discounting?	55
4.2.1 Specific Networks Predict Discounting	56
4.2.2 Specific Nodes Within Networks Predicting Discounting.....	62
4.2.3 Limitations	72
4.3 Summary	73
Chapter 5: General Discussion	76
5.1 Behavioral Results.....	76
5.2 Regions of the Brain Tracking Cognitive Effort	78
5.3 Future Directions	85
References.....	89
Appendix A.....	100
Appendix B	176
Appendix C	179
Appendix D.....	192

List of Figures

Figure 1.1 Subjective value bar plots for two hypothetical participants.....	13
Figure 1.2 Hypothetical effects of effort costs and N-back load on BOLD signal in two task-positive regions	16
Figure 1.3 Combination of discounting and BOLD signal in task-positive networks for specific inferences about the relationship between brain dynamics and subjective motivation for effort, or effort costliness.....	17
Figure 2.1 Schematic of the three phases of the COGED paradigm	20
Figure 2.2 Three session procedure overview schematic	21
Figure 3.1 Group performance d' by load and by AUC group for both Session 2.....	30
Figure 3.2 Lure rates across all loads and both groups for Sessions 2	31
Figure 3.3 μ RT across loads and groups for Session 2.....	32
Figure 3.4 Subjective values across N-back task levels and base offer amounts for COGED decision-making in Session 2.....	35
Figure 3.5 Pairwise COGED AUC plots for Sessions 1 and 2 and Sessions 2 and 3.....	36
Figure 3.6 Averaged 3-session AUC histogram	38
Figure 4.1 GLM β weights for each N-back level averaged across nodes in the DMN...	45
Figure 4.2 Mean GLM β weights for each level of the N-back task averaged across all nodes in the A) FP, B) DorAtt, and C) Sal Networks and a map of nodes included in the respective analyses.....	46
Figure 4.3 Mean GLM β weights for each level of the N-back task averaged across all nodes in the CO network, and a map of nodes included in the analysis.....	47

Figure 4.4 Mean GLM β weights for each level of the N-back in the dACC and the location of the 12 mm spherical node on the medial wall of the left hemisphere	50
Figure 4.5 Mean GLM β weights for each level of the N-back in the dlPFC and the location of the 12 mm node on the left lateral hemisphere surface	51
Figure 4.6 Mean GLM β weights for each level of the N-back in the precuneus.....	52
Figure 4.7 Mean GLM β weights for each level of the N-back in the medial pericingulate	52
Figure 4.8 Load profiles for high (above median 3-back d' scores) and low performers in FP and DorAtt networks	54
Figure 4.9 AUC_{3S} as a function of network-level individual differences in activity in the DorAtt network and activity as a function of load and discounting	58
Figure 4.10 Individual differences in activity in an a priori dACC node as a function of AUC_{3S} and load.....	64
Figure 4.11 Individual differences in average activity in a priori VS nodes predicts AUC_{3S}	65
Figure 4.12 Nodes in which load-independent N-back activity predicts subjective effort costs.....	67
Figure 4.13 Individual differences in activity in a left IPL node as a function of AUC_{3S} and load.....	68
Figure 4.14 Nodes in which load-independent N-back activity predicts subjective effort costs, controlling for performance	71

List of Tables

Table 3.1 Average and standard deviation of median response times by AUC group and by N-back load.....	33
Table 4.1 Linear and quadratic fixed effects of load in networks of interest for N = 1—5 and N = 1—6.....	49
Table 4.2 Repeated Measures ANOVA in N-back regression weights in the DMN, FP, CO, Sal, and DorAtt networks	56
Table 4.3 Effects of load-dependent and load-independent activation on discounting in selected networks	57
Table 4.4 Effects of load-dependent and load-independent activation on discounting in a priori nodes of interest taken from McGuire et al., 2010 and Botvinick et al., 2009	63
Table 4.5 Performance-independent relationship between load-dependent and load-independent activity and discounting in select nodes	69

Acknowledgments

I am grateful to so many students, faculty, and staff in the Department of Psychological and Brain Sciences whom I've come to cherish as colleagues, mentors, and friends. This is certainly true of each of my Committee members, whom I deeply respect and for whom I feel a warm and abiding collegiality.

I am especially grateful to the good people in the Cognitive Control and Psychopathology Lab. To Carol for her thoughtful guidance and ready assistance. To my dear lab mates and friends, who showed me the nuts and bolts and also the big picture. To my advisor-parents Todd and Deanna who not only provided me exceptionally good role-models as scientists, but also as co-workers, educators, mentors, and members of society. And again, to Todd. From our first conversation, I knew I had lucked into finding a mentor with whom I would share great intellectual and personal chemistry, who would provide rich opportunities for learning and exploration, who would show me how to probe the most interesting questions in numerous ways, who would meld tireless curiosity and rigor with playful humanity. I am profoundly grateful for his mentorship and for providing me with a model of the kind of scientist I hope to become.

Above all, I am eternally grateful to my family for their unending love and support. I cannot begin to do justice to how much it means to me.

Finally, I am very grateful to the National Institutes of Health for funding my research and training.

Andrew Westbrook

Washington University in St. Louis

December 2016

ABSTRACT OF THE DISSERTATION
Neural Dynamics of Monitoring and Deciding About Cognitive Effort

by

John Andrew Westbrook

Doctor of Philosophy in Psychological & Brain Sciences

Brain, Behavior, & Cognition

Washington University in St. Louis, 2016

Professor Todd Braver, Chair

What patterns of brain activity reflect engagement with highly demanding cognitive tasks? How do these patterns relate to subjective, phenomenal effort? Answering these questions is critical to understanding what causes some people to experience cognitive tasks as more effortful than others. Subjective experience, in turn, is vital, with trait tendencies to exert effort having been linked to career and academic success. High subjective effort, as in schizophrenia and depression, can thus be extremely problematic. And yet, poor operational definitions have constrained research into basic questions about what neural dynamics track subjective effort. Here, a powerful, new behavioral economic operationalization is employed, in combination with fMRI, to investigate brain dynamics corresponding to subjectively costly cognitive effort.

Brain regions varying in activity by working memory load and cognitive control demands are strong candidates for tracking subjective effort (Westbrook & Braver, 2015). To identify such regions, I examined BOLD data, collected while participants performed a well-established working memory task (the N-back; Kirchner, 1958) that is both subjectively effortful, and for which subjective effort varies as a monotonic function

of load (Westbrook et al., 2013). I focused my search within independently-defined networks of nodes that co-vary (within-network) across a wide range of brain states. Specifically, I examined a subset of a priori “task-positive” networks, as identified by Power et al. (2011), which typically show increasing, and a “task-negative” network which typically shows decreasing activity with greater load. Importantly, variation was examined over N-back loads for which data has never been published, thus the present study reveals novel insights about activity-load functions in independently-defined functional networks from very low ($N = 1$) to very high loads ($N = 6$).

As expected, all task-positive networks showed robustly greater activity during the N-back. However, patterns of variation by load differed by network. While the task-positive fronto-parietal (FP), dorsal attention (DorAtt), and salience (Sal) networks showed inverted-U functions, peaking mid-range (at the 2- or 3-back) and decreasing after, the cingulo-opercular network (CO) showed robust activity that did not further vary by load. Rather than encoding load per se, the CO simply encoded that a participant was performing the N-back. The task-negative default mode network (DMN) was robustly and increasingly de-activated across all load levels examined.

Given that both subjective effort (Westbrook et al., 2013) and DMN deactivation are approximately monotonic functions of load, the DMN is a strong candidate for tracking variation in subjective effort with load. By contrast, inverted-U functions in the FP, Sal, and DorAtt networks do not straightforwardly map to monotonically increasing effort. Performance measures instead suggest that inverted-U functions tracked individual differences in adaptive strategy shifting. Namely, when participants were divided by 3-back performance, better performers showed a pronounced inverted-U (over $N = 1-3$)

while worse performers did not. Interestingly, a similar pattern was found when dividing participants according subjective effort, providing tentative support to a hypothesis that subjective effort acts as a cue to shift strategies adaptively under excessive demands. In any case, surprisingly, in none of the networks did load-specific changes in brain activity predict load-specific changes in subjective effort.

Critically, although load-specific patterns of brain activity did not predict subjective effort, load-independent brain activity predicted individual differences in subjective effort. Namely, higher average brain activity in any of the task-positive networks predicted greater subjective effort. At the sub-network level, this was notably true for two key regions that have been implicated as core components of a cognitive control system, and also hypothesized to track effort costs: the dorsal anterior cingulate cortex (dACC) and the dorsolateral prefrontal cortex (dlPFC) (McGuire et al., 2010). Importantly, after controlling for performance, the dACC remained a reliable predictor of subjective effort, while the dlPFC did not, supporting that the dACC tracks cognitive effort apart from task difficulty (while the dlPFC may not). This is consistent with strong prior theory implicating the dACC in regulating the intensity of cognitive control in response to flagging performance and in proportion to the expected value of doing so (Shenhav et al., 2013). The present results begin to answer basic questions about how the brain tracks subjective effort. They also lay the foundation for future work addressing why subjective effort can be so much greater for some individuals, like those with schizophrenia or depression, and also future work developing interventions for promoting desirable effort expenditure.

Chapter 1: Introduction

Trait tendency to expend cognitive effort reliably predicts academic and career success (Cacioppo, Petty, Feinstein, & Jarvis, 1996; Stumm, Hell, & Chamorro-Premuzic, 2011). There is also emerging evidence of significant differences in willingness to expend cognitive effort between older and younger adults (Hess & Ennis, 2011) and also between healthy individuals and those with depression (Hammar, 2009) and schizophrenia (Culbreth, Westbrook, & Barch, 2016; Gold et al., 2014). Willingness to expend cognitive effort may also be linked with more basic constructs including self-controlled intertemporal choice, proactive over reactive control, and model-based over model-free decision-making (Botvinick & Braver, 2015; Kool, McGuire, Wang, & Botvinick, 2013; A. R. Otto, Skatova, Madlon-Kay, & Daw, 2015). And yet, despite its central importance, we know very little about subjective cognitive effort. For example, what task engaged brain states are subjectively costly? Or, what states are phenomenally effortful? This knowledge gap is critical given that cognitive effort is frequently cited as a mediating factor in individual and group differences in physiological dynamics and task performance.

A reliable, precise, and theoretically consistent operational measure is critical to investigating cognitive effort (Westbrook & Braver, 2015). Such a measure should map closely to subjective experience, and not merely objective quantities of cognitive load, motivation, or performance, without resorting to potentially unreliable introspection and self-report, or circular reference to physiological outcomes about which effort is being inferred in the first place. The recently developed Cognitive Effort Discounting paradigm

(COGED) adopts behavioral economic techniques to quantify trade-offs between reward and effort costs, and is sensitive to numerous state and trait factors which influence subjective effort, and the extent of engagement with cognitively demanding tasks (Westbrook, Kester, & Braver, 2013). This dissertation reports the results of behavioral and fMRI studies that employ the COGED paradigm to investigate how the brain tracks subjective effort costs during engagement with demanding tasks. These results yield new insights about how to infer subjective effort from fMRI data, and lay the groundwork for future studies that may be aimed at targeted interventions for ameliorating deficient effort expenditure, as in psychopathology.

1.1 Brain Regions Tracking Effort

Hypotheses abound regarding the activity profiles (e.g. in fMRI BOLD data) of brain regions involved in tracking effort. These hypotheses stem chiefly from evidence of neural activity scaling with two core dimensions of effort: objective load and incentive motivation. The logic is that motivation should be encoded in dynamics that scale with effort benefits, while load should be encoded in dynamics scaling with effort costs. Here, I consider objective motivation to be the magnitude of an external reward (e.g. dollars), and objective load to be working memory demands needed support rule-guided behavior. Note that a growing body of work supports that tasks are phenomenally effortful when working memory is required for cognitive control (Botvinick, Huffstetler, & McGuire, 2009; Ewing & Fairclough, 2010; Kool, McGuire, Rosen, & Botvinick, 2010; McGuire & Botvinick, 2010; Westbrook et al., 2013). Cognitive control refers top-down signals that bias rule-guided behavior during pursuit of goals (Botvinick, Braver, Barch, Carter, & Cohen, 2001; E. K. Miller & Cohen, 2001; Ruge, Braver, & Meiran, 2009; Sakai,

2008; Yeung, Nystrom, Aronson, & Cohen, 2006). Core functions include updating and maintaining goal-relevant task sets thought to bias the mapping of stimuli to goal-appropriate responses. An example of objective load is N in the N-back task (Kirchner, 1958). The N-back is a well-studied paradigm in which individuals must identify serially-presented items repeated after N positions. Hence, purportedly effortful cognitive control is required to support the correct rule mapping (pressing “target” for arbitrary items repeated after N positions or “non-target” for all other items), and the number of stimuli that the individual must hold in working memory to perform this task scales with N.

It is important to note that while objective load and incentive motivation covary with effort, they should not be taken as synonymous with the construct and are useful only in that they can implicate candidate regions for tracking effort. Similarly, performance measure (errors and response times) also covary with effort (and load and incentive). In fact, by one theory, subjective effort arises from error detection by the dorsal anterior cingulate cortex (dACC), which simultaneously recruits resources to up-regulate control and sends an aversive learning signal in response (Botvinick, 2007). And yet, performance variables are also conceptually distinct from effort. For example, two individuals who are matched on performance, for the same objective load, in pursuit of the same incentive, may still feel differently about the amount of effort involved.

Regions encoding both incentive motivation and objective load are strong candidates for tracking subjective effort. Incentive motivation has wide-ranging effects; consequently, by itself it may offer little specificity in terms of the locations of action. Recent findings have suggested that encoding of reward magnitude can be found virtually throughout the cortex (Vickery, Chun, & Lee, 2011). In contrast, cognitive load, of the

kind described above, has somewhat more specific locations of action, particularly in the parietal and prefrontal cortices (Esterman, Chiu, Tamber-Rosenau, & Yantis, 2009; Etzel, Cole, Zacks, Kay, & Braver, 2015; Kool et al., 2013; McGuire & Botvinick, 2010; Miller, 2000).

The dorsolateral prefrontal (dlPFC) and dorsal anterior cingulate cortex (dACC) are two particularly strong candidates for tracking cognitive effort by these criteria (Botvinick & Braver, 2015; Pessoa & Engelmann, 2010). Both have been implicated as key substrates for cognitive control and working memory. An influential model hypothesizes that the dACC monitors for response conflict (e.g. during difficult discriminations) and flagging performance, and recruits the dlPFC, in proportion to conflict and errors, to broadcast top-down control signals that strengthen lower-level perceptuo-motor association pathways (Botvinick et al., 2001; Miller & Cohen, 2001).

Evidence implicating the dlPFC in working memory and cognitive control includes monkey unit recording studies in which task-relevant information, ranging from concrete stimulus features to abstract task rules, is reliably encoded in sustained delay period activity by dlPFC cells (Dick & Katsuyuki, 2004). Similarly, dlPFC BOLD signals increase reliably and parametrically with set size in the N-back in dozens of fMRI studies across multiple sensory modalities and stimulus types, reviewed in (Owen, McMillan, Laird, & Bullmore, 2005). For example, an early N-back imaging study found that BOLD response amplitudes in the dlPFC (Brodmann areas 46/9 and 44/45), increased with N over the range (N = 1—3) (Braver et al., 1997).

There is also evidence that the dlPFC is recruited volitionally: increasing demands interact with reward magnitude to predict increasing activity (Kouneiher, Charron, &

Koechlin, 2009; Locke & Braver, 2008; Pochon et al., 2002), and both reward and task difficulty showed overlapping but independent contributions to dlPFC recruitment during task engagement (Krebs, Boehler, Roberts, Song, & Woldorff, 2012). The logic here is that independent encoding (or an interaction) of load and incentive is strong evidence for volitional recruitment because recruitment should ramp up when an actor perceives increasing demands and particularly when they are motivated to perform well.

There is also evidence that dlPFC recruitment is phenomenally effortful. In one particularly relevant study, task-based dlPFC activity covaried with subjective ratings of desire to avoid a more demanding over a less demanding version of the task (versions differed by frequency of task-switching) (McGuire & Botvinick, 2010). Importantly, the relationship with avoidance ratings persisted after controlling for performance (RT and error rates), showing that dlPFC recruitment does not merely reflect task difficulty. Of course imperfect perception of performance by participants may account for why performance did not explain shared variance between subjective aversion and dlPFC activity. Or, more intriguingly, the remaining shared variance may reflect non-linear translations of dlPFC recruitment into subjective effort across loads and within individuals or between-individual differences leading some to experience the response to flagging performance as more effortful than others.

There are also numerous reasons to hypothesize that the dACC tracks subjective effort. For one, there are striking neuropsychological accounts linking ACC with phenomenal effort. One ACC-lesioned patient reported no phenomenal sense of effort during performance of a Stroop paradigm that control subjects reported as effortful, despite matching controls' performance and despite showing intact skin conductance

response to task engagement (Naccache et al., 2005). Intriguingly, a recent study found that direct electrical stimulation of the anterior mid-cingulate cortex induced a phenomenal state that participants described as a rising sense of a forthcoming challenge and a concomitant “will to persevere” in the face of that effortful challenge (Parvizi, Rangarajan, Shirer, Desai, & Greicius, 2013). Convergent evidence comes from an EEG study where participants performing a consciously effortful tone discrimination task showed individual differences in N1 amplitudes (localized to ACC) scaling with subjective ratings of effort (Mulert, Menzinger, Leicht, Pogarell, & Hegerl, 2005). Moreover, larger N1 modulation predicted better performance among a group reporting that they increased their effort the most when instructed to do so. Neurophysiological and neuropsychological data like this are also convergent with theories, like the one described above, attributing a role for the dACC in performance monitoring and cognitive control regulation. A recent extension of this theory proposes that the dACC, by virtue of a well-established role in decision-making, regulates the recruitment of cognitive control mechanism, like the dlPFC, in proportion to the expected value of exerting that control (Shenhav, Botvinick, & Cohen, 2013). Importantly, the dACC should, by this theory, track effort costs in the service of determining expected values. Nevertheless, in the McGuire and Botvinick (2010) study described above, although the dlPFC was found to encode subjective desire to avoid a demanding task even when controlling for performance, the dACC was not found to encode that same subjective desire, controlling for performance (McGuire & Botvinick, 2010). Hence, whether the dACC tracks effort during task engagement remains an open question.

Beyond the dlPFC and dACC, objective load and motivation are reflected by marked shifts in the physiological dynamics of wider cortical networks and neurotransmitter systems. Since the recognition of long-range functional coherence spanning the brain (Biswal, Yetkin, Haughton, & Hyde, 1995), considerable evidence has supported the existence of stereotyped functional networks that show coherently increasing or decreasing activity, depending on task engagement. Networks can be categorized broadly as: 1) a “task-positive” in that are more active during task engagement, covering the intraparietal sulcus (IPS), anterior insula (AI), supplementary motor area (SMA), and dorsolateral (dlPFC) and dorsomedial prefrontal cortex, and 2) a “task-negative” in that they are deactivated by external task engagement, covering the medial prefrontal cortex, posterior cingulate cortex / precuneus, and lateral parietal cortex (Fox et al., 2005).

The mean activity level of particular networks corresponds well to the phenomenology of effort. For example, among the 13 canonical brain networks defined by (Power et al., 2011), the task negative DMN, and the task-positive Fronto-Parietal (FP), Dorsal Attention (DorAtt), Salience (Sal), and Cingulo-Opercular (CO) networks, in particular, make strong candidates for tracking effort. First, they are modulated by task demands: regions within the task-positive FP, Sal, DorAtt and CO networks are more active, and within the task-negative DMN are less active with increasing load (e.g. the 2-back working memory task versus the 0-back) (Krebs et al., 2012; McKiernan, Kaufman, Kucera-Thompson, & Binder, 2003; Pyka et al., 2009). Second, task complexity and reward may contribute independently to activity in task-positive regions during task engagement, supporting a role for volitional recruitment (Krebs et al., 2012). Third,

“resting” dynamics following cessation of task engagement is affected by prior task complexity, suggesting a form of recovery that might also track effort costs. Specifically, DMN regions have been found to be more active after disengagement from more demanding tasks (Pyka et al., 2009) and take longer to recover their functional architecture (Barnes, Bullmore, & Suckling, 2009). Fourth, the degree of network modulation in task-positive regions predicts performance, thus providing a potential physiological mechanism linking effort and performance (Kitzbichler, Henson, Smith, Nathan, & Bullmore, 2011; Ossandon et al., 2011). Thus, like patterns of activity in the dACC and dlPFC, the dynamics of these five networks make them strong a priori candidates for tracking subjective effort.

1.1.1 Limitations of Cognitive Effort Literature

As indicated above, most studies of effort have focused on physical effort. A rich literature on human and animal studies have elucidated a network of regions involved in physical effort tracking and decision-making, in great detail. By contrast, only a handful of studies have focused on cognitive effort, and there are preliminary indications of both overlaps (Schmidt, Lebreton, Cléry-Melin, Daunizeau, & Pessiglione, 2012) and distinctions (Hosking, Floresco, & Winstanley, 2014) between the neural substrates mediating physical versus cognitive effort.

Studies investigating dynamics that track cognitive effort are particularly limited. Most prior work has implicated regions in effort incidentally in that they scale with cognitive load, or incentive motivation, as reviewed in detail in (Westbrook & Braver, 2015). For example, one study has provided fMRI evidence implicating the ventral striatum (VS) as a core hub of cognitive and physical incentive motivation. VS was

increasingly active in response to larger incentive cues, it predicted performance in both domains of effort, and the it showed different patterns of functional connectivity depending on the domain (Schmidt et al., 2012). Subjective effort costs, however, were not directly investigated.

The most direct study of regions tracking subjective cognitive effort, to date, was described above in which participants self-reported their desire to avoid a more demanding over a less demanding task (involving more frequent versus less frequent task switching, respectively) (McGuire & Botvinick, 2010). Self-reported desire-to-avoid ratings were then used as a predictor to test what brain activity patterns were related to the intensity of the desire to avoid demand. The key advantage of this methodology is that it investigates relationships between brain dynamics and subjective, phenomenal experience apart from objective load or performance indicators – both of which may covary with, but are conceptually distinct from cognitive effort. The chief limitation of this study, however, was that it relied on self-report measures that make strong assumptions about participants' ability to introspect and report experience in a way that is reliable and consistent across participants. Also, self-report ratings were based on categorical ratings of the intensity of desire to avoid a task, and as such may lack the precision of a continuous, quantitative measure of cognitive effort.

The present study builds on the strengths of that core methodology, by investigating how brain dynamics during effort expenditure map onto a subjective experience. However, as explained in the next section, subjective experience is quantified by patterns of decision-making that are potentially more objective and reliable than self-report. Moreover, this study investigates the relationship between brain dynamics and

subjective experience over a wide range of objective working memory loads. This design permits within-subject analyses about the way subjective experience changes with load, in addition to the between-subjects analyses of how individuals vary in their subjective experiences. Finally, the outcome measure of subjective effort used in this study is a continuous quantity that potentially affords greater precision than categorical distinctions.

1.2 Cognitive Effort Discounting (COGED) Paradigm

A core motivation of this dissertation is to investigate what brain states are experienced as subjectively effortful. The question can be framed in economic terms: what brain states are subjectively costly and thus discount the value of rewards pursued by task engagement? Arguably, subjective costliness is *the* cardinal feature of effort (Westbrook & Braver, 2015) and, as such, any region predicting subjective costliness during task engagement is also tracking subjective effort. Given that subjective effort likely co-varies with objective load, incentive motivation, and performance, the question then amounts to what brain states predict the subjective costliness of task engagement, controlling for objective load, incentive magnitude, and task performance.

Behavioral economists have used reward discounting paradigms to investigate subjective costliness of diverse outcomes. For example, risky rewards are subjectively less valuable than certain rewards. Discounting has been used by both behavioral- and neuro-economists to investigate decision making about risks and also delay and physical effort as well (Du, Green, & Myerson, 2002; Frederick, Loewenstein, & O'Donoghue, 2002; Green & Myerson, 2004; Huettel, Stowe, Gordon, Warner, & Platt, 2006; Jimura et al., 2011; Mitchell & Wilson, 2010; Myerson & Green, 1995; Peters & Buchel, 2010; Prelec & Loewenstein, 1991; Prévost, Pessiglione, Météreau, Cléry-Melin, & Dreher,

2010; Rachlin, Brown, & Cross, 2000; Treadway, Buckholtz, Schwartzman, Lambert, & Zald, 2009; Wardle, Treadway, Mayo, Zald, & De Wit, 2011).

Recently, I adopted this discounting approach by substituting working memory load as a cost factor and showed that preferences regarding cognitive effort can be measured in terms of subjective values (Westbrook et al., 2013). In the COGED paradigm, participants are familiarized with a multiple levels of a demanding cognitive task. They are then presented with a series of paired offers to repeat a more demanding level for more money or a less demanding level for less money. Offers are titrated until participants are indifferent between the offers. Indifference points are critical because they indicate psychophysical equivalence between greater reward (e.g. dollars) and greater effort, thereby rendering effort in terms of a common metric that can be quantified across task features (working memory load), incentive contexts (reward magnitudes), and across participants.

Subjective effort was investigated in this study using the well-studied N-back task (Braver et al., 1997; Jaeggi et al., 2003; Kirchner, 1958; Owen et al., 2005). The N-back is desirable for a number of reasons including that it has parametrically varying working memory load (N) and thus discounting can be quantified as a function of load. Most importantly, the N-back was selected because it is phenomenally effortful, and self-reported effort increases monotonically with N-back load (Ewing & Fairclough, 2010). Also, the N-back is a continuous performance task and thus participants select among levels because load is varying, not because time-on-task is varying. Controlling for time-on-task is critical if the intent is to control for the costliness of time when estimating the costliness of effort. An important caveat is that while task duration is fixed, the amount of

time spent on- versus off-task may vary. For example, a participant may be entirely occupied by a demanding 3-back task, but may afford moments of distraction between stimuli during the 1-back. Thus, response time will be included as a covariate in analyses to further control for load-based and individual differences in time-on-task.

Discounting functions comprise indifference points observed for multiple offer amounts across a range of N-back loads. In the current study, as in prior work, N-back levels 2—6 were used as high load conditions, which during decision-making are pitted against a low load 1-back condition, which is treated as the baseline (Westbrook et al., 2013). The indifference point function plotted against load thus describe the costliness of engaging in each level of the N-back, and moreover can be used to describe how costly participants experience the N-back relative to other participants. An example indifference point function is plotted for two hypothetical participants for loads $N = 2—4$ in Figure 1.1. Note that indifference points normalized by their base offers are converted into subjective value (SV), which ranges from 0 to 1. Also, note that a curve, joining SVs across all load levels, describes average discounting. Area under the curve (AUC) thus describes mean costliness of the N-back task for a given participant.




Figure 1.1 Subjective value bar plots for two hypothetical participants. Area under the curve (AUC) is indicated for each participant by color-coded lines. Participant A (light blue) discounts a task more shallowly and thus has a higher AUC than Participant B (darker blue), who discounts more steeply. Y-axis gives the indifference points (subjective values) and x-axis gives the corresponding N-back task level.

These between- and within-participant measures of effort costs can then be compared to brain activity when participants are engaged with the N-back to determine which regions track effort costs.

1.3 Overview of Study & Predictions

As described above, this dissertation reports the results of behavioral and fMRI studies that employ the COGED paradigm to investigate what brain dynamics track subjective effort during the N-back. The strategy to investigate this question was to first explore how a priori networks of interest, and sub-network loci (e.g. the dIPFC, dACC, and VS) varied as a function of objective load. Given that objective load is a core dimension of subjective effort, regions encoding load were implicated as strong candidates for tracking effort. Next, these regions were tested for whether their N-back activity patterns predicted subjective effort costs (as measured by the COGED paradigm), controlling for load and performance.

1.3.1 Predictions about regions tracking objective load

Block-wise N-back BOLD signal was predicted to vary as a function of load in all a priori regions of interest, including the dlPFC and dACC, as well as task-positive FP, Sal, CO, and DorAtt networks, and the task-negative DMN. The shape of load functions in each of these regions, however, was uncertain. At least three types of load functions were anticipated from prior literature: monotonically increasing (or decreasing, in the case of the DMN) across loads, non-monotonic functions which either asymptote at higher load levels, or show bi-directionality (e.g. inverted-U shapes) across loads, and load-independent functions, which show non-zero, but flat activity across loads. No prior studies have investigated N-back activity functions in a priori functional connectivity networks as defined by Power et al. (2011), nor have they investigated N-back load functions beyond $N = 3$. So, the present dataset provides novel information about the activity of coherent functional networks, across a novel range of load levels. For example, prior evidence suggests that regions overlapping the FP network, particularly in the dlPFC, show inverted-U patterns, while CO network regions show monotonic increases over the range $N = 1-3$ (Callicott et al., 1999; Cappell, Gmeindl, & Reuter-Lorenz, 2010; Jaeggi et al., 2003; 2007). However, it was not clear whether these patterns would hold at the network level. Moreover, regions which appeared monotonic across $N = 1-3$ might have asymptoted at higher loads ($N = 4-6$), and thus proven to be non-monotonic instead. Preliminary predictions were that all task-positive networks would show non-monotonic response profiles across all load levels, with the exception of the CO network, which would increase monotonically with load. The task-negative DMN was predicted to decrease monotonically with load, cf. (Pyka et al., 2009).

The reasons for inverted-U shapes in task-positive regions is unclear, and has been interpreted as indicating excessive demands for cognitive capacity (Callicott et al., 1999; Jaeggi et al., 2003). One hypothesis is that inverted-U shapes in task-positive regions reflect adaptive strategy shifts, when demands become excessive. A key piece of evidence is that participants showing better 3-back performance had sharper inverted-U dlPFC load functions, and thus lower 3-back activity, compared with those showing worse 3-back performance (Jaeggi et al., 2007). As described in the Methods section, participants were pre-selected in the present study to restrict range on N-back performance. Nevertheless, a replication of the prior result was predicted such that a median split on 3-back performance would result in a high performing group showing sharper inverted-U load functions in task-positive networks (particularly the FP, Sal, and DorAtt networks) over $N = 1-3$, as compared with a low performing group, for whom the load function would appear more closely monotonic over the same range.

1.3.2 Predictions about networks tracking subjective effort

As described above, regions tracking subjective effort should covary with the subjective value (SV) of effort-discounted rewards (and conversely effort costs, as measured by the COGED paradigm), controlling for objective load, and performance. For a region to track SV, independently of load and performance, BOLD signal should either be predicted by a separate main effect of effort costs, or an interaction between effort costs and load. Hypothetical effort cost and load effects on BOLD signal are diagrammed in two hypothetical regions, one monotonic and another inverted-U, in Figure 1.2.

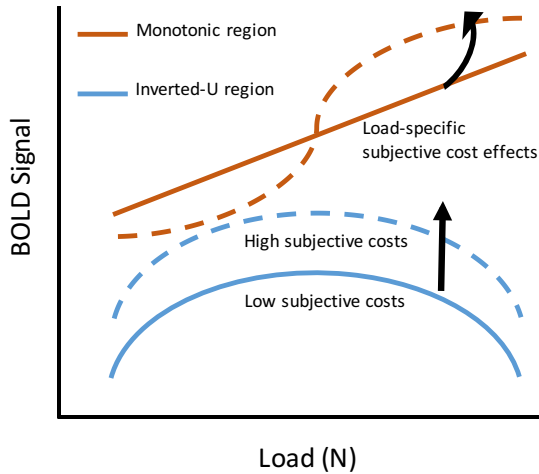


Figure 1.2 Hypothetical effects of effort costs and N-back load on BOLD signal in two task-positive regions. Blue lines correspond to an inverted-U load function and show a main effect of subjective costs on brain activity while orange lines correspond to a monotonic function and show a load by effort cost interaction. Solid lines indicate low subjective effort costs and dashed lines indicate high subjective effort costs.

A hypothetical main effect of subjective costs on BOLD signal in an inverted-U region is diagrammed by the two blue lines in Figure 1.2. The effect obtains as mean, load-independent increases for those experiencing the N-back as more subjectively costly. Note that even though SV is a monotonic function of N-back load (Westbrook et al., 2013), and activity in an inverted-U region is not, this region can still track (load-independent) individual differences in subjective task costliness. A hypothetical interaction of subjective costs and objective load on BOLD signal is diagrammed by the two orange lines. Note that in the example, the interaction is non-linear, reflecting load-specific changes in SV. The effect is such that activity in the monotonic region increases with load, and it does so more steeply as load-specific effort costs increase more steeply (e.g. in the middle of the load range). Also note that the assignment of a load-independent main effect to an inverted-U region, and a load-specific interaction to a monotonic region is arbitrary and for illustrative purposes only. Either type of region could, in principle, show either type of effect.

Discounting measures afford critical, novel inferential traction regarding the relationship between activity dynamics and subjective effort. The idealized predictions made in Figure 1.2 are that greater mean activity, or more steeply increasing activity in task-positive regions during the N-back will relate to greater subjective effort, as defined by COGED AUC. Alternatively, lower activity in task-positive networks among those finding the task costlier (for example) would imply that such networks track diminished willingness to expend costly effort. As shown in Figure 1.3, individual differences task-positive network activity can be combined with individual differences in SV to infer the meaning of those differences in activity for either motivation or subjective costliness. Those results lying across the axis from the upper-left to the lower-right quadrants support an inference about individual differences in brain activity being related to volition, while those on the axis from lower-left to upper-right support inferences linking brain activity to subjective costliness. The current study provided the opportunity to test which of these alternatives best describes individual differences in load-related activation.

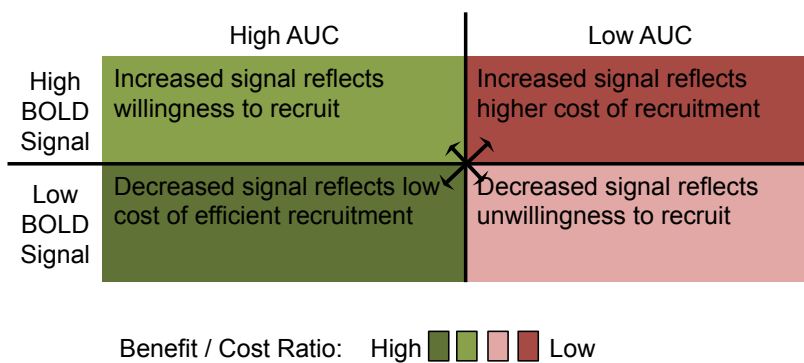


Figure 1.3 Combination of discounting and BOLD signal in task-positive networks for specific inferences about the relationship between brain dynamics and subjective motivation for effort, or effort costliness. Diagonal arrows indicate the volition and cost axes, respectively.

1.3.2 Predictions about subjective effort tracking in the dACC, dlPFC, and VS

At the sub-network scale, the dACC, dlPFC, and VS were also predicted to track subjective effort. As discussed above, prior hypotheses include that the dACC monitors response conflict and performance and, when these are detected, recruits the dlPFC for maintaining task rules and biasing lower-level processing pathways (Botvinick et al., 2001). Moreover, the dACC is hypothesized to regulate dlPFC recruitment based on the expected (and subjective) value of cognitive control, taking into account effort costs (Shenhav et al., 2013). In addition to these cortical loci, the VS is also a strong candidate for tracking subjective effort. VS activity increases monotonically with load for N-back levels $N = 0$ to $N = 2$ (Satterthwaite et al., 2012). Monotonic scaling with working memory load is consistent with the hypothesis implicating the VS in numerous processes including value-learning about states and actions and in value-based gating of items into working memory via cortico-striatal loops and the dorsal striatum (Badre & Frank, 2012; Chatham, Frank, & Badre, 2014; Frank & Badre, 2012; O'Reilly & Frank, 2006). Studies also support cognitive effort costs encoding in the VS: VS activity reflecting receipt of reward after a demanding task was attenuated in proportion to prior cognitive task demands consistent with the hypothesis that the VS encodes cognitive effort-discounted reward values (Botvinick et al., 2009).

SV is a monotonic function of load, and yet the dlPFC and dACC are part of two networks that were predicted to show non-monotonic, inverted-U functions (the FP and Sal networks, respectively). Thus, these regions were anticipated to either show load-specific interactions, or load-independent effects of subjective effort on brain activity (as depicted in Figure 1.2). As described above for networks, the direction of the prediction

was that higher (load-independent) activity, or steeper (load-specific) changes in activity would predict higher subjective effort costs. Load-independent effects, for example, were predicted between individuals such that those finding the N-back task costlier would show greater activity in the dlPFC and dACC. Predictions about variation in VS signal by subjective effort were less certain. One straightforward prediction was that increasing objective demands would be encoded in monotonically increasing VS activity (extending previous observations over $N = 0-2$ (Satterthwaite et al., 2012)), reflecting the intensity of cortico-striatal working memory gating processes. A further prediction was that the monotonically increasing load function would increase more steeply with load-specific increases in subjective effort.

Performance also co-varies with load, thus it was important to test whether these regions predict SV, controlling for performance measures. As described above, a prior study has shown that dlPFC, but not the dACC tracks subjective effort, controlling for performance (McGuire & Botvinick, 2010). Yet, given the hypothesis that dlPFC recruitment scales with declining performance, it was instead predicted that shared variance between dlPFC activity and SV would be explained away by performance measures. Hence, I predicted the opposite result – that the dACC, but not the dlPFC would covary with subjective effort after controlling for performance. Regarding the VS, one prior study has shown that cues indicating accurate responding during the N-back are encoded in increasing VS activity (Satterthwaite et al., 2012). Thus given prior evidence that task-engaged activity scales with performance (like in the dlPFC), VS activity was not predicted to co-vary with subjective effort apart from performance.

Chapter 2: Methods

2.1 Procedure Overview

The core procedure of the COGED paradigm involves having participants experience multiple levels of the N-back task (Phase 1), and then making COGED decisions to estimate SV (and therefore subjective costs) of the task levels they just performed (Phase 2; Figure 2.1). After participants make COGED decisions, one of their choices is selected, at random, for participants to complete again (“N-back re-do”) in exchange for the reward amount selected on that trial (Phase 3).

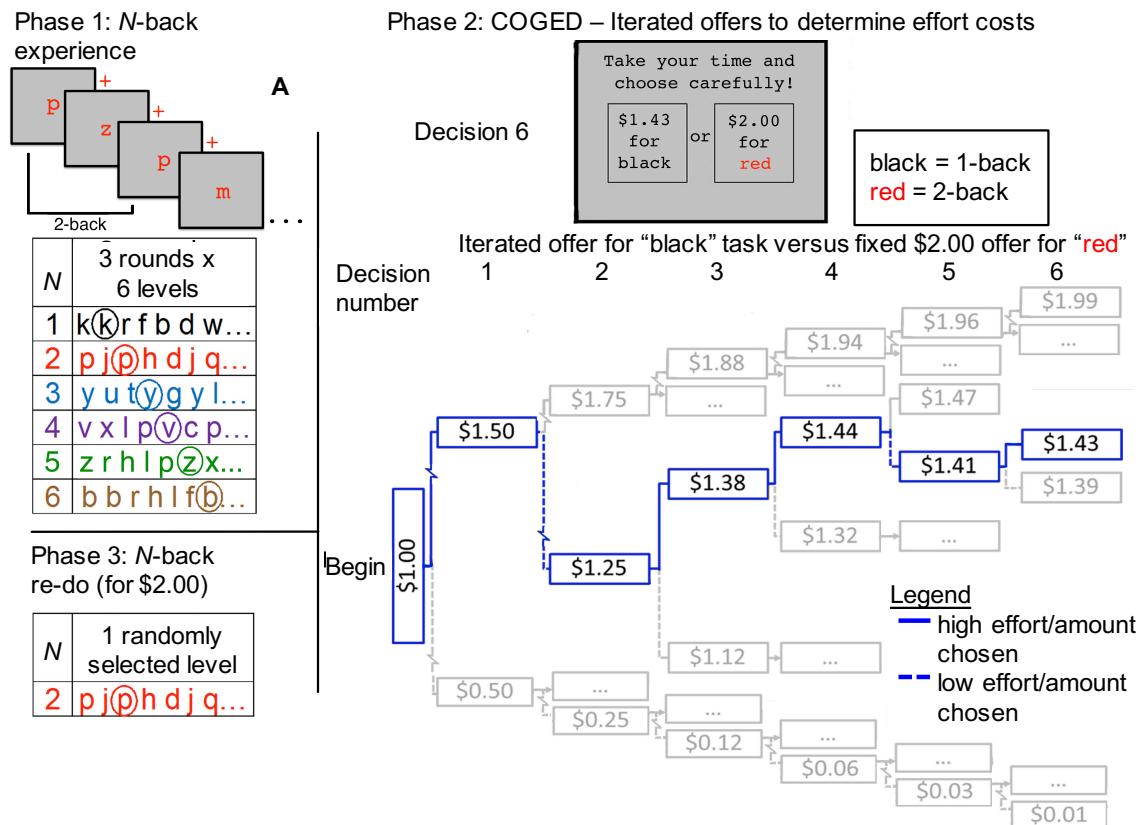


Figure 2.1 Schematic of the three phases of the COGED paradigm.

As shown in Figure 2.1, this same core procedure was repeated for all participants over three separate sessions (Figure 2.2).




Figure 2.2 Three session procedure overview schematic. Components completed in the fMRI scanner are highlighted in red font.

The N-back task was scanned during the second session, and data from that session is the focus of this dissertation. Data from a third session was used in a separate investigation, reported in Appendix A, of brain activity during cognitive effort-based decision-making. The first session was a purely behavioral session designed to identify candidates for the two scanning sessions based on two criteria. Participants were selected based on 1) high N-back performance at all loads and 2) either steep or shallow effort-based discounting (specific thresholds are described below).

2.2 General Task Descriptions

2.2.1 N-back

N-back tasks were presented as a series of consonants in 32 point Arial font in a color corresponding to the level of the task: namely, black (rgb code [0,0,0]) for 1-back, red [240,0,0] for 2-back, blue [0,0,255] for 3-back, purple [95,0,115] for 4-back, green [0,110,0] for 5-back, and brown [102,51,0] for 6-back. Tasks were presented in specific colors so that, during the decision-making phase, they could be referred to in terms of their color rather than the N-back load (e.g., as a choice between performing the red vs. black task, rather than 1-back vs. 2-back). This terminology was chosen to minimize anchoring effects (Chapman & Johnson, 1999). If, for example, the tasks were referred to by their load values (N) during decision-making, this could have biased participants to

demand twice as much for the 2-back versus the 1-back (or three times as much for the 3-back, etc.).

Each trial run began with a fixation symbol (underscore ‘_’) presented centrally in black on a grey screen [200,200,200] for 25 s before each trial run. Next, N-back stimuli were presented in the center of the screen for up to 2 s during which participants could respond by button press whether each stimulus was a “target” or “non-target”. If participants responded in under 2 s, the letter was instantly replaced by the fixation symbol for the remainder of the trial before the next stimulus appeared 2 s after the previous stimulus first appeared (fixed 2 s ISI). If participants did not respond in 2 s, the stimulus was simply replaced by the next item in sequence and the trial was marked as incorrect. N-back lists were 64 items long. They contained 16 targets, and a variable number of lures, depending on the task level (8 for the 1-back, 6 for the 2-back, 5 for the 3-back, and 3 for the 4-, 5-, and 6-back, each) where a lure is considered to be any stimulus repeated within two positions of the target position. The key reason for reducing the number of lures for higher load levels was to attempt to “flatten” performance functions – attenuating differences in performance from lower to higher load levels.

2.1.2 COGED

COGED trials were presented as a series of offers designed to identify subjective indifference points (and therefore SV; Figure 2.1). The COGED paradigm identifies the point of subjective indifference between a larger offer for doing a more difficult task, and a smaller offer for doing a less difficult task, using a procedure identical to that described in Westbrook, Kester, and Braver (2013). That is, indifference points were identified by stepwise titration of offers, in which stepwise adjustments were half as large on each

subsequent trial for a total of five adjustments for each task-amount pair. In the first session (Figure 2.2), base amounts were drawn from the set [\$2, \$4, \$5], depending on the particular experimental parameters in which the participant was first run. In the second and third sessions, base amounts were [\$2, \$3, \$4], for a total of 3 amounts by 5 tasks, or 15 task-amount pairs. 5 trials are used to identify indifference points for each task-amount pair for a total of 75 decision trials.

2.2 Participants

Twenty-five participants were recruited from the Washington University community, primarily through Experimentix. All participants were healthy, young (ages: 18 – 40), right-handed, neurologically normal, and not taking any psychoactive medications. All participants also had normal-to-corrected vision. Participants were further selected to fit one of two group profiles: 1) steep discounters (13 participants; $AUC < 0.55$) or 2) shallow discounters (12 participants; $AUC > 0.83$). AUC cutoffs were derived from the distribution of AUC values from prior COGED studies, conducted on similar populations (primarily Washington University undergraduates), such that they reflect the upper and lower tertiles of the typical AUC distribution. Prior COGED studies conducted without AUC-based selection have shown the typical AUC distribution to be strongly skewed negatively, with most participants showing relatively shallow effort discounting (high AUC values), and a smaller proportion with steep effort discounting (low AUC values). Thus it was hoped that pre-selection of an equal number of high and low AUC individuals for recruitment into the study would lead to a more symmetric (and bimodal) distribution and also afford extreme-groups contrasts. As discussed further below, and in Chapter 3, this approach was only partially successful, due to intersession

variability in discounting. Subsequent descriptions (e.g. in the Results section) of “high AUC / shallow discounters” or “low AUC / steep discounters” refer to participants above or below a median split in observed AUC, averaged across all three sessions, rather than to the initial selection criteria.

Note that for the purposes of investigating the encoding of subjective costs, discounting was averaged across all three sessions in Figure 2.2 to enhance the reliability of SV and AUC as trait measures of subjective effort. The hypothesis that SV and AUC reflect trait sensitivity to cognitive effort costs is supported by prior work showing that AUC relates to other trait variables including aging, Need for Cognition, and delay discounting (Westbrook et al., 2013), and also negative symptoms in schizophrenia (Culbreth et al., 2016).

Participants were further selected for having high levels of performance across all levels of the N-back task. Performance was quantified by the discrimination index d' and were based on (arbitrary) thresholds of $d' \geq 1.0$ (6-back), 1.25 (5-back), 1.5 (4-back) and 1.75 (3-back). The purposes of these thresholds were to ensure that all participants were fully engaged with all levels of the N-back task and to restrict performance differences between individuals and groups.

While the selection criteria (participants showing universally high N-back performance and either steep or shallow discounting behavior) were artificial, they were intended to increase power to demonstrate the principle that individuals could differ in terms of subjective effort, while being matched in terms performance. However, the non-random selection process was, by design, unrepresentative of the wider population, thus limiting generalizability of conclusions.

It is also important to note that although one of the core questions of this study is how brain activity differs among individuals, the number of participants is not powered for a single individual differences fMRI study. As mentioned above, this study investigates a dataset collected as part of a larger series of three sessions (one behavioral and two imaging sessions), that was originally designed to investigate both how the brain tracks subjective effort and how it supports cognitive effort-based decision-making. In the larger series of studies, the same set of participants participated in all three sessions. Thus when examining whether brain activity tracking effort is related to brain activity during effort-based decision-making, participants could be compared to themselves across sessions. The number of participants was thus optimized for the original design in which there would be 75 participant-sessions (3 sessions for 25 participants), and 50 participant-imaging sessions. Follow-up studies would benefit from a larger sample size to increase power to detect individual differences for a single imaging session.

2.3 Imaging Procedure

Following consent and screening to ensure MR compatibility, participants were stripped of any metal and brought into the scanning room. T1 and T2 anatomical scans followed localization and alignment scans. Then participants underwent the first of two 8.5 min resting state scans (which were not within the scope of the dissertation, and so were not analyzed or reported here). Immediately after the first resting state scan, the first of six runs devoted to N-back performance were initiated. All blocks of a given N-back level were performed within the same scanning run. Each run consisted of alternating N-back task (64 items \times 2 s + 5 sec post-block performance feedback, indicating accuracy and RT for a total 133 sec) and resting fixation (30 sec) blocks. Three task blocks and 4

fixation blocks were included in each run, for a total run duration of 519 seconds, or 8.65 min.

Between each run, participants were reminded about load-specific task instructions. N-back tasks were always performed in order of increasing demand so that all participants would have the same experience with each level in terms of sequential order. Note that fixed task order introduces confounds for between-load comparisons. The tradeoff is that fixed order ensures, to the extent possible, that all participants have same experience for a given level of the task. This was done so that endogenous individual differences were emphasized relative to differences in external features of the paradigm including, for example, differences in practice effects or accumulated fatigue for a given N-back level. After completing three blocks of each level, a second 8.5 min resting scan was conducted, participants were removed from the scanner, and asked to complete a round of COGED to establish their indifference points for each of the N-back levels for that session.

2.4 Scanning Parameters

All fMRI data were collected in a 3 Tesla Siemens Trio scanner. Anatomical T1 images were collected in 176 frames of $1 \times 1 \times 1$ mm voxels using 2.4 s TRs, and spin-echo times of 3,080 ms, and an 8 degree flip angle. Anatomical T2 images were also collected in 176 frames of $1 \times 1 \times 1$ mm voxels using 3.2 s TRs, spin-echo times of 455 ms, and a 120 degree flip angle. Functional imaging sequences during resting state and task engaged state scans were collected in $4 \times 4 \times 4$ mm voxels using a 256×256 voxel field of view, 2,000 ms TRs, 27 ms spin-echo times, and 90 degree flip angles. For N-back imaging runs, 260 volumes were collected for each level of the N-back task.

2.5 Image Processing

All images were processed and statistical analyses conducted using AFNI. Raw DICOM images were first converted to NIFTI format using the Freesurfer `mri_convert` function, and the AFNI `3dSkullstrip` function was used to mask brain tissue from the surrounding skull.

Functional images were concatenated using `3dTcat`, aligned from oblique to cardinal orientation using the `3dWarp` function, and then upsampled from $4 \times 4 \times 4$ mm voxels to $3 \times 3 \times 3$ mm voxels and aligned across all functional runs to the first run. Next, parameters for registration of functional volumes with anatomical T1 images were computed for each participant separately. Precise registration was verified visually for every participant and cost functions were tailored to optimize registration for each participant. Then, parameters for warping participant-specific anatomical images to a standard MNI space (`MNI152_T1_2009c+tlrc`) were computed. All registration and warping parameters were concatenated using the `cat_matvec` function, and applied as a single transformation to aligned functional image volumes using the `3dAllineate` function. Extents masks were computed for each participant as a further check to ensure good resulting alignment of the transformed datasets.

Following these transformations, functional images intensities were scaled such that each voxel had a mean value of 100 and were restricted to the range $[0, 200]$. Specifically, after normalizing all time points for a given voxel by 100, step functions were multiplied by all values across the time series, thus constraining the products to a maximum of 200 and a minimum of 0.

Next, functional images were smoothed using an 8.0 mm FWHM kernel and the `3dmerge` function. Then, subject brain masks were computed from anatomical images

using the 3dmask_tool and dilate and erode parameter values of 5 voxels each, while also filling holes. This mask was then correlated with functional image intensities to further ensure good alignment.

General linear models (GLMs) were used to investigate the relationship between voxel intensities and task events using the 3dDeconvolve function. Specifically, GLMs modeled activity during performance each level of the N-back task as a block design, with boxcar functions each spanning the 128 sec duration of a given N-back stimulus list, convolved with a gamma function. Note that GLMs incorporated motion censoring with a frame displacement threshold of 0.3 mm as well as 6 motion regressors: roll, pitch, yaw, and x, y, and z translations. All runs were inspected manually to confirm satisfactorily low levels of motion. The mean fraction of censored frames was 5.1% with a median of 3.1% and a range of 0.5% to 19.1% across 24 participants. GLMs also incorporated polynomial regressors to control for low-frequency trends, depending on the duration of a particular run. Note that although datasets are concatenated for preprocessing in AFNI, separate sets of polynomial regressors are applied for each imaging run. Thus, low-frequency trends are not confounded with block order or N-back load.

Chapter 3: Behavioral Results

3.1 N-back Performance

A hallmark feature of the N-back task is that objective load can be varied parametrically by N. Thus, performance was predicted to decline with increasing N. Specifically, the ability of participants to distinguish targets and non-targets should diminish, and participants should respond more slowly as demand rises. Importantly, while performance was predicted to vary with N, performance was predicted to be unrelated to subjective effort (e.g. as measured by AUC). As described in the Methods, participants were intentionally selected to constitute a sample with a wide range in discounting variance, but to restrict range on individual differences in performance. Although SV, as measured by COGED, has been related to performance when participants are selected pseudo-randomly (Westbrook et al., 2013), the intent of this study was to demonstrate that differences in subjective effort could exist apart from differences in performance.

N-back performance remained consistently high across all levels of the N-back task, yet performance patterns reveal that task load parameter N influenced task demands as anticipated. Figure 3.1 shows that the performance measure d' is both high and monotonically decreasing with N. Importantly, High AUC (shallow discounting) and Low AUC (steep discounting) participant groups were well matched on performance across all task loads in terms of d' . There are no reliable group differences at any of the N-back task loads. Formally, a Load x Group ANOVA indicates that d' varies linearly with Load ($F_{1,140} = 204.8, p < 0.01$) but does not vary by Group ($F_{1,140} = 0.20, p = 0.66$)

and there is not a reliable Group x Load interaction ($F_{1,140} = 0.45, p = 0.50$). Note that d' measures were adjusted by the log-linear transform to account for extreme hit rate and false alarm proportions (Hautus, 1995). This transformation adds small, non-zero amounts to every target hit or false alarm when determining hit and false alarm rates so that perfect hit rates, or false alarm rates (otherwise equal to 1 and 0, respectively) would not equal +/- infinity when converted into the inverse of the cumulative Gaussian distribution function. The transform has little effect on d' scores, except at these limits.




Figure 3.1 Group performance d' by load and by AUC group for both Session 2.

While d' provides a response bias-free measure of performance, performance on “lure” trials may be particularly revealing as to strategy (e.g. explicit, phonological loop representations of stimuli sequences versus familiarity-based recognition, cf. (Juvina & Taatgen, 2007), and such strategy shifts are likely to occur at the high load levels ($N \geq 3$) used in these experiments. As show in Figure 3.2, the percentage of false alarm lure trials (“lure rate”), varies with load ($F_{5,137} = 8.70, p < 0.01$), appearing to increase up to $N = 3$, and then remaining constant after. Post-hoc pairwise comparisons reveal that the only reliable differences were between a higher lure rate for levels $N \geq 2$ as compared with $N = 1$ ($p < 0.01$). Also, the lure rate is higher for $N = 3$ than $N = 2$ ($p = 0.03$), and the lure

rate is trending higher for $N = 6$ than $N = 2$ ($p = 0.08$). Interestingly, though the lure rate rises up to $N = 3$, it does not rise after that, suggesting that participants relied on familiarity to the same degree at extremely high load levels. Indeed, overall, lure rates are on average very low, across all load levels. Given that a familiarity-based strategy is susceptible to lures, low lure rates suggest that participants are not relying solely on familiarity to perform the N-back at any level. As noted in the Methods, however, there was a confound in that the number of lures decreased (from 8 at $N = 1$ to 3 by $N = 4$ —6). As such, there were fewer opportunities for participants to make lure errors at higher loads. This confound prevents strong inferences about the response to lures across loads. Importantly, however, as was shown for d' , there is no reliable group difference in lure rates ($p = 0.27$) supporting that participants in both groups relied on familiarity to a similar extent.

The fact that d' continued to fall at higher loads ($N > 3$) while the lure rates stalled indicates that the falling discrimination index reflects a progressive decrease in target hit rates.




Figure 3.2 Lure rates across all loads and both groups for Sessions 2.

Response times (RT) analyses (of correct trials) also support that higher load levels are more demanding, and also that there were no systematic differences in how High and Low AUC groups performed. RT distributions are typically skewed, and thus not well described by a Gaussian curve. Ex-Gaussian analyses fit a three parameter model to RT distributions, describing the central-tendency with the parameter μ_{RT} , the variability around that point with σ_{RT} , and the magnitude of the right-skewed tail with τ_{RT} . Critically, while these parameters are sensitive to Load, there were no differences by Group. μ_{RT} was statistically indistinguishable across Groups ($F_{1,137} = 0.14, p = 0.71$). This is also true when using the tail of each groups' RT distribution: τ_{RT} , ($F_{1,137} = 1.35, p = 0.25$). σ_{RT} also does not vary by Group ($F_{5,143} = 1.56, p = 0.18$).




Figure 3.3 μ_{RT} across loads and groups for Session 2.

To demonstrate that the Ex-Gaussian analysis does not distort the result that groups did not differ in terms of response times, the average and standard deviation of median response times for every load and both groups is provided in Table 3.1. Mirroring the Ex-Gaussian analysis, the groups are also very well matched in terms of median RT values.

		Average (and SD) of RT Median by N-back Load (s)					
		1-back	2-back	3-back	4-back	5-back	6-back
Low AUC		0.55	0.61	0.65	0.64	0.62	0.61
		(0.07)	(0.13)	(0.15)	(0.17)	(0.12)	(0.13)
High AUC		0.55	0.68	0.69	0.65	0.61	0.60
		(0.08)	(0.13)	(0.17)	(0.16)	(0.13)	(0.12)

Table 3.1 Average and standard deviation of median response times by AUC group and by N-back load.

This dataset successfully demonstrates the principle that subjective effort can be decoupled from, and is thus not redundant with, performance. In other words, some find the N-back task to be more subjectively costly, even if they perform the task equally well. This dissociation of performance and subjective effort was also critical for subsequent analyses examining individual differences in brain activity that are associated with differential subjective effort, controlling for performance.

It is important to note, however, that the dissociation of subjective effort and performance obtained, in part, because of an artificial selection process maximizing one sort of variance (in subjective effort) and minimizing another (in performance). Thus, these data do not support the inference that subjective effort and performance are generally unrelated in the wider population. Indeed, as noted above, performance and subjective effort, as measured by the COGED paradigm, have been related in other studies (Westbrook et al., 2013). Also, as discussed in the Introduction, others have hypothesized that flagging performance is related to the subjective effort costs because the detection of errors and cognitive conflict yields an aversive learning signal that could drive learning of an effort cost function (Botvinick, 2007).

An important question is whether performance changed when participants completed “re-do” trials (see Figure 2.1), following COGED decision-making, when they were asked to repeat one level of a task based on random selection from among their

choices. Given that participants were instructed that they would be paid “regardless of performance” and that they must simply “maintain their effort” from prior levels, there was some chance that participants selected a more demanding level for higher pay simply because they did not believe that their effort level would be monitored, and therefore they could earn pay without actually exerting effort. This could confound COGED analyses, as indifference points would not accurately measure a participants’ actual willingness to repeat the N-back for monetary reward. However, participants, on average, performed better on re-do trials than they did in the prior practice. On average, d' increased on re-do trials: 0.58 units across the low AUC group, and 0.22 across the high AUC group. Performance remained high across all levels that were repeated. Importantly, high and low AUC groups did not differ in how their performance shifted between early N-back exposure and re-do trials ($p = 0.19$). This result supports that high AUC individuals took re-do performance just as earnestly as low AUC individuals, and contradicts the hypothesis that high AUC individuals showed little discounting because they did not intend to exert effort on N-back re-do trials.

3.2 Decision-Making Behavior

Based on prior studies, three kinds of results were expected: 1) participants would discount rewards for performing all higher levels of the N-back ($N = 2\text{—}6$) relative to the lowest level ($N = 1$), 2) the SV of an offer to perform a given level of the N-back to earn a reward decreases as load increases, and 3) the SV increases as the base reward amount (offered for performing the harder task) increases (Westbrook et al., 2013). As shown in Figure 3.4, discounting was reliable and monotonic, such that SV reliably decreases with load across all levels of the N-back task. That is, participants discounted monetary

rewards at all levels of the N-back ($N \geq 2$), relative to the 1-back. Moreover, the decrease in SV with load demonstrates that, as expected, participants found increasing N-back demands to be increasingly costly. Though the data suggest a small trend of increasing SV with larger amounts, the anticipated increase in SV with amounts is not reliable. Formally, Load x Amount ANOVAs reveal reliable effects of Load ($F_{4,352} = 50.6$, $p < 0.01$), but no effects of Amount ($F_{2,352} = 0.37$, $p = 0.69$). Amount effects have been documented in other domains like in delay discounting (Estle, Green, Myerson, & Holt, 2006; Green, Myerson, Oliviera, & Chang, 2013). Amount effects were not reliable, however, so discounting rates were collapsed across offer amounts for subsequent analyses.

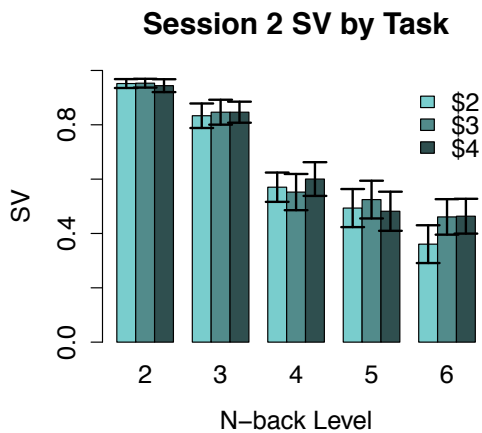


Figure 3.4 Subjective values across N-back task levels and base offer amounts for COGED decision-making in Session 2.

AUC quantifies the area under line segments connecting SV across N-back levels, and hence provide a single averaged discounting measure for a subject (see Figure 1.1). As Figure 3.5 shows, AUC varies across sessions for participants – AUC values do not lie on a line with unity slope. However, the AUC of one session was generally predictive of the AUC of the next session. A linear model of Session 1 AUC predicting Session 2

AUC gives a reliable regression coefficient ($B = 0.62$, $p < 0.01$). Likewise, Session 2 reliably predicts Session 3 AUC ($B = 0.57$, $p < 0.01$). The Pearson correlation (including all 25 participants) between Session 1 and Session 2 AUC is $\rho = 0.61$, between Session 2 and Session 3 is $\rho = 0.64$, and between Session 1 and Session 3 is $\rho = 0.28$. The ICC for all three sessions is thus 0.47 with 95% CI of [0.23, 0.69]. One interpretation of the observation that high AUC participants tend to remain high AUC participants, and low AUC participants tend to remain low AUC participants, is that AUC reflects a trait measure of subjective effort costliness on the N-back task. However, this trait measure is not perfectly predictive; state also influences the extent of a participant's discounting in a particular session. This is not entirely surprising given that at least one study has shown that sleep deprivation state can affect cognitive effort discounting (Libedinsky et al. 2013). Of course it is possible that variability might also reflect measurement noise. To my knowledge, cross-session discounting stability in other domains, (e.g. delay discounting), has rarely if ever been studied. As such, it is unclear how reliable are discounting paradigms.

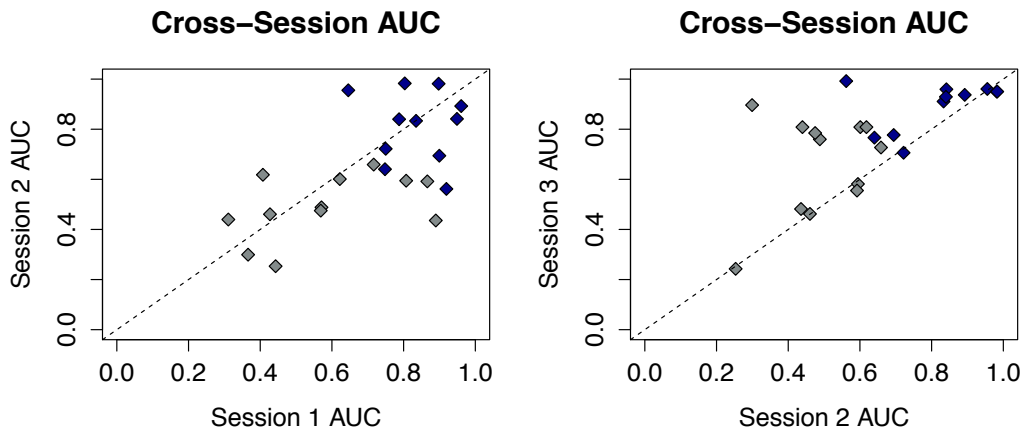


Figure 3.5 Pairwise COGED AUC plots for Sessions 1 and 2 and Sessions 2 and 3. Color indicates membership into High AUC_{3S} (navy) and Low AUC_{3S} (grey) groups.

Interestingly, one hint of state influences on discounting is that discounting was reliably shallower (AUC values are higher), across participants, on Session 3 than on Session 2 ($p_{\text{paired}} < 0.01$; other inter-session comparisons are non-significant). Participants experienced 2 runs of each N-back level in Session 3 versus 3 runs in Session 2 and Session 1. Thus, the observed AUC difference is consistent with past piloting projects that have also yielded shallower discounting following more brief, in-session N-back exposure, and supports the hypothesis that fatigue may be an important state factor influencing discounting behavior in COGED. Given inter-session variability in AUC, and to maximize the reliability of AUC as a trait measure, all subsequent analyses averaged AUC values across all three sessions (AUC_{3S}), to provide an AUC score for each participant, except where explicitly noted (as AUC). Subject assignment to High AUC and Low AUC groups depends on whether a participant's AUC_{3S} falls above or below the median AUC_{3S} . Note that in Figure 3.5, final group assignment is color-coded (grey for Low AUC_{3S} and navy for High AUC_{3S}).

Although participants were selected to show either steep or shallow discounting in an extreme-groups design, the final distribution (of AUC_{3S} values) was uni- rather than bi-modal. Figure 3.6 provides a histogram of the cross-session averaged AUC_{3S} to show that the final group distribution was not bimodal as originally intended. Moreover, the low AUC group is skewed, with most participants close to the sample median. This does suggest some caution in interpreting group analyses based on AUC_{3S} (i.e., high vs. low). Consequently, subsequent analyses supplanted group-based designations with correlations that treated AUC_{3S} as a continuous variable.




Figure 3.6 Averaged 3-session AUC histogram. Blue dashed line indicates the sample median used for high / low AUC split.

3.4 Summary

Selection of participants with a restricted range of (high) performance was successful, as demonstrated by high d' across all levels of the task. More importantly, steep (low AUC) and shallow (high AUC) discounters showed practically identical performance, ensuring that differences in discounting did not stem from individual differences in performance alone. This was true across all measures examined including the discrimination index d' , Ex-Gaussian RT parameters, lure rates, etc. This supports the assumption that participants vary in discounting because of intrinsic cognitive effort-related cost functions, and not because of task-performance differences.

There is evidence that discounting rates were influenced by both state and trait factors, as has been observed previously (Westbrook et al., 2013). Evidence for trait effort discounting is seen in the inter-session correlation of individual differences in AUC. Steep discounters tend to remain steep discounters across all three sessions. State factors influencing discounting include cognitive load and base offer amount (though this effect was not reliable), and also, potentially, the number of N-back rounds participants completed in a session before engaging in decision-making (fewer rounds yields shallower discounting).

Either because of state-dependent variation, measurement error in the COGED procedure, or both, intersession discounting was not perfectly stable within participants. Although steep discounters tended to remain steep discounters, and vice versa, fluctuations in session-specific AUC values meant that a clear group distinction between steep and shallow discounters were not achieved. To emphasize trait experience of subjective effort, given intersession variation in discounting, AUC was averaged across all three sessions (as AUC_{3S}) for subsequent analyses. Also, individual difference analyses consider AUC_{3S} as a continuous variable rather than defining strict group membership.

Chapter 4: Brain Regions Tracking

Cognitive Effort

The core question of this dissertation was what brain regions track subjective cognitive effort, during a demanding working task like the N-back. Brain activity tracking effort should vary by load. However, it should also vary by how participants feel about performing high load tasks. Here, COGED was used to quantify subjective effort costs participants experienced while completing the N-back. Thus, to investigate which brain regions track subjective effort is to ask which brain regions track both load and discounting while participants are engaged with the task.

In this section, whole-brain BOLD data, collected while participants performed increasingly demanding N-back loads were fit by GLMs with convolved boxcar regressors spanning the duration of each block of N-back stimuli. Separate regressors were fitted for each N-back level, such that regression weights (beta parameters) indicate the amplitude of tonic activity while participants were engaged in each N-back load level. The resulting regression weights were then tested at the group level to investigate regions of the brain tracking objective load (varying by N in the N-back), and also discounting measures, collected with the COGED task.

An important concern is whether instrumental, free-choice COGED decision-making can be informative about patterns of brain activity observed during mandatory N-back performance. That is, why should subjective effort costs, as measured by COGED, be reflected in brain data while participants are performing a non-instrumental, mandatory task? The logic of the experimental design is that, given that all participants

were required to perform the N-back to the best of their ability, incentives for performance are arguably consistent across participants, and only subjective effort costs should differentiate participants' experience. Hence, when presented with instrumental COGED decision-making, the subjective costs experienced during the prior N-back practice alone would influence participants' willingness to pursue reward.

4.1 Brain Regions Varying by Load

As discussed in the Introduction, a network of “task-positive” regions, especially including the dACC, the dlPFC, the AI, and the IPS have been implicated in supporting tasks requiring effortful working memory and cognitive control, as reviewed in (Westbrook & Braver, 2015). Activity in these regions typically increases with load, and their recruitment predicts performance. Hence, they are prime targets for tracking cognitive effort. Also, as discussed in the Introduction, decreasing activity in a “default-mode” network of regions, especially including the vmPFC, and the PCC, has been shown for increasing load and decreased activity therein predicts more consistent task performance, as reviewed in (Westbrook & Braver, 2015). Hence default mode regions may also be prime targets.

Though task-positive and task-negative regions may vary monotonically with load, there is also evidence of non-monotonic, inverted-U functions that reflect supra-capacity demands, shifting cognitive strategies at excessive loads, or some combination of factors (Jaeggi et al., 2003; 2007). Given the limited range of N-back loads tested in prior studies (typically $N = 0$ —2 or 3), it is unclear whether inverted-U patterns are specific to certain regions, and whether they persist at higher levels. For example, (Callicott et al., 1999) found evidence of linearly rising activity in a “percingulate” area

roughly corresponding to the preSMA and dACC over $N = 1-3$, relative to $N = 0$ (identifying every instance of a pre-defined target letter, e.g.: “x”), and an inverted-U function in the dlPFC over that same range. Although an inverted-U pattern might indicate disengagement from an overly effortful task, the steadily rising pericingulate activity contradicts this hypothesis. Yet, it is unknown what would have happened at higher load levels. A key advantage of the study design is that very high demand levels ($N = 4-6$) have been included, thus giving the opportunity to investigate a wider range of load-response functions. If all a priori regions of interest showed inverted-U load functions (or U-shaped, in the case of the task-negative regions) across the range such that $N = 6$ activity profiles resembled resting levels of activity, this would support that participants are disengaging from the task.

Another hypothesis is that inverted-U functions reflect adaptive strategy shifts (Jaeggi et al., 2007). The key evidence for this hypothesis, as discussed in the Introduction, is that participants performing the 3-back better showed a bigger drop in activity from the 2-back, in task-positive regions (and thus a sharper inverted-U), than those performing the 2-back worse. That is, a more pronounced inverted-U pattern was associated with better performance, suggesting the inverted-U reflects an adaptive shift in strategies. Evidence for this hypothesis would include 1) a replication of the result that better 3-back performers show a more pronounced inverted-U load function over $N = 1-3$ and 2) that other regions show either flat or monotonic load functions over the full range ($N = 1-6$). The logic of the second piece of evidence is that patterns of activity across $N = 1-3$ that are maintained across $N = 4-6$ would suggest equal (or greater)

engagement at higher loads, and consequently would contradict the hypothesis of disengagement.

4.1.1 Load Modulation by Network

In recent years, it has become well-established that the brain can be segmented into a canonical set of functionally-coupled networks that coactivate within the network and show distinct responses to brain states across networks. By examining patterns of underlying functional connectivity – pairwise voxel time series correlations in fMRI data – Power et al. (2011), among others, have identified an intrinsic architecture of the brain. This architecture can be described by a standard set of nodes that functionally couple to varying degrees with strong intra-network connectivity and weak inter-network connectivity. Task-positive control and working-memory regions include nodes that largely fall within the dorsal attention (DorAtt), fronto-parietal (FP), salience (Sal) and cingulo-opercular (CO) networks, while task-negative the default mode (DMN) regions comprise their own network.

The following analysis capitalizes on the intrinsic functional architecture of the brain, analyzing load functions first at the level of networks, and then at nodes within networks. Nodes are defined from a standard set of 264, with each node represented as a 12 mm sphere belonging to one of 13 different networks (Power et al., 2011). These nodes, and the networks to which they belong, are considered the basic unit of analysis rather than individual voxels. The biggest advantages of this approach are that: a) it provides a much smaller set of statistical comparisons, for which correction is much less exacting than voxel-wise analyses; and b) node selection and network assignment are based on independently defined criteria (namely: intrinsic functional connectivity and

stereotyped network architecture across a wide range of cognitive states). Moreover, the current analyses focuses on a subset of 5 networks hypothesized to track the effort-related load: the DMN, DorAtt, FP, CO and Sal networks, further reducing the search space. Complementary analyses for all networks can be found in the Appendix.

The biggest disadvantage of node-based analyses is that the specific extents and loci may not coincide with regions defined across the whole brain by the specific design of this study, and as such either dilute or miss the strongest local effects. A related disadvantage of the network-based analysis is that network membership is defined by resting-state time series correlations (Power et al., 2011) rather than mean, task-related activity and hence such independently-defined networks may contain nodes with differing (block-level activity) load functions. Moreover, networks were not defined with regard to the N-back task, in particular. Hence, it is possible that nodes within a network could show differing load functions that would cancel each other out when averaged across nodes. To address these potential disadvantages, complementary voxel-wise, whole-brain analyses were conducted and reported in Appendix C. These complementary analyses essentially recapitulate the major findings described in the next section, supporting the utility of the node-based approach.

4.1.1.1 Linear and Non-Linear Load Functions Among Hypothesized Networks

Among a priori networks of interest, only one appears to vary approximately monotonically with load, mirroring monotonic discounting functions. As shown in Figure 4.1, the DMN appears to decrease monotonically with load. The function is approximately monotonic in that only $N = 1$ shows an obviously smaller response than other levels. Nevertheless, a repeated measures ANOVA confirms variation by load over $N = 2$ — 6 ($p < 0.01$) and there are pairwise, trending differences with $N = 5$ and $N = 2$ (p

= 0.08) and $N = 5$ and $N = 3$ ($p = 0.05$). There is no reliable difference between $N = 2$ — 3 and $N = 4$ — 6 , however ($p = 0.32$). It is clear, in any case, that the pattern of activity does not support disengagement at high loads as the extent of deactivation at $N = 6$, for example, is equal to, or greater than that at any of the lower loads. The DMN has been observed to vary as a decreasing function of load in prior reports (Pyka et al., 2009), cf. (Jansma, Ramsey, de Zwart, Van Gelderen, & Duyn, 2007; McKiernan et al., 2003), but never at the very high demand levels explored in this study.

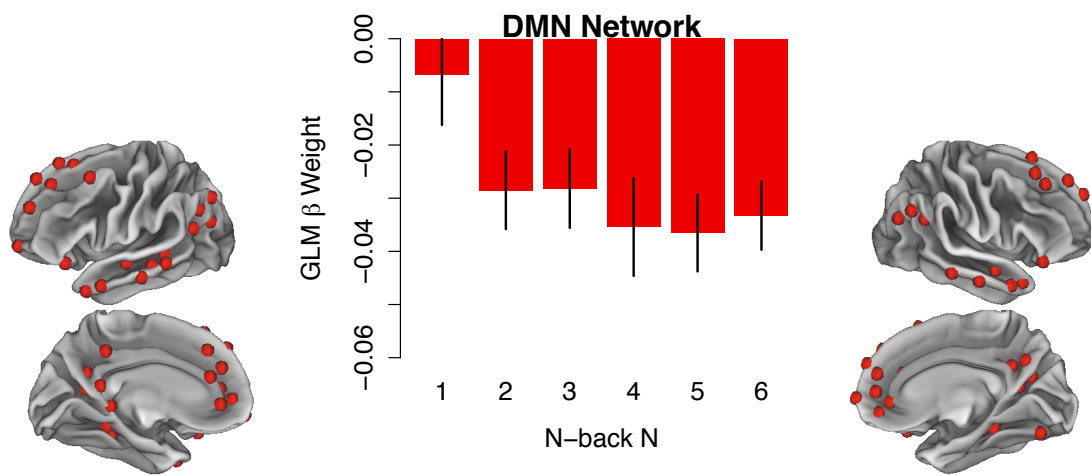


Figure 4.1 GLM β weights for each N-back level averaged across nodes in the DMN. Lines indicate SEM. Nodes are 12 mm in diameter.

A more common load function, however, is quadratic. As shown by others investigating load effects of the N-back, regions that correspond with the task-positive networks, in particular, show an inverted-U shaped pattern, on average, increasing up to approximately $N = 3$ and then decreasing after. This is most notable in the FP, DorAtt, and Sal networks, shown in Figure 4.2.




Figure 4.2 Mean GLM β weights for each level of the N-back task averaged across all nodes in the A) FP, B) DorAtt, and C) Sal Networks and a map of nodes included in the respective analyses. Lines indicate SEM.

Finally, the CO network is essentially flat across load levels, with no real change across loads with the exception of a dip at $N = 5$ (mirroring what is found in other regions). Unlike the monotonic or inverted-U patterns, the CO network appears to encode simply that the participant is engaged in the task, rather than the objective load.




Figure 4.3 Mean GLM β weights for each level of the N-back task averaged across all nodes in the CO network, and a map of nodes included in the analysis. Lines indicate SEM.

One unexpected effect is that activity deviates from the inverted-U or flat load patterns for $N = 6$, increasing from $N = 5$ for some networks. This may result from some interesting, as yet unknown shift in the way participants handle the 6-back (or 5-back), in particular, or it may result from something less interesting, like order effects. As described in the Methods, participants completed all 6 levels of the N-back in order of increasing demand. This method was adopted so that all participants would experience all levels of the N-back identically (the 3-back after the 2-back, etc.), enabling better controlled comparisons of individual differences in discounting. As discussed in the Methods section, the downside of a fixed order, of course, is that load is confounded with order. For example, after a long run of $N = 1$ — 5 , participants may essentially “perk up” when they know they are completing the final level ($N = 6$), due to revival of motivation.

Alternatively, participants may engage less than would otherwise be expected on the 5-back, due to fatigue factors, and the knowledge that an even more demanding load level (i.e., 6-back) is left to perform. Unfortunately, the fixed task order precludes resolving these questions.

A formal test of networks showing linear and non-linear effects of load is accomplished by multi-level models including both first- and second-order predictors of Load, and for which all predictors were further allowed to vary by participant j . Multi-level models are useful because they can accommodate effects of multiple load levels nested within participants, by properly assigning between- and within-individual variance. This is also more powerful in that it allows modeling the common and unique effects of linear and quadratic trends across participants simultaneously.

$$\beta_i = B1_{j[i]} + B2_{j[i]} \text{Load} + B3_{j[i]} \text{Load}^2 + \varepsilon_i \quad (4.1)$$

$$B1_j = \gamma_{1,0j} + \eta1_j \quad (4.2)$$

$$B2_j = \gamma_{2,0j} + \eta2_j \quad (4.3)$$

$$B3_j = \gamma_{3,0j} + \eta3_j \quad (4.4)$$

Given the unexpected bump in recruitment for $N = 6$, quadratic models may fit better for $N = 1—5$ better than across the entire range of Loads. Thus, for exploratory purposes, the same set of analyses was conducted for both $N = 1—5$ and $1—6$. The results of these model fits, describing fixed linear and quadratic effects, in particular, are given in Table 4.1. Note that a full table of all networks is provided in Appendix C.

Network	Linear effects		Quadratic effects	
	B2*10 ⁻³	p-value	B3*10 ⁻³	p-value
Across Loads N = 1—5				
CO	-4.1	0.05	-3.0	0.05
DMN	-6.7	0.01	2.4	0.23
FP	1.7	0.45	-6.6	<0.01
Sal	-1.7	0.40	-4.2	0.02
DorAtt	-2.2	0.34	-4.5	0.01
Across Loads N = 1—6				
CO	-2.7	0.11	-0.2	0.82
DMN	-4.7	0.01	2.1	0.10
FP	0.9	0.62	-3.0	<0.01
Sal	-1.3	0.42	-1.3	0.23
DorAtt	-2.2	0.09	-1.7	0.27

Table 4.1 Linear and quadratic fixed effects of load in networks of interest for N = 1—5 and N = 1—6. Shading: light grey for p < 0.10, medium for p < 0.05, and dark for p < 0.01.

Consistent with an apparent monotonic function in the DMN, a reliable negative linear effect of load was observed, regardless of the range considered. Tentative negative linear effects were observed in the DorAtt and CO networks, depending on the load range. Negative linear effects were not anticipated for any region except the DMN, as a function of load alone. These negative linear effects may reflect accommodation, fatigue, or any other number of confounds resulting from a fixed task order.

Riding on top of linear effects are quadratic effects, particularly in the task-positive, working memory and control-related networks, like the FP network, and for N = 1—5, the Sal, and DorAtt networks. The fact that most quadratic effects are negative is consistent with previously observed inverted-U functions of load observed on the N-back across N = 1—3 (Callicott et al., 1999; Jaeggi et al., 2003; 2007). For the Sal and DorAtt networks, quadratic effects are attenuated with the inclusion of N = 6. Also, across N = 1—6, the CO network is a flat function, as described above, with neither linear nor quadratic effects of load.

Prior studies have not investigated patterns of recruitment of these regions beyond $N = 3$, and so the current study provides novel information about what happens in these regions at extremely high load levels. Notably, for example, while dACC (approximated by nodes within the Sal and CO networks) was shown to increase monotonically with load over lower loads, when examining beyond $N = 3$, this region was found to show a clear inverted-U profile (i.e., for $N = 1$ — 6). The dACC region of interest is best characterized through Node #213, which matches closely in anatomic location to prior studies examining dACC effects in N-back, cognitive control, and effort-based decision-making tasks (MNI coordinates -1, 15, 44; see Figure 4.4). Consequently, this node will be the primary target of subsequent analyses focused on the dACC.

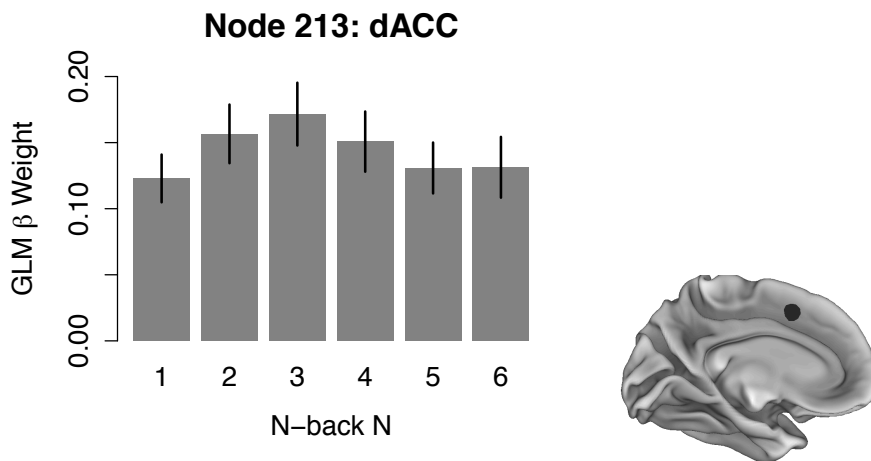


Figure 4.4 Mean GLM β weights for each level of the N-back in the dACC and the location of the 12 mm spherical node on the medial wall of the left hemisphere.

As mentioned above, an early N-back study identified three load functions at lower load levels: $N = 1$ — 3 , (Callicott et al., 1999), which were characterized as an inverted-U pattern in a dlPFC cluster, a flat pattern in a precuneus cluster, and a monotonic pattern in a medial pericingulate cluster which rose monotonically up to $N =$

3. The dACC node is a clear example where, if only examined across $N = 1-3$, it would have appeared to be monotonic, but across $N = 1-6$ proves to be inverted-U.

An examination of other nodes approximating loci of interest in the early study demonstrates that other regions also require updating when examined across the full load range. A FP node (#201; $x = -42, y = 25, z = 30$; 14 mm from the dlPFC cluster center reported in Callicott et al.) also shows an inverted-U profile, just as the dlPFC did in the earlier study (Figure 4.5).

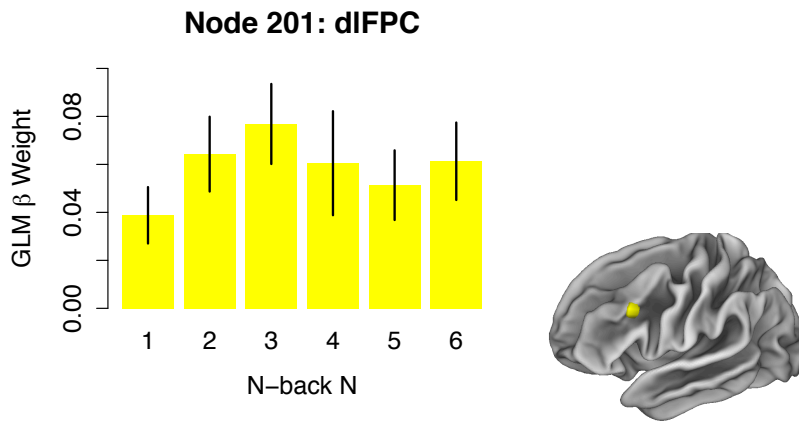


Figure 4.5 Mean GLM β weights for each level of the N-back in the dlPFC and the location of the 12 mm node on the left lateral hemisphere surface.

The decline, however, only starts at $N = 4$ in the present dataset. By contrast, a DorAtt node approximating the precuneus cluster showing a flat load function (#251; $x = 10, y = -62, z = 61$; 9 mm from the precuneus cluster reported in Callicott et al.) also shows an inverted-U profile rather than a flat load function.

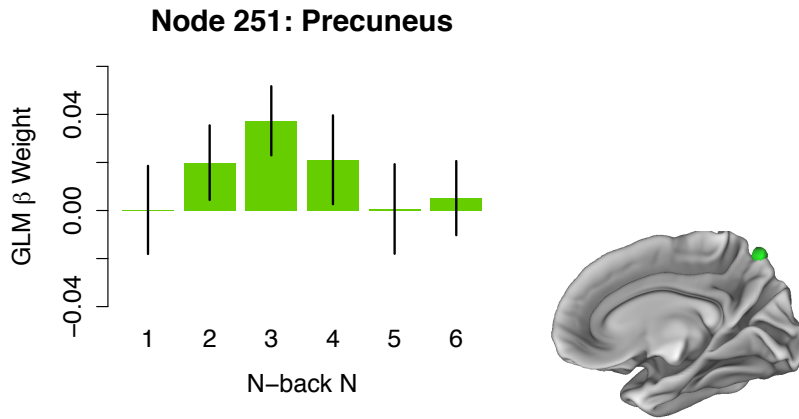


Figure 4.6 Mean GLM β weights for each level of the N-back in the precuneus.

The new result may be different because of slight distinctions in location; but it could also have to do with the fact that a larger range of load levels were examined in the present study. Because the prior study only examined up to $N = 3$, what looked like a flat function in the previous study actually resulted from a range restriction on load. Although a flat load function was not observed in the DorAtt node, a flat function was observed in other locations. Notably, a CO node approximating the earlier monotonic cluster, medial pericingulate cluster (#53; $x = 13, y = -1, z = 70$; 13 mm away from the region described in Callicott et al.) showed a flat load function.

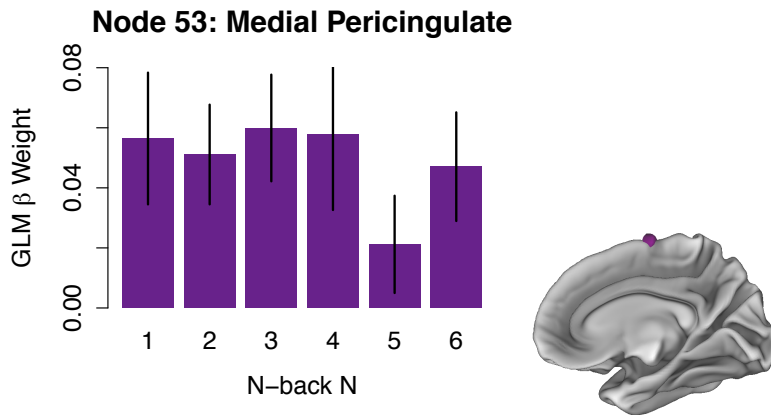


Figure 4.7 Mean GLM β weights for each level of the N-back in the medial pericingulate.

That is, rather than showing an unconstrained rise across all load levels ($N = 1\text{---}6$), this node showed essentially flat behavior, like the rest of the CO network. Again, what looked like monotonic behavior in this region in the early study may have resulted from the smaller range of loads examined in that study.

It is important to point out that differences in the precise timing or trial sequential structure of the N-back stimuli presented in the earlier study, or even differences in participant samples could have also led to differences between the results of Callicott et al. (1999) and the load functions observed here. This study was not designed to test for a precise replication of the load functions observed in the earlier study. Instead, the key point was to investigate Callicott et al.'s descriptive load patterns across a wider range of loads in independently defined networks. Thus the key results here are that a priori task-positive networks of interest all show what was originally characterized as inverted-U patterns with the exception of the CO network, which showed a flat pattern, while the DMN is the only network with an arguably monotonic pattern.

As described above, one hypothesis about inverted-U load functions is that they reflect adaptive strategy shifting. In a key study, participants showing the sharpest inverted-U pattern – decreasing most at $N = 3$ – showed the best performance. Hence adaptive strategy shifts are reflected both in better performance and a more pronounced inverted-U across $N = 1\text{---}3$ (Jaeggi et al., 2007). Using the same criterion employed to divide participants by performance (3-back d'), the current dataset replicates the results of Jaeggi et al. (2007). Namely, across $N = 1\text{---}3$, individuals above a median split in 3-back performance (“high performers”) had a drop in activity in FP and DorAtt networks for the 3-back relative to the 2-back, while those below the performance split (“low

performers”) showed yet higher activation for the 3-back. That is, high performers showed a more pronounced inverted-U pattern across $N = 1—3$.




Figure 4.8 Load profiles for high (above median 3-back d' scores) and low performers in FP and DorAtt networks.

Given that both the FP and DorAtt networks have inverted-U load functions at a wider load range, both high and low performers show a peak and decline. However, because the high performers have a more pronounced inverted-U over the lower range $N = 1—3$, the low performers, by comparison, peak at a higher load. That is, while high performers show their peak earlier at $N = 2$ in these regions, low performers peak later, at $N = 3$.

The current design, which varied load from $N = 1$ to $N = 6$, revealed a pattern of extended load functions across brain networks. Given that subjective effort increased monotonically with load (i.e., in terms of the COGED discounting pattern; see Figure 3.6) a network like the DMN, which has monotonically decreasing activity with load, may straightforwardly track effort as hypothesized. Monotonic load functions in the DMN and flat functions in the CO network support that participants remain engaged at very highly demanding levels ($N = 4—6$). Inverted-U functions, on the other hand are not

immediately reconcilable with monotonically increasing effort. It is important to note that even if inverted-U shapes reflect shifting strategies, participants may still find the strategies to which they shift more subjectively costly, if necessary, to deal with very high demands. The problem is that load functions alone do not strictly indicate whether a network tracks effort. It is further necessary to ask whether networks vary by an operational measure of subjective costs – effort discounting.

4.2 Do Brain Networks Vary by Load and Discounting?

A key question is whether networks show patterns of activity that vary as a function of subjective sensitivity to cognitive effort. That is, do individuals who vary in terms of subjective effort show differences in activity across loads? Do their load functions vary systematically? And also, do these effects vary by network?

To answer these questions, a 3-way Network x Load x AUC_{3S} repeated measures ANOVA of N-back β weights (averaged across all nodes in each network), treating Load and Network as within-participant variables, and AUC_{3S} as a between-participant variable, was computed to test for the influence of critical variables on the hypothesized networks. Note that Load is treated as a factor with multiple levels rather than a linear predictor (given that multiple networks are already known to show non-linear load functions). As described previously, AUC_{3S} is the three session-averaged discounting value, capturing participants' mean sensitivity to effort costs across all N-back levels. The results of that test (Table 4.2) demonstrate that there is a reliable effect of AUC_{3S} , and that, mirroring the previous analysis, mean activity and load functions both vary by network. There does not appear to be a reliable difference in the effect of AUC_{3S} across networks, implying that the effect is instead consistent across networks.

Effect	DF	F	p	
AUC _{3S}	1, 22	5.71	0.02	*
Network	4, 88	71.8	<0.01	**
Load	5, 110	1.78	0.12	
AUC _{3S} x Network	4, 88	1.57	0.15	
AUC _{3S} x Load	5, 110	1.35	0.25	
Network x Load	20, 440	5.31	<0.01	**
AUC _{3S} x Network x Load	20, 440	0.62	0.89	

Table 4.2 Repeated Measures ANOVA in N-back regression weights in the DMN, FP, CO, Sal, and DorAtt networks. (**: $p < 0.01$, *: $p < 0.05$, .: $p < 0.10$)

4.2.1 Specific Networks Predict Discounting

The previous ANOVA supports investigating the relationship between discounting and brain activity further, and also investigating load functions in each network separately. To explore activity-discounting relationships, a multi-level model was fit to N-back activity levels (reflected in β weights), nested within participants to predict SV_{3S} values: one for each participant, at each load level. A multi-level approach permits modeling relationships across all participants and loads simultaneously, assigning variances appropriately, without resorting to aggregate discounting measures (AUC_{3S}) that sacrifice potentially useful (between-load) information.

Note that to model a fixed effect of β weights, while properly accounting for the fact that different participants “experience” different β values, β weights are centered within participants (β_{ctr}) and treated as a predictor at the load level of the model. This predictor describes how SV varies with respect to (load-specific) changes in participants’ brain activity (cf. orange lines in Figure 1.2). In addition, another predictor (β_{avg}) describes the load-independent effects of individual differences in participants’ mean regression weights (cf. blue lines in Figure 1.2). In this model, only the intercept is allowed to vary randomly across participants, and β_{avg} is a predictor of this intercept. Note that the load-independent predictor β_{avg} describes how a subject’s SV intercept

varies, and thus describes a main effect of individual differences in cross-load averaged brain activity on mean discounting. β_{ctr} , on the other hand, accounts for the way that SV is predicted to change as a function of changes in brain activity with load, and thus is akin to an interaction of load and brain activity as a predictor of SV, independent of objective load.

$$SV_{3Si} = B1_{j[i]} + B2_{j[i]} \text{Load} + B3_i \beta_{ctr} + \varepsilon_i \quad (4.5)$$

$$B1_j = \gamma 1_{0j} + \gamma 1_{1j} \beta_{avg,j} + \eta 1_j \quad (4.6)$$

The last two columns of Table 4.3 provide the estimated load-specific β_{ctr} and load-independent β_{avg} fixed effects modeled separately for each a priori network of interest.

Network	Activity effects *10 ⁻¹ (p-value)	
	Load-specific B3	Load-independent $\gamma 1_1$
FP	-0.6 (0.93)	-16.5 (<0.01)
Sal	21.1 (0.70)	-16.1 (<0.01)
CO	3.6 (0.63)	-15.5 (<0.01)
DMN	0.7 (0.92)	-8.3 (0.21)
DorAtt	1.6 (0.81)	-24.5 (<0.01)

Table 4.3 Effects of load-dependent and load-independent activation on discounting in selected networks. Shading in the table indicates significance level with light grey for $p < 0.10$, medium for $p < 0.05$, and dark for $p < 0.01$.

In none of the networks does the load-specific β_{ctr} predict activity when controlling for the objective load. Load-independent β_{avg} , on the other hand, predicts SV_{3S} in all networks except the DMN. The fact that the load-specific predictor does not explain variance in SV_{3S} indicates that there is insufficient information about SV in the way that activity changes over task load levels, apart from objective load (N) itself.

Importantly, for task-positive networks, activity is reliably smaller with increasing AUC_{3S} indicating that these networks not only vary with load, but they also vary such that steeper discounters recruit most networks more robustly, independently of task level.

In other words, individuals that discount cognitive effort more steeply (i.e., for whom the task is subjectively costlier) show overall greater activity in the N-back than individuals who are shallow discounters. These AUC_{3S} effects are found in all of the task-positive (CO, FP, Sal, and DorAtt) networks, but not in the task-negative DMN. Figure 4.9 shows an example of the discounting and load effects in the DorAtt network. As is readily apparent, lower load-averaged regression weights predict higher AUC_{3S} values while regression weights are higher at all load levels for a group with below-median AUC_{3S} values than it is for those with above-median AUC_{3S} values.




Figure 4.9 AUC_{3S} as a function of network-level individual differences in activity in the DorAtt network and activity as a function of load and discounting. A) AUC_{3S} as a function of activity averaged across task load levels, with task loads modeled separately. Solid lines give linear regressions, and dashed lines give the 95% CI. B) Activity across loads, separated by whether participants have above (Hi AUC) or below (Lo AUC) median AUC_{3S} values. Vertical bars give the SEM.

Aside from the main effect of discounting on activity in the DorAtt network, there also appears to be a difference in the shape of load functions for steep and shallow effort discounters as shown in Figure 4.9B. Namely, while both steep and shallow discounters (those below and above the median AUC_{3S} value, respectively) show an inverted-U profile, the profile also appears shifted, such that steep discounters peak at $N = 2$, while shallow discounters peak at $N = 3$. Further evidence is that the load function is inverted-U for the steep discounters over $N = 1$ – 3 , while it is not for the shallow discounters. A

multi-level, variable-intercept model (varying by participants in each of two groups) reveals that across $N = 1-3$, there is a reliable positive linear effect ($p = 0.01$) and negative quadratic effect ($p = 0.02$) of load for the steep discounters, but not the shallow discounters ($p = 0.95$ and $p = 0.89$, respectively). Note that this formal test complements the visual observations of an inverted-U pattern for steep discounters across $N = 1-3$, and no inverted-U pattern for shallow discounters across this range, although it does not demonstrate that one group shows a reliably stronger inverted-U than another.

Lower mean activity and later peaking (at $N = 3$ rather than $N = 2$) among shallow discounters support a “neural efficiency” hypothesis (Jaeggi et al., 2007) – that some individuals operate more efficiently and do more with less. This corresponds to the axis of Figure 1.3 from lower-left to the upper-right quadrant. By this interpretation, peaking at higher loads among shallow discounters reflects greater efficiency, allowing for still higher activity at $N = 3$ relative to $N = 2$ whereas reduced efficiency among steep discounters requires them to bring certain resources fully to bear at $N = 2$. While these results are consistent with a neural efficiency hypothesis, the conclusion is tentative, being based on comparison of two load levels only. Lower load-independent activity among shallow discounters, on the other hand, is a robust effect at all loads and strongly contradicts the alternative interpretation (on the upper-left to lower-right axis of Figure 1.3) that individual differences in mean activity in task-positive are related to volitional will to recruit task-positive networks. If that were true, load-independent activity in task-positive regions would have been *higher* not *lower* for those finding the N-back less costly.

Interestingly, the shifted load function for shallow versus steep discounters mirrors the same pattern described above for low versus high 3-back performers, and thus is consistent with the hypothesis (Jaeggi et al., 2007) that subjective effort is a cue to shift strategies adaptively. Just like those showing worse 3-back performance, shallow discounters also show lower activity at lower loads, and peak later (at $N = 3$). Also, just like those showing better 3-back performance, steep discounters show higher activity at lower loads, and peak earlier (at $N = 2$). In other words, higher activity at low loads ($N < 3$) and earlier peaking are associated with both greater subjective effort and adaptive strategy shifting (that preserves performance for $N = 3$). This coincidence supports a hypothesis, proposed by Jaeggi et al. (2007), that subjective effort is used as a cue that adaptive strategy shifting is needed (e.g. to maintain performance). As noted, performance is unrelated to discounting in our sample (AUC_{3S} is neither a linear predictor of performance at $N = 3$, $p = 0.64$ nor $N = 2$, $p = 0.57$), likely because of the non-random way in which participants were selected for this study (restricting range on performance while maximizing range on subjective effort). Indeed, I have shown previously that SV and performance are related when participants are selected pseudo-randomly, via self-selection (Westbrook et al., 2013). Thus the current dataset cannot be used to test directly whether subjective effort drives adaptive strategy shifting. Moreover, an interpretation that two patterns are qualitatively similar because of a shared mechanism of subjective effort driving strategy shifting is admittedly post-hoc. Further study is needed to whether performance and AUC are related at specific load levels (like the 3-back) due to adaptive strategy shifting.

Complementary analyses in which the relationships between discounting and BOLD signal in all 13 networks defined by Power et al. (2011) are provided in Appendix C. For example, one complementary approach examines how AUC_{3S} predicts brain activity in multi-level models in which dummy-coded load predicts brain activity, along with AUC_{3S} as an intercept predictor. The key result confirms a main effect between AUC_{3S} and load-independent BOLD signal in multiple networks (Table C.2, C.3 and Figure C.4). Furthermore, when each load-specific dummy code is further allowed to vary by AUC_{3S} , AUC_{3S} is never a reliable predictor of the load effect for any load or for any network. This result confirms that discounting does not interact with load to predict brain activity, consistent with the observation that load-specific changes in brain activity do not predict changes in SV_{3S} , independently of objective load.

While a reliable load-independent effect of activity on subjective effort is both intriguing and confirms a prior prediction, the results do not support the prediction that load interacts with activity to predict discounting. That is, although subjective effort varies by objective load, that variation does not further vary by load-specific changes in network activity. It was not the case, for example, that task-positive networks showing increasing activity across $N = 1-3$, increased *more steeply* for those who find the N -back costlier (refer to the orange lines, Figure 1.2). Despite some suggestive results (e.g., earlier peaking, Figure 4.9B), in no regions was the shape of the load-related activity pattern reliably changed as a function of AUC. Put differently, after controlling for the objective load level, load-specific brain activity did not explain additional variance in discounting. It is important to consider that because this was an unpredicted null result, interpretations are necessarily post-hoc and caution is warranted. In particular, the null

result may have to do with the limited sample size and thus sensitivity to detect load-specific effects apart from objective load. It may also relate to the fact that block regression weights mask underlying dynamics that would differentially predict discounting at each load level. Future focus on smaller-scale features of load-specific BOLD signal, like trial-wise, event-related (rather than block-wise) N-back activity or changes in network properties as a function of load and their relationship to subjective effort, for example. These possibilities will be elaborated in the General Discussion.

4.2.2 Specific Nodes Within Networks Predicting Discounting

Reliable individual difference effects of discounting were reflected in the activity level of task-positive networks, when characterized at the whole network level. It is also likely that specific nodes within these networks show particularly strong activity-discounting relationships. Moreover, as discussed above, a prior study has identified specific regions of interest in the dlPFC and dACC that may play particularly important roles in encoding subjective effort during engagement with a challenging cognitive task (McGuire & Botvinick, 2010). Likewise, other studies have implicated the ventral striatum (VS) for encoding effort costs when cued with a reward earned for exerting prior cognitive effort (Botvinick et al., 2009), or encoding motivation during (Schmidt et al., 2012) and prior to (Schouppe, Demanet, Boehler, Ridderinkhof, & Notebaert, 2014; Vassena et al., 2014) exertion of a cognitively effortful task. Hence, certain nodes within networks might relate to discounting measures particularly strongly, and identifying those nodes would afford greater specificity about which brain regions track cognitive effort.

The following analysis considers whether activity in specific nodes encodes the subjective cost of effort on the N-back among two sets of nodes: first, among a target set of a priori nodes identified in prior literature for tracking cognitive motivation or subjective cognitive effort or, and second, an exploratory set, including all nodes in the a priori networks of interest. A priori nodes include two FP nodes, which are 8 mm and 9 mm from peak voxels identified by (McGuire & Botvinick, 2010) for encoding subjective effort in the lateral PFC, and a medial PFC node, part of the Sal network, which is 5 mm from a key loci of interest in the dACC. In addition to these, two VS nodes are also included. The nodes were not part of the original 264 node set, since the Power et al. (2011) node set does not adequately target subcortical structures.

As in the network-level models, fitted models describe, for each node, the estimated load-specific β_{ctr} and load-independent β_{avg} effects on SV. Table 4.4 gives both effects for the five a priori nodes of interest.

Network	Node	Description	MNI (LPI)			Activity effects * 10^{-1} (p-value)	
			x	y	z	Load-specific B3	Load-independent $\gamma 1_1$
FP	176	l MFG	-47	11	23	0.7 (0.87)	-6.6 (<0.01)
	186	r IFG	47	10	33	2.3 (0.66)	-7.8 (0.02)
Sal	213	l dACC	-2	15	43	-0.2 (0.97)	-7.0 (<0.01)
Other	265	r VS	-12	12	-6	0.1 (0.98)	-18.4 (<0.01)
	266	l VS	12	10	-6	2.6 (0.45)	-13.8 (<0.01)

Table 4.4 Effects of load-dependent and load-independent activation on discounting in a priori nodes of interest taken from McGuire et al., 2010 and Botvinick et al. 2009. Significance level indicated light $p < 0.10$, medium $p < 0.05$, and dark shading for $p < 0.01$.

In none of the nodes does the load-specific β_{ctr} predict activity when controlling for the objective load, mirroring the network-level results. Again, this seems to indicate that there is insufficient information about SV in the way that activity changes over task load levels, apart from objective load (N) itself. Load-independent β_{avg} , on the other hand,

predicts SV in all a priori nodes – an effect that also mirrors the priori network-level analyses.

Among the a priori nodes, all show reliable negative effects of load-independent activity on SV_{3S}. Of particular interest is that both dlPFC and dACC loci identified for encoding subjective, self-reported effort in a priori study (McGuire et al., 2010) also predict subjective effort in this dataset. In that study, greater activity during performance of a cognitively demanding task was related to increased within- and between-subject self-reported desire to avoid that task, presumably because participants found the tasks to be costlier. Here, we find a concordant between-subjects relationship in all three regions (Node #s, 176, 186, and 213) whereby higher, load-independent activity predicts greater subjective effort costs. Figure 4.10 highlights the pattern observed in the dACC region, indicating a similar profile to what was observed in the DorAtt network, i.e., greater overall activity in steeper discounters and a subtle change in the load function (peaking at N=2 among steep discounters and N=3 among shallow discounters).

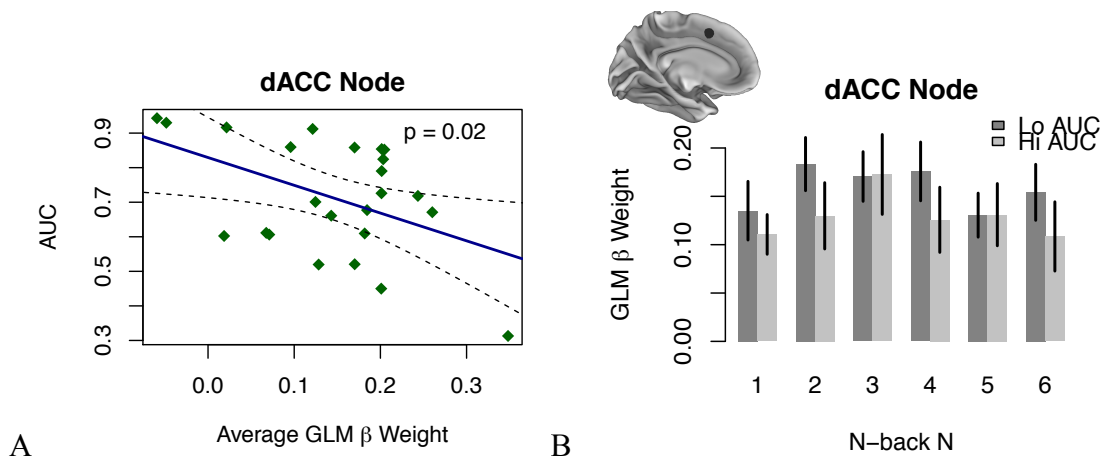


Figure 4.10 Individual differences in activity in an a priori dACC node as a function of AUC_{3S} and load. A) AUC_{3S} as a function of activity averaged across task load levels with task loads modeled separately. Solid lines give linear regressions, and dashed lines give the 95% CI. B) Activity across loads, separated by whether participants have above (Hi AUC) or below (Lo AUC) median AUC_{3S} values. Vertical bars give the SEM. Location of the dACC node is also plotted.

In addition to the lateral and medial PFC, the bilateral VS nodes also explained individual differences in discounting. This result is consistent with other studies showing the encoding of effort and, conversely, motivation in the VS (Botvinick et al., 2009; Schmidt et al., 2012; Schouppe et al., 2014; Vassena et al., 2014). The direction of the effect in these nodes – higher activity among those who found the task more effortful (Figure 4.11) – was not necessarily anticipated.




Figure 4.11 Individual differences in average activity in a priori VS nodes predicts AUC_{3S} . The average regression weight predicts AUC_{3S} in both the right ($p < 0.01$) and left ($p = 0.02$) VS, negatively.

In the prior studies, the common finding was that during effort anticipation, higher activity was observed when individuals anticipated lower effort. Yet, one study

showed that this may depend on whether a participant is engaged in forced-choice or free decision-making about task engagement (Schouppe et al., 2014). Moreover, the comparison to decision-making studies is complicated by the fact that the present results were obtained during task performance itself (rather than during decisions about task performance). Two studies have examined objective load encoding in the VS during task performance. One of these studies did not find reliable load encoding (Schmidt et al., 2012), while the other did find load encoding (Satterthwaite et al., 2012), yielding an indeterminate conclusion. A more obvious interpretation of the sign observed here is that, rather than reflecting motivation, greater VS activity among those finding the N-back to be costlier reflects more vigorous task-coordination processes in the striatum that have also been linked to cognitive control, such as working memory gating in cortico-striatal loops (Frank, Loughry, & O'Reilly, 2001).

Across all nodes in the DMN, FP, CO, Sal, and DorAtt networks, a larger exploratory set show significant load-independent relationships between β_{avg} and SV_{3S} . Although this analysis regards nodes within a priori networks that moreover show, in the case of task-positive networks, reliable network-level relationships, it is ultimately exploratory in nature (among 126 nodes). A False Discovery Rate procedure (Benjamini & Hochberg, 1995) is used to limit family-wise error. 49 Nodes surviving a corrected $p < 0.05$ threshold are mapped in Figure 4.12. Surviving nodes cluster along the left dlPFC, and also the bilateral IPS, dACC, pre-SMA, anterior insula, and vIPFC. In addition to these task-positive regions, a specific subset of DMN nodes show reliable relationships to SV, particularly in the bilateral vmPFC and right temporal pole.




Figure 4.12 Nodes in which load-independent N-back activity predicts subjective effort costs. All nodes survive FDR corrected $p < 0.05$. Colors indicate network membership and sign: negative effects are red for DMN, black for Sal, green for DorAtt, and purple for CO; positive effects are dark red for DMN and dark yellow for FP.

Surviving nodes all show greater activity for those with greater subjective effort costs, with four exceptions. As indicated by darker colors, three surviving DMN nodes (dark red) on the right temporal pole and right supramarginal gyrus, as well as one FP node (dark yellow) on the left inferior parietal lobule (IPL) show the reverse effect. In these four nodes, load-independent activity is *lower* for steeper discounters. Closer inspection of the three exceptional DMN nodes, however, reveals that they do not have reliably non-zero regression weights, unlike the wider DMN, and are thus difficult to interpret. Namely, while the network average load function is robustly deactivated, and approximately monotonically decreasing across loads, these three DMN nodes do not show reliable deactivation at any load level. The only exceptions are that the low AUC group alone shows deactivations in the right supramarginal gyrus node at $N = 5$ ($\beta_{avg} = -0.05$, $p < 0.01$) and $N = 6$ ($\beta_{avg} = -0.04$, $p = 0.01$). Given that these three DMN nodes do

not reliably deactivate or track load like the rest of the wider network, it is unclear how to interpret load-independent variation by AUC_{3S} .

By contrast, the exceptional FP node showed reliably positive, non-monotonic behavior, just like the wider FP network. However, unlike the wider FP network for which activity is higher among those for whom the N-back is costlier, in this single node, the pattern is reversed and activity is lower among those for whom the N-back is costlier (Figure 4.13). Given that this effect obtains in only a single node, it is entirely possible that the effect is simply a matter of chance. However, if this node showed a consistently opposite result in a replication sample, it would support that the particular left inferior parietal lobule node ($x = -42, y = -55, z = 45$) tracks individual differences in subjective effort during the N-back in a unique way. For example, higher load-independent activity in this region could reflect greater volition to perform the N-back (rather than greater subjective costliness).




Figure 4.13 Individual differences in activity in a left IPL node as a function of AUC_{3S} and load. A) AUC_{3S} as a function of activity averaged across task load levels with task loads modeled separately. Solid lines give linear regressions, and dashed lines give the 95% CI. B) Activity across loads, separated by whether participants have above (Hi AUC) or below (Lo AUC) median AUC_{3S} values. Vertical bars give the SEM.

4.2.2.1 Relationships Between Node-Level Activity and Discounting, Controlling for Performance

Activity-discounting relationships support that a node’s activity level relates to individual differences in subjective cognitive effort. However, as other authors have noted (Kool et al., 2010; McGuire & Botvinick, 2010), demanding tasks are subjectively aversive (and thus may engender discounting) for a number of reasons apart from effort per se. In particular, performance of demanding tasks is typically associated with slower responding and higher error rates. Thus, it is possible that estimates of subjective costliness may implicitly factor in these performance variables, and thus reflect a form of both delay discounting (longer time for trial completions) and probability discounting (lower likelihood of successful performance). Hence is it important to demonstrate directly that variance among individual differences in brain response to demanding tasks explained by COGED cannot be explained by individual differences in task performance metrics.

To examine this question, the same multi-level modeling approach was used, with the addition of two performance predictors: d' for quantifying accuracy and μRT for response time.

$$SV_{3Si} = B1_{j[i]} + B2_{j[i]} \text{Load} + B3_i \beta_{ctr} + B3_i \mu RT + B4_i d' + \varepsilon_i \quad (4.7)$$

$$B1_j = \gamma 1_{0j} + \gamma 1_{1j} \beta_{avg,j} + \eta 1_j \quad (4.8)$$

Network	Node	Description	MNI (LPI)			Activity effects * 10 ⁻¹ (p-value)	
			x	y	z	Load-specific B3	Load-independent $\gamma 1_1$
FP	176	l MFG	-47	11	23	1.1 (0.70)	-6.7 (0.12)
	186	r IFG	47	10	33	2.2 (0.54)	-6.8 (0.28)
Sal	213	l dACC	-2	15	43	0.0 (0.99)	-6.6 (0.04)
Other	265	r VS	-12	12	-6	0.1 (0.98)	-18.8 (0.02)
	266	l VS	12	10	-6	2.6 (0.30)	-13.1 (0.08)

Table 4.5 Performance-independent relationship between load-dependent and load-independent activity and discounting in select nodes. Five a priori nodes of interest (c.f. McGuire et al., 2010 and Botvinick et al.

2009) are reported. Shading in the table indicates significance level with light grey for $p < 0.10$, medium for $p < 0.05$, and dark for $p < 0.01$.

Although reliabilities (p-values) are attenuated, the load-independent predictor β_{avg} remained reliable (at $p < 0.05$) in the right VS node, and also the dACC after the inclusion of performance variables. That is, for most of these nodes, performance does not explain away the individual differences relationships between mean cross-level activity and subsequent discounting. Notably, β_{avg} in the dlPFC nodes are no longer reliable predictors. These results stand in contrast with those of McGuire et al. (2010) who found the exact opposite: the dlPFC, but not dACC remained reliable predictors of subjective experience after controlling for performance. While a role for the dlPFC in encoding subjective effort was predicted from a literature linking cognitive control processes to cognitive effort, and also, in turn, to dlPFC representation (Botvinick & Braver, 2015; Kool et al., 2013; Westbrook & Braver, 2015), the lack of evidence implicating the dACC was surprising. In contrast, our results support that the dACC, but not the dlPFC best encode subjective cognitive effort, controlling for performance.

Our result confirms the strong prior hypotheses implicating the dACC in tracking cognitive effort for the purposes of value-based regulation of cognitive control (Shenhav et al., 2013). Our results are also consistent with a well-supported hypothesis about the respective roles played by the dlPFC and dACC. According to the hypothesis, the dACC monitors performance and recruits dlPFC to support more intensive cognitive control when performance is in decline (Botvinick et al., 2001). Thus, shared variance (explained away when controlling for performance) between performance measures and mean dlPFC activity in our data are consistent with the hypothesis that dlPFC is recruited to support cognitive control in proportion to declining performance. It is also possible that we found

evidence implicating the dACC, where the earlier study did not for methodological reasons. For example, a continuous effort-discounting measure of effort costs may simply be more precise than categorical self-report measures used in the previous study.

Among the wider set of nodes spanning the DMN, CO, FP, Sal, and DorAtt networks, many remain reliable predictors at $p < 0.05$. However, p-values are also attenuated when performance predictors are included and none survive FDR correction, as a consequence. For exploratory purposes, nodes with significant effects (at $p < 0.05$, uncorrected) are mapped in Figure 4.14. These nodes cluster mostly in the left dACC, and the bilateral vIPFC and IPS.




Figure 4.14 Nodes in which load-independent N-back activity predicts subjective effort costs, controlling for performance. All nodes are reliable at uncorrected $p < 0.05$. Colors indicate network membership and effect sign: red for DMN, black for Sal, green for DorAtt, and purple for CO, and dark yellow for a FP node with a positive load-independent effect on SV_{3S}.

The exploratory analyses (across all five a priori networks of interest) have largely recapitulated the a priori node tests, showing performance-independent effects of

average activity on subjective effort costs in the dorsal medial PFC, but not the dlPFC. In addition to these regions, a set of nodes in the vlPFC, AI and IPS also showed significant (if uncorrected) effects, spanning the FP, CO, Sal, and DorAtt networks.

Although none of the exploratory nodes survived correction after controlling for performance, in none of the models were performance variables themselves (d' or μRT ; all p 's ≥ 0.22) reliable predictors of SV_{3s} . This suggests that the attenuation of the reliability of the predictors of interest stems from the additional complexity of the larger models rather than shared variance, per se. In fact, a series of nested model comparisons reveals that additional performance predictors do not explain sufficient variance to justify the degrees of freedom for any node (all $X^2 p$'s ≥ 0.16). Also note that the use of other Ex-Gaussian response time parameters (σ or τ) did not fundamentally alter the pattern of results.

4.2.3 Limitations

As described above, one of the key limitations of the present experimental design is that task order was fixed and, as such confounded with load level. This may drive some of the unexpected results like the negative linear slopes across load, or the drop in activity at $N = 5$ relative to $N = 6$ in task-positive networks. Future designs could incorporate load order counterbalancing to resolve this confound. Nevertheless, the fixed load order was an intended aspect of the design, implemented to reduce between-subject variability (unrelated to effort) in the subjective experience of each load. In a counterbalanced design, a participant performing the $N = 6$ load at the very beginning of the session might have a different experience than one that performs $N = 6$ at the very end of the session, for reasons that have nothing to do with subjective effort, per se. Hence, fixed order

conferred validity benefits as well. Another limitation of this design was the rapid, fixed pacing of the N-back stimuli, which limit the ability to resolve event-related neural activity to different trial types (lures, targets, and non-targets). Again, however, this was also an intended feature of the experiment, which not only simplified the analytic approach, but also ensured that the cognitive engagement was maximally constant and continuous across each block. Nevertheless, future experiments could explore different approaches, such as optimizing the design for event-related analyses (or mixed event-block designs) and the contrasting of different trial types, by jittering stimulus presentation rates.

4.3 Summary

The hypothesis that a given region tracks cognitive effort during task engagement requires not only that activity in the region scales with objective cognitive load, but also that it scales with subjective experience. In this chapter, the encoding of objective load and subjective effort during N-back performance was investigated in four task-positive networks: the FP, Sal, CO, and DorAtt networks, and the task-negative DMN.

Across a wide range of loads ($N = 1$ — 6), all five networks showed distinct load functions, including monotonic decline in the DMN, inverted-U response profiles in the FP, Sal, and DorAtt networks, and flat, positive activity in the CO network. This set of patterns, especially including a flat function in the CO and decreasing activation in the DMN, along with performance statistics, supports the hypothesis that, rather than disengaging from very highly demanding loads, participants may have shifted their response strategies in an adaptive fashion. Particularly strong evidence against the disengagement hypothesis is that those performing the 3-back better have a more

pronounced inverted-U load function (across $N = 1-3$) as compared with those performing the 3-back worse, cf. (Jaeggi et al., 2007). The observed pattern of load functions also supports the interpretation that the FP, Sal, DorAtt, and DMN networks, in particular, track objective load and are thus also candidates for tracking cognitive effort.

Regarding subjective experience, the FP, Sal, DorAtt, and CO networks all showed higher levels of activity, across loads, for steeper discounters. That is, participants finding the N-back to be more effortful also showed more activity in these networks, regardless of load. However, contrary to the original hypothesis, no networks showed activity by load interactions predicting subjective effort. Thus although networks vary by load, this variation does not, itself, vary as a function of subjective effort. One tentative exception to this null result is that steep effort discounters showed higher activity at low loads ($N < 3$) and earlier peaking (at $N = 2$ versus $N = 3$), relative to shallow effort discounters. Interestingly, this distinction mirrored the performance distinction such that the load profile of steep effort discounters resembled that of better 3-back performers while shallow effort discounters resembled that of worse 3-back performers. This resemblance supports a prior hypothesis that high subjective effort constitutes a cue to shift strategies adaptively (thus maintaining performance at the demanding 3-back).

At the sub-network level, higher load-independent activity in a set of five a priori nodes of interest, including the dACC, and bilateral dlPFC and VS predicted greater subjective effort. This result mirrored the findings in the wider set of a priori networks. Interestingly, after controlling for performance, the dACC (along with the bilateral VS) remained a reliable predictor of subjective effort, though the dlPFC nodes did not. While

this perfectly contradicts a prior study showing effects in the dlPFC but not dACC (McGuire & Botvinick, 2010), it is consistent with strong prior hypothesis implicating the dACC, in particular, in tracking subjective effort.

Finally, a wider, exploratory investigation of all nodes across the five networks of interest revealed a set of 49 nodes surviving multiple comparisons correction in the left dlPFC, and also the bilateral IPS, dACC, pre-SMA, AI, and vlPFC. While load-independent activity in almost all of these nodes negatively predicted SV_{3S} , it positively predicted SV_{3S} in three DMN nodes and one FP node. The positively related DMN nodes mostly did not also encode load, and showed almost no reliable task-based deactivation, making their relationship with individual differences in subjective effort difficult to interpret. The FP node, by contrast, showed the same qualitative pattern as the wider FP network, with the particular distinction of showing lower, rather than higher load-independent activity among the low AUC group. Given that it was the only node showing an effect in this direction, stronger conclusions require additional study and replication. Finally, a smaller set of nodes, mostly in the vlPFC, IPS, and dACC/pre-SMA showed significant (if uncorrected) effects after controlling for performance.

Chapter 5: General Discussion

Cognitive effort is an attribute of everyday experience with important consequences for decision-making and action in normal behavior, and also in disorders ranging from schizophrenia to depression. Yet, despite its importance, very little is understood about what the brain is doing when a cognitive task feels effortful. The experiments described here were conducted to fill these gaps, by addressing two broad questions: 1) what brain regions track objective working memory load; and 2) what regions track subjective effort beyond variation by load.

To investigate these questions, subjective effort was operationalized in terms of effort-based reward discounting (COGED: the COGNitive Effort Discounting task). This novel approach conceptualizes effort as a cost that can be tracked during engagement with a demanding task. This cost is thought to be cached as an accessible quantity that can be recalled when presented with future opportunities to expend effort to earn reward. This type of operationalization avoids problems associated with traditional effort measures like self-report (which relies on potentially unreliable self-introspection and veridical report) or indirect measures like objective load, or physiology (which can only be linked to effort circularly; i.e., if one assumes they are linked to effort). The operationalization also provides conceptual clarity, so that the questions addressed by this work can be asked more precisely, in terms of what pattern of brain activity during N-back engagement reflects subjective costliness. Below, I briefly recapitulate the key results of the study and their implications for our understanding of the neural basis of cognitive effort representation.

5.1 Behavioral Results

The behavioral results of the current study confirm what has been observed in extensive piloting and two recent publications (Culbreth, Westbrook, Braver, & Barch, 2015; Westbrook et al., 2013), namely that individuals discount monetary rewards more strongly with increasing load in the N-back task. Increased discounting with increasing load supports the hypothesis that the experience of cognitive effort is state-dependent: some tasks are subjectively more costly than others. The observed cross-session stability in the extent to which individuals discount rewards in exchange for completing N-back tasks is a finding that supports the hypothesis that subjective experience of effort is also trait-like. Some individuals tend to find the tasks more costly than others, and this pattern of individual differences appears to be relatively stable across time. However, there was also variability in participants' discounting across sessions suggesting that other state factors, along with measurement noise, play a role as well. Though not investigated in the current study, the validity of the effort discounting as a trait construct has been supported by evidence that individual differences in discounting rates predict negative symptoms in schizophrenia (Culbreth et al., 2016), cognitive aging, and Need for Cognition (Westbrook et al., 2013). In the current study, the intentional restriction of the variance in N-back performance facilitates a clear demonstration that individual differences in subjective effort exist apart from differences in N-back performance or ability. Performance rates, in terms of d' and response times were practically identical between groups defined by a median split of average discounting rates. As noted above, the artificial selection of participants showing uniformly high performance, and also either very steep or very shallow discounting enables the current dataset to be used to demonstrate that differences in subjective effort can exist apart from differences in

performance. Moreover, the artificially selected sample confers the additional benefit for imaging analysis by emphasizing differences in brain activity that distinguish steep and shallow effort discounters apart from differences in performance. It is important to note, however, that the artificiality of the selection process also limits the study's generalizability. Indeed, pseudo-random selection in a prior study (via self-selection among undergraduates and older adults) has shown that performance is typically related to subjective effort (Westbrook et al., 2013).

5.2 Regions of the Brain Tracking Cognitive Effort

A number of regions were hypothesized to track cognitive effort. At the broadest level, a set of task-positive networks, including the FP, DorAtt, CO, and Sal networks are more active when participants are engaged in a demanding task, while the DMN is less active under such conditions (Braver et al., 1997; Callicott et al., 1999; Jaeggi et al., 2003; 2007; D. C. Park & Reuter-Lorenz, 2009). Moreover, activity in regions of these networks have been shown to vary with load (McKiernan et al., 2003; Pyka et al., 2009). Within these networks, particular nodes are thought to be especially relevant for tracking subjective effort because they have been hypothesized to support cognitive control processes (including regulating the intensity of control signals and maintaining task rules for biasing behavior), and these cognitive control processes, in turn are thought to be subjectively effortful (Westbrook & Braver, 2015). Canonical models of cognitive control implicate the dlPFC for working memory maintenance of rules for guiding behavior and the dACC for detecting control demands by tracking errors or conflict and upregulating control accordingly. Hence, these frontal regions were particularly strong candidates for tracking cognitive effort (McGuire & Botvinick, 2010).

To assess whether a brain region tracks cognitive effort, the first step was to identify regions that responded to task engagement and also varied by objective N-back load. Numerous prior studies of N-back activity have identified a core set of load functions across $N = 0$ — 3 (Owen et al., 2005). As characterized by one early study, these include inverted-U shapes, monotonic patterns, and flat functions that reflect task engagement but do not vary by load (Callicott et al., 1999). Yet, these load functions have never been examined beyond $N = 3$, nor have they been investigated in a set of networks independently defined, by their functional connectivity as in (Power et al., 2011). So, the present design generated novel data about load functions extended to extremely highly demanding load levels in independently-specified functional networks.

All three kinds of load functions characterized by Callicott et al. (1999) were also found in the present dataset, with some differences in the precise dynamics and locations. Namely, most task-positive networks showed inverted-U patterns, including the FP, Sal, and DorAtt networks, as defined by (Power et al., 2011). These networks span the very dlPFC and dACC regions hypothesized to be central to tracking effort. Inverted-U patterns in the task-positive networks peaked at $N = 3$ and so would have appeared as monotonic load functions by previous study designs. One caveat is that there was an unexpected rise in activity from $N = 5$ to $N = 6$ for most task-positive networks. Fixed task order precludes careful investigation of this effect, which could have arisen from some particular cognitive strategy that participants used to complete the 6-back, or something less interesting like rebound from fatigue, or participants' motivation from knowing that the 6-back was the ultimate load level.

The CO network, by contrast, showed a flat load function. Activity in this network was robustly positive across all levels of the task, but did not vary reliably across loads (there were no linear or quadratic trends of load across $N = 1-6$). This result suggests that the CO network simply encodes that a participant is engaged in a task or not, rather than cognitive load per se. Finally, approximately monotonically decreasing behavior was observed in the DMN. Decreasing activity with load has been observed in prior studies of the DMN (McKiernan et al., 2003; Pyka et al., 2009), but never at this wide range of loads. The fact that approximately monotonic behavior was observed across all levels $N = 1-6$ in the DMN is striking and implicates the DMN in monotonic load functions much more strongly than if only $N = 1-3$ had been investigated.

Inverted-U patterns present two complications for interpretation. First, they suggest the possibility that participants may be disengaging when task conditions become overly demanding. This hypothesis, however, is contradicted by evidence that performance remains high across all N-back loads, and also by the monotonic and flat load functions observed in the DMN and CO networks, respectively. The fact that activity patterns are not diminished at $N > 3$ in the CO and DMN implies that participants are remaining deeply engaged with the task, but instead might be shifting strategies when the N-back becomes very demanding. This hypothesis was originally suggested by a prior study which found that the highest 3-back performers actually showed a more pronounced inverted-U (bigger drop from a peak of activation for the 2-back) as compared with the lowest performers (Jaeggi et al., 2007). It was suggested that they performed better because they shifted strategies adaptively in a way that resulted in a stronger inverted-U pattern. The present dataset replicates this earlier result, supporting

that rather than disengaging (at, say $N = 3$ relative to $N = 2$), participants shifted strategies adaptively, helping them maintain higher 3-back performance.

The second complication with inverted-U patterns is that they are not readily reconcilable with monotonically increasing effort costs (discounting) with load. It is entirely possible that even though participants are shifting strategies at higher load levels, they still find those high-load strategies more subjectively costly. Nevertheless, strategy selection is a degree of freedom that cannot be accounted for by this experimental design. The best candidate for tracking effort by load alone, by this criterion, is therefore the DMN. The DMN, like SV is also monotonically decreasing as load increases. It is possible that subjective, phenomenal effort indexes the degree to which the DMN is suppressed while individuals are engaged with external tasks. This hypothesis is consistent with the observations that low-effort mind-wandering is associated with greater DMN activity (Schooler et al., 2011), and lapses of attention indexed by greater response time variability is associated with the extent to which the DMN and task-positive regions are anti-correlated (Kelly, Uddin, Biswal, Castellanos, & Milham, 2008).

Variation by load is only one piece of evidence that a region tracks cognitive effort; a region should also encode subjective effort beyond objective load. In the current study, subjective effort was defined by the extent to which individuals discounted rewards for completing N-back tasks. According to this definition, regions in which activity co-varies with discounting would be strongly implicated in tracking subjective effort. Interestingly, all a priori task-positive networks are good candidates in that their load-independent activity encoded SV_{3S} . Namely, greater load-independent activity predicted greater subjective effort, as quantified by steeper discounting. While this was

true in all task-positive networks, there was no relationship between load-independent activity in the DMN and discounting. This result is surprising and not easy to reconcile with the interpretation that the approximately monotonic DMN load function otherwise makes it a good candidate for tracking effort. That is, a monotonic load function suggests the DMN tracks (monotonically increasing) effort, but there is no evidence that the DMN tracks individual differences in subjective effort – even though the task-positive networks appear to track individual differences. Part of this discrepancy in whether a network tracks individual differences may have to do with functional heterogeneity in the DMN. While other networks showed consistently negative relationships between activity and AUC_{3S} , the DMN showed some nodes with negative, and three statistically reliable nodes with positive relationships. Hence, because these effects cancel each other out at the network level, the network-wide relationship would be weaker than more homogeneous networks.

One surprising and unpredicted finding of the study was that while load-independent activity predicted subjective effort, load-dependent changes in activity did not, controlling for load. That is, accounting for the relationship between objective load and SV_{3S} , no additional variance was explained by load-dependent changes in brain activity. This result may reflect insufficient power to detect a load-dependent effect.

Relatedly, there was no statistically reliable evidence of activity by load interactions in discounting. That is, none of the networks of interest evinced steeper (or shallower) load functions among those who find the N-back more effortful. Again, the lack of evidence for activity by load interactions may reflect insufficient power. Indeed, this interpretation is supported by the finding of subtler trends in key networks and

regions such as the DorAtt and dACC, for which activity in steep discounters numerically peaked at $N=2$, whereas shallow discounters' activity peaked at $N=3$.

Interestingly, this pattern of steep discounters peaking earlier mirrors the pattern of high 3-back performers peaking earlier, and provides tentative support to the hypothesis that high subjective effort is a cue to shift strategies adaptively at very high cognitive demands. This pattern is also intriguing in light of data on differential dlPFC load functions and subjective effort in older adults. Namely, higher load-independent activity and earlier peaking among those experiencing greater subjective effort mirrors findings that older adults show over-recruitment, at low loads, and under-recruitment at high loads, in the lateral PFC during working memory tasks (D. C. Park & Reuter-Lorenz, 2009; Reuter-Lorenz & Cappell, 2008), and moreover supports the widespread assumption older adults experience greater subjective effort for the same cognitive tasks (Hess, 2014). There is a striking convergence between the lines of evidence that 1) older adults experience the N-back as costlier, as we have shown using the same effort-discounting paradigm (Westbrook et al., 2013), and 2) that when comparing lateral PFC activity during a working memory task across age groups (Cappell et al., 2010) to our FP and DorAtt load functions, older adults look like our high-cost group while younger adults look like our low-cost group.

Follow-up studies are needed to investigate the lack of stronger evidence for load by activity interactions predicting subjective effort. Also, as mentioned above, it may be that the best indices of effort are not univariate block regression weights, but rather event-related responses to task events, or some combination of block and event-related responses (as probed by “state-item” response models). Additionally, the intensity of

engagement / effort may track load-specific changes in multivariate patterns of activity that are distinct from mean signal. These possibilities are discussed further under Future Directions.

Although load-dependent changes in brain activity did not relate to subjective effort, load-independent activity robustly predicted individual differences in subjective effort in all task-positive networks of interest and also the dlPFC and dACC regions of interest. Importantly, these findings also place a distinct interpretation on prior results in this domain. In particular, a previous, influential study found that self-reported desire to avoid demanding tasks measures related to brain activity in both of these regions, but not the dACC, when controlling for performance measures (McGuire & Botvinick, 2010). The McGuire & Botvinick result was surprising, in that it contradicted strong prior hypotheses regarding the role of the dACC in actively tracking control demands. Likewise the wealth of evidence regarding dACC and cognitive effort led to a recent, influential theoretical account in which the dACC up-regulates control in proportion to the expected costs and benefits of doing so (Shenhav et al., 2013). It is particularly interesting, therefore, that the present data showed just the opposite of the McGuire & Botvinick (2010) finding: here, discounting related to mean activity in the dACC, but not dlPFC nodes after controlling for performance. The lack of a relationship with dlPFC activity, after controlling for performance, is consistent with a well-supported hypothesis that the dlPFC is recruited to support cognitive control in proportion to flagging performance (Botvinick et al., 2001). One potential explanation for why we detected a performance-independent effect in the dACC when the prior study did not is that the

COGED measure might be more reliable or provide more precision for discriminating between individuals than categorical self-report ratings.

Finally, an exploratory analysis reveals a cluster of nodes in the bilateral vIPFC, IPS, and dACC / pre-SMA for encoding subjective effort apart from performance measures. While these nodes did not survive multiple-comparisons correction, the evidence that the performance measures themselves were not significant predictors of subjective effort suggests that these nodes did not survive correction chiefly because of limited power (for investigating individual differences). That is, it is more likely that the addition of performance predictors attenuated p-values for the load-independent activity predictor because they consumed degrees of freedom, rather than explaining shared covariance with subjective effort. As such, the set of nodes and regions implicated by this analysis warrant future studies targeted to investigate the encoding of subjective effort therein.

5.3 Future Directions

There are a number of future directions to take the question of what brain dynamics track subjective effort during task engagement, using the COGED paradigm. First, future experimental designs might be optimized to examine event-related responses to N-back stimuli rather than a block design. The block design is a reasonable first approach to quantifying task engagement, but this approach revealed that most block-wise load effects are clearly not monotonic in the brain the way they are in discounting (i.e., the monotonic COGED functions shown in Figure 3.4). Perhaps event-related activity patterns might exhibit monotonic load functions. Alternatively, there might be an important link between cognitive effort discounting and the relationship between

sustained and transient N-back activity, as could be observed in mixed block/event-related designs, e.g., (Reynolds, West, & Braver, 2009). For example, subjective effort may track closely with norepinephrine function, which, according to an influential theory, shows a higher ratio of event-related to tonic responses when reward frequency higher (Aston-Jones & Cohen, 2005). Event related analyses may also reveal information about strategy shifts and thus enrich our understanding of what happens at supra-capacity N-back loads. For example, shifting to a familiarity-based strategy could result in a progressively smaller distinction between event-related responses to lures and targets. Controlling for individual differences in strategies applied at a given task level could also increase power to detect individual differences in subjective cognitive effort.

It is also possible that load-specific changes might be more robustly encoded as multivariate patterns of activity in key cognitive control and working memory regions like the dlPFC, cf. (Etzel et al., 2015). In particular, there is a growing appreciation that working memory content related to cognitive control might be more reliably encoded in terms of multivariate activity patterns rather than in the mean amplitude of load-related activity (Riggall & Postle, 2012). Such multivariate patterns might also relate strongly to subjective effort, such that they could be used to decode differences in subjective effort between individuals or across loads. Importantly, effort might not be encoded in terms of mean activity across a block, but rather might be more sensitively detected as multivariate patterns of activity and their load-dependent changes across the block.

Another future direction involves looking at network connectivity dynamics during N-back engagement. For example, network connectivity properties could be used to investigate the hypothesis that conscious attention corresponds to a global workspace

of dynamically configured network components that are functionally coupled (Kitzbichler et al., 2011). As such, measures of increasing functional integration may covary with the intensity of focused attention brought to bear on a task, or the objective load of the task. These effects may relate to subjective cognitive effort in key ways. Less modularity and more global integration may obtain among shallow effort discounters who find attentional focus less costly, independent of objective load. Furthermore, given evidence that DMN deactivation is monotonically decreasing in load (like SV) and that it may reflect individual differences, it is reasonable to suspect that it is not just load-independent activity in task-positive networks that is costly, but the strength of the anti-correlation between task-positive and task-negative networks that tracks subjective effort.

The greater value of the questions investigated in this research lies in understanding how effort is tracked and how that information comes to influence cognitive motivation. As these are particularly consequential for disorders of anhedonia and anergia, another important future direction is to investigate how subjective effort is differentially encoded among those with, for example, depression and schizophrenia. A greater range in discounting or sharper group differences would afford greater power to detect subjectivity effects, while also elucidating the key nodes in the pathway for experiencing, learning, and deciding about cognitive effort. More importantly, such work could inform hypotheses regarding the neural circuits and pathways that are dysfunctional in psychopathology. This may also lead to future interventions to promote desirable effort. For example, neurostimulation methodologies such as TMS and tDCS might be useful in this domain, as a technique to “trick” the brain into thinking that it is experiencing less subjective effort than expected, i.e., by targeting key ROIs such as

dACC or AI. Or, another approach might be to focus on dopaminergic receptors that are preferentially expressed in the brain regions associated with integrating effort costs into action selection (e.g., dACC, DLPFC, VS). Of course, creating targeted interventions of this kind will require much more precise information about the how the brain tracks and makes decisions about cognitive effort. This dissertation provides an important first step in that direction.

References

- Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: Adaptive Gain and Optimal Performance. *Annual Review of Neuroscience*, 28(1), 403–450.
<http://doi.org/10.1146/annurev.neuro.28.061604.135709>
- Badre, D., & Frank, M. J. (2012). Mechanisms of Hierarchical Reinforcement Learning in Cortico-Striatal Circuits 2: Evidence from fMRI. *Cerebral Cortex*, 22(3), 527–536.
<http://doi.org/10.1093/cercor/bhr117>
- Barnes, A., Bullmore, E. T., & Suckling, J. (2009). Endogenous Human Brain Dynamics Recover Slowly Following Cognitive Effort. *PloS One*, 4(8), e6626.
<http://doi.org/10.1371/journal.pone.0006626.t001>
- Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, 76, 412–427.
<http://doi.org/10.1016/j.neuroimage.2013.02.063>
- Basten, U., Biele, G., Heekeren, H. R., & Fiebach, C. J. (2010). How the brain integrates costs and benefits during decision making. *Proceedings of the National Academy of Sciences of the United States of America*, 107(50), 21767–21772.
<http://doi.org/10.1073/pnas.0908104107>
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1), 289–300.
- Biswal, B., Yetkin, F., Haughton, V., & Hyde, J. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine*, 34(4), 537–541.
- 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., & Braver, T. S. (2015). Motivation and Cognitive Control: From Behavior to Neural Mechanism. *Annual Review of Psychology*, 66(1), 83–113.
<http://doi.org/10.1146/annurev-psych-010814-015044>
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108(3), 624–652.
<http://doi.org/10.1037//0033-295X.108.3.624>
- Botvinick, M. M., Huffstetler, S., & McGuire, J. T. (2009). Effort discounting in human nucleus accumbens. *Cognitive, Affective, & Behavioral Neuroscience*, 9(1), 16–27.

<http://doi.org/10.3758/CABN.9.1.16>

- Braver, T. S., Cohen, J. D., Nystrom, L. E., Jonides, J., Smith, E. E., & Noll, D. C. (1997). A parametric study of prefrontal cortex involvement in human working memory. *NeuroImage*, 5(1), 49–62.
- Cacioppo, J., Petty, R., Feinstein, J., & Jarvis, W. (1996). Dispositional differences in cognitive motivation: The life and times of individuals varying in need for cognition. *Psychological Bulletin*, 119(2), 197.
- Callicott, J. H., Mattay, V. S., Bertolino, A., Finn, K., Coppola, R., Frank, J. A., et al. (1999). Physiological characteristics of capacity constraints in working memory as revealed by functional MRI. *Cerebral Cortex*, 9(1), 20–26.
- Cappell, K. A., Gmeindl, L., & Reuter-Lorenz, P. A. (2010). Age differences in prefrontal recruitment during verbal working memory maintenance depend on memory load. *Cortex*, 46(4), 462–473. <http://doi.org/10.1016/j.cortex.2009.11.009>
- Chapman, G. B., & Johnson, E. J. (1999). Anchoring, activation, and the construction of values. *Organizational Behavior and Human Decision ...*
- Chatham, C. H., Frank, M. J., & Badre, D. (2014). Corticostriatal Output Gating during Selection from Working Memory. *Neuron*, 81(4), 930–942. <http://doi.org/10.1016/j.neuron.2014.01.002>
- Cowen, S. L., Davis, G. A., & Nitz, D. A. (2012). Anterior cingulate neurons in the rat map anticipated effort and reward to their associated action sequences. *Journal of Neurophysiology*, 107(9), 2393–2407. <http://doi.org/10.1152/jn.01012.2011>
- Craig, A. D. (2002). How do you feel? Interoception: the sense of the physiological condition of the body. *Nature Reviews Neuroscience*, 3(8), 655–666.
- Crosson, P. L., Walton, M. E., O'Reilly, J. X., Behrens, T. E. J., & Rushworth, M. F. S. (2009). Effort-based cost-benefit valuation and the human brain. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, 29(14), 4531.
- Culbreth, A., Westbrook, A., & Barch, D. (2016). Negative Symptoms Are Associated With an Increased Subjective Cost of Cognitive Effort. *Journal of Abnormal Psychology*. <http://doi.org/10.1037/abn0000153>
- Culbreth, A., Westbrook, A., Braver, T. S., & Barch, D. M. (2015). Cognitive Effort in Schizophrenia: A Behavioral Economic Approach (Vol. 41, p. S72). Presented at the Schizophrenia Bulletin, Oxford.
- Dick, P., & Katsuyuki, S. (2004). The prefrontal cortex and working memory: physiology and brain imaging. *Current Opinion in Neurobiology*, 14(2), 163–168. <http://doi.org/10.1016/j.conb.2004.03.003>
- Du, W., Green, L., & Myerson, J. (2002). Cross-cultural comparisons of discounting

- delayed and probabilistic rewards. *Psychological Record*, 52(4), 479–492.
- Esterman, M., Chiu, Y., Tamber-Rosenau, B., & Yantis, S. (2009). Decoding cognitive control in human parietal cortex. *Proceedings of the National Academy of Sciences*, 106(42), 17974. Retrieved from <http://www.pnas.org/content/106/42/17974.full.pdf>
- Estle, S., Green, L., Myerson, J., & Holt, D. (2006). Differential effects of amount on temporal and probability discounting of gains and losses. *Memory & Cognition*, 34(4), 914–928.
- Etzel, J. A., Cole, M. W., Zacks, J. M., Kay, K. N., & Braver, T. S. (2015). Reward Motivation Enhances Task Coding in Frontoparietal Cortex. *Cerebral Cortex*. <http://doi.org/10.1093/cercor/bhu327>
- Ewing, K., & Fairclough, S. (2010). The impact of working memory load on psychophysiological measures of mental effort and motivational disposition. In D. de Waard, A. Axelsson, M. Berglund, B. Peters, & C. Weickert (Eds.), *Human Factors: A system view of human, technology and organisation* (pp. 1–11). Maastricht: Shaker Publishing.
- Figner, B., Knoch, D., Johnson, E. J., Krosch, A. R., Lisanby, S. H., Fehr, E., & Weber, E. U. (2010). Lateral prefrontal cortex and self-control in intertemporal choice. *Nature Publishing Group*, 1–24. <http://doi.org/10.1038/nn.2516>
- Floresco, S. B., & Ghods-Sharifi, S. (2006). Amygdala-Prefrontal Cortical Circuitry Regulates Effort-Based Decision Making. *Cerebral Cortex*, 17(2), 251–260. <http://doi.org/10.1093/cercor/bhj143>
- Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences of the United States of America*, 102(27), 9673–9678.
- Frank, M. J., & Badre, D. (2012). Mechanisms of Hierarchical Reinforcement Learning in Corticostriatal Circuits 1: Computational Analysis. *Cerebral Cortex*, 22(3), 509–526. <http://doi.org/10.1093/cercor/bhr114>
- Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between frontal cortex and basal ganglia in working memory: a computational model. *Cognitive, Affective, & Behavioral Neuroscience*, 1(2), 137–160.
- Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2), 351–401.
- Gold, J. M., Kool, W., Botvinick, M. M., Hubzin, L., August, S., & Waltz, J. A. (2014). Cognitive effort avoidance and detection in people with schizophrenia. *Cognitive, Affective, & Behavioral Neuroscience*. <http://doi.org/10.3758/s13415-014-0308-5>
- Green, L., & Myerson, J. (2004). A Discounting Framework for Choice with Delayed and

- Probabilistic Rewards. *Psychological Bulletin*, 130(5), 769–792.
<http://doi.org/10.1037/0033-2909.130.5.769>
- Green, L., Myerson, J., Oliviera, L., & Chang, S. E. (2013). Delay discounting of monetary rewards over a wide range of amounts. *Journal of the Experimental Analysis of Behavior*. <http://doi.org/10.1002/jeab.45>
- Hammar, Å. (2009). Cognitive functioning in major depression – a summary. *Frontiers in Human Neuroscience*, 3. <http://doi.org/10.3389/neuro.09.026.2009>
- Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values of d' . *Behavior Research Methods*.
- Hess, T. M. (2014). Selective Engagement of Cognitive Resources: Motivational Influences on Older Adults' Cognitive Functioning. *Perspectives on Psychological Science*, 9(4), 388–407. <http://doi.org/10.1177/1745691614527465>
- Hess, T. M., & Ennis, G. E. (2011). Age differences in the effort and costs associated with cognitive activity. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 67(4), 447–455.
- Holroyd, C. B., & Yeung, N. (2012). Motivation of extended behaviors by anterior cingulate cortex. *Trends in Cognitive Sciences*, 16(2), 122–128.
<http://doi.org/10.1016/j.tics.2011.12.008>
- Hosking, J. G., Floresco, S. B., & Winstanley, C. A. (2014). Dopamine Antagonism Decreases Willingness to Expend Physical, but not Cognitive, Effort: a Comparison of Two Rodent Cost/Benefit Decision-Making Tasks. *Neuropsychopharmacology*, 40(4), 1005–1015. <http://doi.org/10.1038/npp.2014.285>
- Huettel, S. A., Stowe, C. J., Gordon, E. M., Warner, B. T., & Platt, M. L. (2006). Neural Signatures of Economic Preferences for Risk and Ambiguity. *Neuron*, 49(5), 765–775. <http://doi.org/10.1016/j.neuron.2006.01.024>
- Jaeggi, S., Buschkuhl, M., Etienne, A., Ozdoba, C., Perrig, W., & Nirrko, A. (2007). On how high performers keep cool brains in situations of cognitive overload. *Cognitive, Affective, & Behavioral Neuroscience*, 7(2), 1–15. Retrieved from http://www.apn.psy.unibe.ch/unibe/philhuman/psy/apn/content/e5616/e5621/e6314/e6333/files6334/Jaeggi_et_al2007_ger.pdf
- Jaeggi, S., Seewer, R., Nirrko, A., Eckstein, D., Schroth, G., Groner, R., & Gutbrod, K. (2003). Does excessive memory load attenuate activation in the prefrontal cortex? Load-dependent processing in single and dual tasks: functional magnetic resonance imaging study. *NeuroImage*, 19(2), 210–225.
- Jansma, J. M., Ramsey, N. F., de Zwart, J. A., Van Gelderen, P., & Duyn, J. H. (2007). fMRI study of effort and information processing in a working memory task. *Human Brain Mapping*, 28(5), 431–440. <http://doi.org/10.1002/hbm.20297>

- Jimura, K., Myerson, J., Hilgard, J., Keighley, J., Braver, T. S., & Green, L. (2011). Domain independence and stability in young and older adults's discounting of delayed rewards. *Behavioural Processes*, *87*(3), 253–259. <http://doi.org/10.1016/j.beproc.2011.04.006>
- Juvina, I., & Taatgen, N. A. (2007). Modeling control strategies in the N-back task. *Proceedings of ICCM - 2007 - 8th International Conference on Cognitive Modeling*, 73–78.
- Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, *10*(12), 1625–1633. <http://doi.org/10.1038/nn2007>
- Kahnt, T., Heinzle, J., Park, S. Q., & Haynes, J.-D. (2011). Decoding different roles for vmPFC and dlPFC in multi-attribute decision making. *NeuroImage*, *56*(2), 709–715. <http://doi.org/10.1016/j.neuroimage.2010.05.058>
- Kelly, A. M. C., Uddin, L. Q., Biswal, B. B., Castellanos, F. X., & Milham, M. P. (2008). Competition between functional brain networks mediates behavioral variability. *NeuroImage*, *39*(1), 527–537. <http://doi.org/10.1016/j.neuroimage.2007.08.008>
- Kennerley, S. W., Behrens, T. E. J., & Wallis, J. D. (2011). Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nature Neuroscience*, *14*(12), 1581–1589. <http://doi.org/10.1038/nn.2961>
- Kirchner, W. (1958). Age differences in short-term retention of rapidly changing information. *Journal of Experimental Psychology*, *55*(4), 352.
- Kitzbichler, M. G., Henson, R. N. A., Smith, M. L., Nathan, P. J., & Bullmore, E. T. (2011). Cognitive Effort Drives Workspace Configuration of Human Brain Functional Networks. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *31*(22), 8259–8270. <http://doi.org/10.1523/JNEUROSCI.0440-11.2011>
- 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), 665–682. <http://doi.org/10.1037/a0020198>
- Kool, W., McGuire, J. T., Wang, G. J., & Botvinick, M. M. (2013). Neural and Behavioral Evidence for an Intrinsic Cost of Self-Control. *PloS One*, *8*(8), e72626. <http://doi.org/10.1371/journal.pone.0072626.s001>
- Kouneiher, F., Charron, S., & Koechlin, E. (2009). Motivation and cognitive control in the human prefrontal cortex. *Nature Neuroscience*, *12*(7), 939–945. <http://doi.org/10.1038/nn.2321>
- Krebs, R. M., Boehler, C. N., Roberts, K. C., Song, A. W., & Woldorff, M. G. (2012). The involvement of the dopaminergic midbrain and cortico-striatal-thalamic circuits

in the integration of reward prospect and attentional task demands. *Cerebral Cortex*, 22(3), 607–615. <http://doi.org/doi:10.1093/cercor/bhr134>

- Kurniawan, I. T., Guitart-Masip, M., Dayan, P., & Dolan, R. J. (2013). Effort and Valuation in the Brain: The Effects of Anticipation and Execution. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, 33(14), 6160–6169. <http://doi.org/10.1523/JNEUROSCI.4777-12.2013>
- Kurniawan, I. T., Seymour, B., Talmi, D., Yoshida, W., Chater, N., & Dolan, R. J. (2010). Choosing to Make an Effort: The Role of Striatum in Signaling Physical Effort of a Chosen Action. *Journal of Neurophysiology*, 104(1), 313–321. <http://doi.org/10.1152/jn.00027.2010>
- Leech, R., Braga, R., & Sharp, D. J. (2012). Echoes of the Brain within the Posterior Cingulate Cortex. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, 32(1), 215–222. <http://doi.org/10.1523/JNEUROSCI.3689-11.2012>
- Levy, D. J., & Glimcher, P. W. (2012). The root of all value: a neural common currency for choice. *Current Opinion in Neurobiology*, 22(6), 1027–1038. <http://doi.org/10.1016/j.conb.2012.06.001>
- Locke, H. S., & Braver, T. S. (2008). Motivational influences on cognitive control: Behavior, brain activation, and individual differences. *Cognitive, Affective, & Behavioral Neuroscience*, 8(1), 99–112. <http://doi.org/10.3758/CABN.8.1.99>
- Massar, S. A. A., Libedinsky, C., Weiyan, C., Huettel, S. A., & Chee, M. W. L. (2015). Separate and overlapping brain areas encode subjective value during delay and effort discounting. *NeuroImage*, 36(6), 899–904. <http://doi.org/10.1016/j.neuroimage.2015.06.080>
- 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.
- McKiernan, K. A., Kaufman, J. N., Kucera-Thompson, J., & Binder, J. R. (2003). A parametric manipulation of factors affecting task-induced deactivation in functional neuroimaging. *Journal of Cognitive Neuroscience*, 15, 394–408. <http://doi.org/10.1162/089892903321593117>
- Meyniel, F., Sergent, C., Rigoux, L., Daunizeau, J., & Pessiglione, M. (2013). Neurocomputational account of how the human brain decides when to have a break. *Proceedings of the National Academy of Sciences*, 110(7), 2641–2646. <http://doi.org/10.1073/pnas.1211925110>
- Miller, E. (2000). The prefrontal cortex and cognitive control. *Nature Reviews Neuroscience*, 1(1), 59–65.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function.

Annual Review of Neuroscience, 24(1), 167–202.

- Mitchell, S. H., & Wilson, V. B. (2010). The subjective value of delayed and probabilistic outcomes: Outcome size matters for gains but not for losses. *Behavioural Processes*, 83(1), 36–40. <http://doi.org/10.1016/j.beproc.2009.09.003>
- Montague, P. R., King-Casas, B., & Cohen, J. D. (2006). Imaging valuation models in human choice. *Annual Review of Neuroscience*, 29, 417–448.
- Montague, P., Dayan, P., & Sejnowski, T. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *The Journal of Neuroscience*, 16(5), 1936.
- Mulert, C., Menzinger, E., Leicht, G., Pogarell, O., & Hegerl, U. (2005). Evidence for a close relationship between conscious effort and anterior cingulate cortex activity. *International Journal of Psychophysiology*, 56(1), 65–80.
- Myerson, J., & Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the Experimental Analysis of Behavior*, 64(3), 263.
- Naccache, L., Dehaene, S., Cohen, L., Habert, M., Guichart-Gomez, E., Galanaud, D., & Willer, J. (2005). Effortless control: executive attention and conscious feeling of mental effort are dissociable. *Neuropsychologia*, 43(9), 1318–1328.
- Niv, Y. (2009). Reinforcement Learning in the Brain. *Journal of Mathematical Psychology*, 53(3), 139–154. <http://doi.org/10.1016/j.jmp.2008.12.005>
- Ossandon, T., Jerbi, K., Vidal, J. R., Bayle, D. J., Henaff, M.-A., Jung, J., et al. (2011). Transient Suppression of Broadband Gamma Power in the Default-Mode Network Is Correlated with Task Complexity and Subject Performance. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, 31(41), 14521–14530. <http://doi.org/10.1523/JNEUROSCI.2483-11.2011>
- Otto, A. R., Skatova, A., Madlon-Kay, S., & Daw, N. D. (2015). Cognitive Control Predicts Use of Model-based Reinforcement Learning. *Journal of Cognitive Neuroscience*, 27(2), 319–333. http://doi.org/10.1162/jocn_a_00709
- Otto, T., Zijlstra, F. R. H., & Goebel, R. (2014). Neural correlates of mental effort evaluation-- involvement of structures related to self-awareness. *Social Cognitive and Affective Neuroscience*, 9(3), 307–315. <http://doi.org/10.1093/scan/nss136>
- Owen, A. M., McMillan, K. M., Laird, A. R., & Bullmore, E. (2005). N-back working memory paradigm: A meta-analysis of normative functional neuroimaging studies. *Human Brain Mapping*, 25(1), 46–59. <http://doi.org/10.1002/hbm.20131>
- O'Reilly, R. C., & Frank, M. J. (2006). Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation*, 18(2), 283–328.

- Padoa-Schioppa, C., & Assad, J. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature*, *441*(7090), 223–226.
- Park, D. C., & Reuter-Lorenz, P. (2009). The Adaptive Brain: Aging and Neurocognitive Scaffolding. *Annual Review of Psychology*, *60*(1), 173–196. <http://doi.org/10.1146/annurev.psych.59.103006.093656>
- Parvizi, J., Rangarajan, V., Shiner, W. R., Desai, N., & Greicius, M. D. (2013). The Will to Persevere Induced by Electrical Stimulation of the Human Cingulate Gyrus. *Neuron*, *80*(6), 1359–1367. <http://doi.org/10.1016/j.neuron.2013.10.057>
- Pearson, J. M., Heilbronner, S. R., Barack, D. L., Hayden, B. Y., & Platt, M. L. (2011). Posterior cingulate cortex: adapting behavior to a changing world. *Trends in Cognitive Sciences*, *15*(4), 143–151. <http://doi.org/10.1016/j.tics.2011.02.002>
- Pessoa, L., & Engelmann, J. B. (2010). Embedding reward signals into perception and cognition. *Frontiers in Neuroscience*, *4*. <http://doi.org/10.3389/fnins.2010.00017>
- Peters, J., & Buchel, C. (2010). Neural representations of subjective reward value. *Behavioural Brain Research*, *213*(2), 135–141. <http://doi.org/10.1016/j.bbr.2010.04.031>
- Pine, A., Seymour, B., Roiser, J. P., Bossaerts, P., Friston, K. J., Curran, H. V., & Dolan, R. J. (2009). Encoding of Marginal Utility across Time in the Human Brain. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *29*(30), 9575–9581. <http://doi.org/10.1523/JNEUROSCI.1126-09.2009>
- Pine, A., Shiner, T., Seymour, B., & Dolan, R. J. (2010). Dopamine, Time, and Impulsivity in Humans. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *30*(26), 8888–8896. <http://doi.org/10.1523/JNEUROSCI.6028-09.2010>
- Pochon, J. B., Levy, R., Fossati, P., Lehericy, S., Poline, J. B., Pillon, B., et al. (2002). The neural system that bridges reward and cognition in humans: an fMRI study. *Proceedings of the National Academy of Sciences of the United States of America*, *99*(8), 5669–5674.
- Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., et al. (2011). Functional Network Organization of the Human Brain. *Neuron*, *72*(4), 665–678. <http://doi.org/10.1016/j.neuron.2011.09.006>
- Prelec, D., & Loewenstein, G. (1991). Decision making over time and under uncertainty: A common approach. *Management Science*, 770–786.
- Prévost, C., Pessiglione, M., Météreau, E., Cléry-Melin, M.-L., & Dreher, J.-C. (2010). Separate valuation subsystems for delay and effort decision costs. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *30*(42), 14080–14090. <http://doi.org/10.1523/JNEUROSCI.2752-10.2010>

- Pyka, M., Beckmann, C. F., Schöning, S., Hauke, S., Heider, D., Kugel, H., et al. (2009). Impact of Working Memory Load on fMRI Resting State Pattern in Subsequent Resting Phases. *PloS One*, 4(9), e7198. <http://doi.org/10.1371/journal.pone.0007198.t001>
- Rachlin, H., Brown, J., & Cross, D. (2000). Discounting in judgments of delay and probability. *Journal of Behavioral Decision Making*, 13(2), 145–159.
- Reuter-Lorenz, P. A., & Cappell, K. A. (2008). Neurocognitive Aging and the Compensation Hypothesis. *Current Directions in Psychological Science*, 17(3), 177–182. <http://doi.org/10.1111/j.1467-8721.2008.00570.x>
- Reynolds, J. R., West, R., & Braver, T. (2009). Distinct Neural Circuits Support Transient and Sustained Processes in Prospective Memory and Working Memory. *Cerebral Cortex*, 19(5), 1208–1221. <http://doi.org/10.1093/cercor/bhn164>
- Riggall, A. C., & Postle, B. R. (2012). The Relationship between Working Memory Storage and Elevated Activity as Measured with Functional Magnetic Resonance Imaging. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, 32(38), 12990–12998. <http://doi.org/10.1523/JNEUROSCI.1892-12.2012>
- 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(9), 1161–1168. <http://doi.org/10.1038/nn1756>
- Ruge, H., Braver, T. S., & Meiran, N. (2009). Attention, intention, and strategy in preparatory control. *Neuropsychologia*, 47(7), 1670–1685.
- 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(9), 410–417. <http://doi.org/10.1016/j.tics.2004.07.009>
- Sakai, K. (2008). Task Set and Prefrontal Cortex. *Annual Review of Neuroscience*, 31(1), 219–245. <http://doi.org/10.1146/annurev.neuro.31.060407.125642>
- Satterthwaite, T. D., Ruparel, K., Loughead, J., Elliott, M. A., Gerraty, R. T., Calkins, M. E., et al. (2012). Being right is its own reward: Load and performance related ventral striatum activation to correct responses during a working memory task in youth. *NeuroImage*, 61(3), 723–729. <http://doi.org/10.1016/j.neuroimage.2012.03.060>
- Schmidt, L., Lebreton, M., Cléry-Melin, M.-L., Daunizeau, J., & Pessiglione, M. (2012). Neural Mechanisms Underlying Motivation of Mental Versus Physical Effort. *PLoS Biology*, 10(2), e1001266. <http://doi.org/10.1371/journal.pbio.1001266.g006>
- Schooler, J. W., Smallwood, J., Christoff, K., Handy, T. C., Reichle, E. D., & Sayette, M. A. (2011). Meta-awareness, perceptual decoupling and the wandering mind. *Trends in Cognitive Sciences*, 15(7), 319–326. <http://doi.org/10.1016/j.tics.2011.05.006>

- 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 : the Official Journal of the Society for Neuroscience*, 34(6), 2148–2154. <http://doi.org/10.1523/JNEUROSCI.1214-13.2014>
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593–1599.
- 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. <http://doi.org/10.1016/j.neuron.2013.07.007>
- Shenhav, A., Straccia, M. A., Cohen, J. D., & Botvinick, M. M. (2014). nn.3771. *Nature Neuroscience*, 17(9), 1249–1254. <http://doi.org/10.1038/nn.3771>
- Stumm, Von, S., Hell, B., & Chamorro-Premuzic, T. (2011). The Hungry Mind: Intellectual Curiosity Is the Third Pillar of Academic Performance. *Perspectives on Psychological Science*, 6(6), 574–588. <http://doi.org/10.1177/1745691611421204>
- Treadway, M. T., Buckholtz, J. W., Cowan, R. L., Woodward, N. D., Li, R., Ansari, M. S., et al. (2012). Dopaminergic Mechanisms of Individual Differences in Human Effort-Based Decision-Making. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, 32(18), 6170–6176. <http://doi.org/10.1523/JNEUROSCI.6459-11.2012>
- Treadway, M., Buckholtz, J., Schwartzman, A., Lambert, W., & Zald, D. (2009). Worth the “EEfRT?” The Effort Expenditure for Rewards Task as an Objective Measure of Motivation and Anhedonia. *PloS One*, 4(8), e6598.
- Vassena, E., Silvetti, M., Boehler, C. N., Achten, E., Fias, W., & Verguts, T. (2014). Overlapping Neural Systems Represent Cognitive Effort and Reward Anticipation. *PloS One*, 9(3), e91008. <http://doi.org/10.1371/journal.pone.0091008.t002>
- Vickery, T. J., Chun, M. M., & Lee, D. (2011). Ubiquity and Specificity of Reinforcement Signals throughout the Human Brain. *Neuron*, 72(1), 166–177. <http://doi.org/10.1016/j.neuron.2011.08.011>
- 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. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, 23(16), 6475.
- Walton, M. E., Groves, J., Jennings, K. A., Crosson, P. L., Sharp, T., Rushworth, M. F. S., & Bannerman, D. M. (2009). Comparing the role of the anterior cingulate cortex and 6-hydroxydopamine nucleus accumbens lesions on operant effort-based decision making. *European Journal of Neuroscience*, 29(8), 1678–1691. <http://doi.org/10.1111/j.1460-9568.2009.06726.x>

- 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*(8), 1302–1314.
- Wardle, M. C., Treadway, M. T., Mayo, L. M., Zald, D. H., & De Wit, H. (2011). Amping Up Effort: Effects of d-Amphetamine on Human Effort-Based Decision-Making. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *31*(46), 16597–16602. <http://doi.org/10.1523/JNEUROSCI.4387-11.2011>
- Westbrook, A., & Braver, T. S. (2015). Cognitive effort: A neuroeconomic approach. *Cognitive, Affective, & Behavioral Neuroscience*, *15*(2), 395–415. <http://doi.org/10.3758/s13415-015-0334-y>
- 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. <http://doi.org/10.1371/journal.pone.0068210.t004>
- Yeung, N., Nystrom, L. E., Aronson, J. A., & Cohen, J. D. (2006). Between-Task Competition and Cognitive Control in Task Switching. *The Journal of Neuroscience*, *26*(5), 1429–1438. <http://doi.org/10.1523/JNEUROSCI.3109-05.2006>

Appendix A

The main text described an investigation of brain activity patterns tracking subjective effort during task engagement. This is relevant for behavior because, as the COGED paradigm itself demonstrates, subjective effort costs influence cognitive motivation. That is, subjective effort costs learned during task performance influence the subjective value of an external reward and therefore an individual's drive to pursue the reward. In this Appendix, a closely related question is asked: what brain regions support effort-based decision-making? To address this broad question, two more focused questions are asked. 1) What brain regions are engaged by close offer comparison relative to trivial selection, and 2) what regions encode dimensions of choice like subjective value and its objective dimensions, reward and load magnitude?

A.1 Regions Engaged by Effort-Based Decision-Making

Most effort-based decision-making data comes from animal lesion and unit-recording studies while animals make decisions about physical effort, and a handful of brain-wide imaging studies of physical effort studies among humans. Very few studies address the neural mechanisms of cognitive effort-based decision-making directly. While caution is warranted in extrapolating results to cognitive effort, there may be considerable overlap in the systems mediating either (Schmidt et al., 2012). As suggested in the main text, the ACC, and the dorsal region in particular (dACC), appear to play a central role in effort-based decision making. Other key regions include the ventral striatum (VS), particularly a subregion of the VS known as the nucleus accumbens (NAcc), as well as the ventromedial prefrontal cortex (vmPFC, and orbitofrontal cortex), and insular cortex.

A diverse array of methodologies, paradigms, and species implicate the ACC in effort-based decision-making. One emerging theory is that the ACC is critical for selecting and motivating series of effortful actions by maintaining action-outcome associations in pursuit of valuable goals (Cowen, Davis, & Nitz, 2012; Holroyd & Yeung, 2012; Kennerley, Behrens, & Wallis, 2011; Rushworth, Walton, Kennerley, & Bannerman, 2004; Shenhav et al., 2013). Evidence that the ACC supports overcoming a prepotent bias against effort includes that ACC lesions have been shown repeatedly, in both rats and monkeys, to yield a bias against effort (Floresco & Ghods-Sharifi, 2006; Rudebeck, Walton, Smyth, Bannerman, & Rushworth, 2006; Walton et al., 2009; Walton, Kennerley, Bannerman, Phillips, & Rushworth, 2006). In a common paradigm, rats navigating a T-maze may select either a high reward arm with a large barrier to climb or another with a low reward and low barrier, thereby expressing their preference for effort. Rats that start with a preference for the high reward / high effort option will switch to the low reward / low effort option following ACC lesion. Lesioned rats will also switch back if effort is equalized, demonstrating that apparent preference is not the result of difficulties in learning, or decrements in perceiving reward value or reward-based decision-making, e.g. (Walton, Bannerman, Alterescu, & Rushworth, 2003). Thus the ACC appears critical for overcoming a bias against effort. Human fMRI studies also provide evidence that the ACC encodes demand during effort-based decision making (Croxson, Walton, O'Reilly, Behrens, & Rushworth, 2009; Kurniawan, Guitart-Masip, Dayan, & Dolan, 2013; Kurniawan et al., 2010; Prévost et al., 2010). In one study, human participants were cued to squeeze a handgrip either with high or low effort for a fixed period in order to make them eligible to win a probabilistic reward, or avoid a

probabilistic loss (Kurniawan et al., 2013). Critically, activity in the ACC and VS signaled anticipation of effort, and did so independently of whether the trial was gain or loss. Also, the inclusion of catch trials signaling effort, but not requiring subsequent response, enabled the investigators to rule out that the anticipation of effort in the ACC reflected motor preparation.

The dACC also features prominently in cognitive control and thus is likely play an important role in effort-based decision-making that involves cognitive control demands. An influential theory implicates the dACC in monitoring response conflict and up-regulating control signals in the dlPFC to resolve conflict (Botvinick et al., 2001). As described in the main text, demands for cognitive control and dlPFC recruitment may be closely linked with cognitive effort. Thus, assuming the dACC is responsible for regulating the intensity of control, it is well positioned to track the cost of cognitive effort. According to a more recent proposal, the dACC plays a more decision-making style role in that it selects among task sets based on the expected value of the outcomes of those task sets, using inputs from valuation regions like the vmPFC, and integrates the cost of control (effort costs) to determine the most valuable task sets (Shenhav et al., 2013).

The nucleus accumbens in the ventral striatum (VS) and VS dopamine (DA) may also be critical for effort-based decision-making. Phasic DA from midbrain neurons to targets in the VS has been hypothesized to encode unexpected rewards and reward cues for use in reinforcement learning about optimal behavior (Montague, Dayan, & Sejnowski, 1996; Schultz, Dayan, & Montague, 1997). VS activity thus encodes whether on-going events are better or worse than expected and the degree to which they are better

or worse. VS BOLD signals have correlated with reward value during decision-making and at reward receipt in numerous studies, see (Bartra, McGuire, & Kable, 2013) for meta-analysis). Also, these signals appear to incorporate costs, responding less vigorously to reward cues when costs are high in terms of both physical (Croxson et al., 2009; Kurniawan et al., 2010; 2013; Schmidt et al., 2012) and cognitive effort (Botvinick et al., 2009; Schmidt et al., 2012). For example, VS responses to reward cues were smaller following a series of trials with a large amount of task switching (high cognitive effort) than after a series of trials with little task switching (low cognitive effort) (Botvinick et al., 2009).

The insular cortex has been identified as component to a network of cortical and subcortical regions responsive to aversive stimuli known as the “pain matrix”, and may also respond to aversive cognitive effort (Meyniel, Sergent, Rigoux, Daunizeau, & Pessiglione, 2013; Prévost et al., 2010; Treadway et al., 2012). A recent conjoint fMRI-MEG study in which humans were allowed to squeeze a handgrip and rest as they pleased, but were rewarded for the duration of squeezing, implicates the insula in effort-based decisions, and specifically the decision to take a break (Meyniel et al., 2013). The investigators found evidence of an accumulating cost signal in the insula with time squeezing and that participants would take a break when this signal reached a decision bound. They further found that the slope of accumulation was greater when greater force was required and less when greater rewards were offered. Finally, they found that the duration of voluntary rest periods in between squeezing was impacted by reward magnitude, but not by force requirements.

Apart from these effort-specific hypotheses, there are more general decision-making regions that likely support cost-benefit decisions as well, especially including the vmPFC, the VS, the intraparietal sulcus (IPS) and the posterior cingulate cortex (PCC). The vmPFC has been implicated in encoding choice variables during economic decision-making (Montague, King-Casas, & Cohen, 2006). In the neighboring orbitofrontal cortex, neurons has been shown to scale with economic value, for example, integrating dimensions of both cost and benefit into a single common currency reflecting value during decision making (Padoa-Schioppa & Assad, 2006). A recent meta-analysis of human fMRI studies has identified the vmPFC, along with the VS, and PCC as part of a core valuation network: a network of regions encoding the subjective value of diverse rewards, whether primary, monetary, delayed, physically effortful, risky, etc. (Bartra et al., 2013). Finally, the IPS has been implicated in the accumulation of cost-benefit difference information (Basten, Biele, Heekeren, & Fiebach, 2010), much as it has historically been implicated in the accumulation of stimulus intensity differences during perceptual decision-making, suggesting it may play a role in decisions balancing reward value against effort.

A.2 Limitations of prior studies on cognitive effort-based decision-making

In the realm of valuation and decision-making there are, as mentioned above, numerous neuroeconomic studies that have investigated the brain regions encoding key choice dimensions that are components of the valuation process. The subjective value (SV) of rewards discounted by costs like delay, risk, and physical effort has been shown in meta-analyses to be encoded in a core valuation network encompassing the VS,

vmPFC, and PCC in particular (Bartra et al., 2013; Levy & Glimcher, 2012).

Nevertheless, it is currently unclear whether these same regions will also encode the SV of rewards discounted by cognitive effort costs. Regarding cognitive effort, there have been studies investigating the encoding of dimensions relevant to SV like cognitive load and incentive amount during effort anticipation, e.g. (Vassena et al., 2014), but these data do not speak directly to explicit valuation or decision-making.

Valuation of cognitive effort has been investigated indirectly in a study in which participants are cued that they will receive a reward (or not) after performing a high (frequent task switching), or low (infrequent switching) demand task (Botvinick et al., 2009). Interestingly, VS activity, which was higher for cues indicating reward versus no reward, was also lower on trials following high versus low cognitive demands. This pattern was interpreted as evidence that cognitive effort is encoded as a cost that discounts the value of a reward, and this discounted value is represented in the VS. While the interpretation is intriguing, it relies on key assumptions regarding the nature of VS activity during passive receipt of a reward cue, as opposed to during instrumental decision-making. Specifically, it assumes that during passive reward receipt, such reward cues evoke implicit valuation processes that incorporate cost information about recently exerted effort. It seems clear that such an assumption needs to be more directly tested within the context of cognitive effort-based decision-making

One study that did examine cognitive effort-based decision-making found evidence that striatum may represent both motivation and outcome value with regard to cognitive effort (Schouppe et al., 2014). The key finding was that when presented with a choice to perform a high- versus a low-effort demanding task, in the absence of external

reward, activity was higher in the striatum (and possibly the ACC – though the exact loci was ventral and anterior to classic dACC) when participants elected the high demand option. On forced choice trials, this pattern reversed, and activity was lower when participants were forced to choose the high demand option. The interpretation of this pattern is that it reflected high intrinsic motivation when participants freely selected the high effort option, but high costs when subjects were forced to select the high effort option. While this interpretation is consistent with other literature showing value coding in the VS, the lack of external rewards presents a confound in that the participants' motivation state must be inferred from the fact that they chose the high effort option. It is equally possible, for example, that high effort selection trials were unrelated to intrinsic motivation and yet showed higher VS encoding (under the free-choice condition) because they did not carefully consider the costs involved. Without explicit information about costs and benefits, it is hard to make strong inferences about what this pattern of VS and ACC brain responses represent.

Only one study, to date, has investigated SV encoding of cognitively effortful rewards during effort-based decision-making where the SV is determined independently from choice behavior, and used as a regressor during decision-making (Massar, Libedinsky, Weiyan, Huettel, & Chee, 2015). In this study, seven small clusters correlating positively with SV of discounted reward were identified in locations ranging from the cingulate gyrus to the IPL to the IFG. Some of the peak voxels, however, fell in white matter regions. More importantly, they did not report whether the clusters were encoding the task demands, offer amounts, or both dimensions. Thus it is unclear whether the clusters implicated by this study reflect SV or covarying dimensions (e.g. reward

amount). Also, during scanning, both offers were presented simultaneously, so it is unclear whether the regions identified truly encode SV, or just some combination of offers dimensions. As described in the next section, the core methodology of the present study follows the approach of Massar et al. (2015) most closely, while making key improvements that enhance the reliability of the results and the strength of the inferences that can be made.

A.3 Predictions about regions supporting effort-based decision-making

While participants are evaluating offers in the context of COGED-like decision-making, the brain should encode key choice features relevant to the computation of subjective value. The broader neuroeconomics literature provides a number of predictions based on the domain-general encoding of subjective value when evaluating offers, e.g. (Bartra et al., 2013; Levy & Glimcher, 2012). In particular, it is hypothesized that a core valuation network implicated in representing value in other domains will also support valuation process regarding cognitive effort. As such it is predicted that the SV of effort-discounted rewards will be reflected in activation patterns of the VS, the vmPFC, and the PCC (i.e., with activation in these regions tracking trial-by-trial fluctuations in SV).

During valuation of effortful rewards, there should also be a network of regions involved in integrating expected costs that communicates with regions tracking effort during effortful engagement to access cached information about expected costs. That is, regions identified for encoding subjective costliness during N-back task engagement should show some correspondence with regions encoding dimensions of SV (e.g. load) during effort-based decision-making. Since it is hypothesized that the ACC and DMN

will both track subjective effort monotonically, it is predicted that individual differences in mean N-back activity or load functions during N-back engagement in these regions will anticipate the encoding of costs during subsequent offer valuation. Specifically, since SV is encoded positively in core valuation regions (VS, vmPFC, PCC), stronger N-back responses in the ACC and DMN should relate to more negative deflections in valuation-related activity in the core valuation regions.

It is further hypothesized that regions otherwise engaged by negative or aversive outcomes (in particular the AI and ACC) will encode the costs of effort for integration into a valuation process (Crosson et al., 2009; Kurniawan et al., 2013; Prévost et al., 2010; Schmidt et al., 2012). According to the hypothesis that the dACC biases the selection of effortful rewards in general, and more particularly that the dACC regulates the selection and intensity of cognitive control signals in accordance with the expected value of control model (Shenhav et al., 2013) the dACC should encode both reward and load information during evaluation of effort-demanding rewards. Thus it is predicted that the ACC generally will encode cognitive load, but the dACC, in particular, will encode reward and load together during offer valuation.

Finally, apart from the specific encoding of SV during valuation, there should be regions specifically engaged in decision-processes that do not explicitly encode SV, but nevertheless support decision-making. Based on prior hypotheses that the ACC and IPS are involved in close offer comparison and cost-benefit integration, respectively, it is predicted that these regions will show more activity during difficult choices (between two offers close in value) than easy choices (when one offer is clearly superior). Note that the ACC is thus predicted to encode two distinct forms of cognitive load during decision

making – higher anticipated N-back effort negatively, and higher decision difficulty positively during valuation and decision-making.

A.4 Methods

A.4.1 In-scanner COGED trials

COGED trials were presented either as a series designed to identify subjective indifference points (and therefore SV) in all three sessions and also, during scanning in the third session, as a series of randomly-ordered trials that orthogonalized the variables: base amounts, N-back load level, and proximity to indifference points following indifference point identification. This last variable was employed to systematically vary the difficulty of the decision, as described further below.

The COGED procedure was modified during the third session, to include offers that varied with specific proximity to indifference points. Specifically, ten proximity settings were designed to yield trials in which offers were close to indifference points, or far away from indifference points. A proximity parameter defined the percent difference between the estimated indifference point and either \$0 or the base amount. Proximity parameter settings, ranging from -1.0 to 1.0 were $\{-1, -0.2, -0.2, -0.1, -0.1, 0.4, 0.4, 0.6, 0.6, 1.0\}$ where 0.0 would represent an offer for the easier 1-back that was equivalent to subjective indifference. At the limits, the offer for the easier 1-back task is \$0 (proximity = -1) or the full base amount (proximity = 1). These settings yield trivial decisions and are thus considered “catch” trials as opposed to regular decision trials. Note that regular decision trial proximity values are repeated twice each to reflect that these proximity settings are used twice as often as each of the catch trial settings. Also, note that regular decision trial proximity values are asymmetric with respect to the indifference point (0.0).

The decision to use asymmetric proximity values was made in piloting, based on the observation that pilot participants showed a bias towards choosing the hard task / large amount offer slightly more frequently than expected from their originally estimated indifference points. Thus, the positive proximity settings (0.4, 0.6) for which the 1-back offers is above indifference, are slightly larger in magnitude relative to the negative proximity settings (-0.1, -0.2) for which the 1-back offer is below indifference.

Asymmetric proximity values were designed as a counter-bias to yield the desired proportion of trials in which participants would select the option to which they are being counter-biased by the current offer (the easy task / small offer for positive values and the hard task / large offer for negative values). Crossing ten proximity settings with 15 task-amount pairs produced 150 trials in-scanner trials.

A.4.2 TEMPD Trials

Temporal discounting (TEMPD) was also conducted with a decision trial structure mirroring that used in COGED. The intent of scanning TEMPD trials was to provide a baseline for comparison with effort discounting results. Delay discounting is much more well studied, and provides a rich prior body of work against which to compare effort-based decision-making results.

Like COGED, three different amounts {\$10, \$15, \$20} were offered in TEMPD trials, at five different delays {2 weeks, 1 month, 2 months, 6 months, 1 year} versus cash payment today, for a total of 15 delay-amount pairs. In the first session, 5 decision trials for each delay-amount pair were presented to estimate each indifference point, for a total of 75 decision trials. TEMPD was repeated again in the third session immediately prior to scanning, and, like COGED, was also modified for scanning, with 150 decision trials

performed in the scanner. As with COGED, these trials were also defined by 10 different proximity settings $\{-1, -0.2, -0.2, -0.1, -0.1, 0.4, 0.4, 0.6, 0.6, 1.0\}$ for each of the 15 indifference points for 15 delay-amount pairs. A key difference between the paradigms is that in COGED participants were paid for repeating a task, based on the instruction that they “maintain their effort” (though in reality, all participants were paid for every repetition) while in TEMPD participants were either paid immediately with cash (in case an immediate choice was randomly selected from among all choices made by the participant), or at some delay by Amazon gift card delivered electronically at the corresponding delay (in case a delayed choice was randomly selected to be paid).

A.4.3 In-Scanner Delay- and Effort-Based Decision Trials

Outside the scanner, participants were given unlimited time to make each choice during indifference point identification. In scanner, by contrast, task trials were adapted to the constraints of the fMRI design, and to facilitate independent analyses of brain activity during key trial phases. Specifically, decision-making was constrained by a particular temporal structure designed to separate pure offer evaluation from comparative decision-making processes. Namely, each fixed duration (13 sec) trial began with exactly 6 sec of a “pure evaluation” phase in which only the high effort (or delayed) reward was presented (e.g. \$4 for the red task or \$15 in 6 months). Participants were explicitly instructed to consider “how much [they] like the offer by itself”, but not to respond otherwise. Next, was a decision-making phase in which the second offer (always the easiest 1-back “black” task in COGED, or an immediate payout in TEMPD, for an amount specified by the indifference point and proximity parameter for that trial) was presented alongside the first offer, for up to 5.25 seconds during which participants

selected their preferred choice. After making their response, or at the end of the decision-making window, participants received feedback either of their choice on that trial in the form of a black box surrounding their selection, or text indicating that they missed the response window for that trial. Trials in which participants missed the response window were not repeated. The non-response rate was very low with a median of 0% all participants, and the mean was 1.15% for COGED and a median of 0% and a mean of 0.37% for TEMPD.

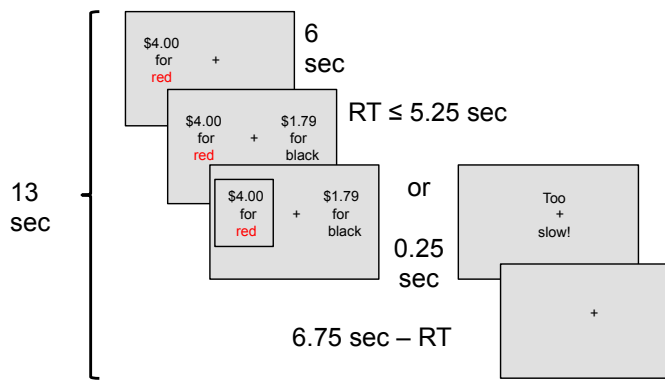


Figure A.1 Decision trial stimulus and stimulus timing for in-scanner decision trials.

A.4.4 Procedure

In the third session, participants were re-screened and consented. Then they conducted a shortened version of each level of the N-back task outside of the scanner. Specifically, they completed two rounds of each N-back level (instead of three) to re-familiarize them with the task. While still outside the scanner, participants next made a series of 75 COGED choices to establish indifference points for 15 task-amount pairs. Then participants made a series of 75 TEMPD choices to establish indifference points for 15 delay-amount pairs. After indifference points were established, participants were brought into the scanner.

As described above, the third session consisted of 6 scanning runs, broken up into a set of 3 COGED and 3 TEMPDP runs each, with short duration breaks (minimum 2 minutes) of rest given between each run (with the exception of the first 9 participants – see above for full details). These breaks were either filled with T1 or T2 scans, or rest. Across the 6 runs (or 2 runs for the first cohort), 150 COGED trials and 150 TEMPDP trials were performed, each generated by crossing 15 indifference points with 10 proximity settings. Although the precise order of scanning runs varied across participants, all participants completed COGED trials before performing any TEMPDP trials. For the main cohort, each of the 6 runs consisted of 345 scans (790 sec; 13.16 min) and approximately 50 decision-making trials. After the session, a single COGED trial was randomly selected to identify which task would be repeated for pay. Likewise, a single TEMPDP trial was randomly selected, identifying at which delay participants would receive their reward.

For COGED and TEMPDP imaging runs there were two cohorts that had slightly different settings. In an earlier cohort (the first 9 participants) all COGED trials were collected in a single run of 1019 volumes, (all TEMPDP trials were also collected in a single run of 1019 volumes). For the subsequent cohort, a decision was made to break up the decision-trials and provide more between trial rest. Consequently, for the subsequent cohort (the last 16 participants), decision trials were broken up into three runs for COGED and another three runs for TEMPDP, for a total of six imaging runs, each of which comprised 345 volumes for all remaining participants.

A.4.5 Image processing and modeling

As with the N-back image data, motion censoring was applied prior to modeling GLMs. After evaluation motion parameters, one participant was excluded because of excessive motion – 40% of their frames exceeded the motion censoring threshold. Mean fraction of censored frames was 5.5% across the remaining 24 participants, with a median of 5.5%, and a range of 0% to 16.8%.

For modeling activity during effort-based decision-making, GLMs used stick functions designating the onset of three types of events: evaluation phases, decision phases, and infrequent text displays of the menu of cost settings that participants might encounter, in order of increasing cost (e.g. a list with “black, red, blue ...” tasks, or a list with “2 weeks, 1 month, 2 months ...”). These menu displays were intended to remind participants of the potential costs they encounter on subsequent trials and their ordering in terms of cost magnitude. Stick functions were variously convolved with a simple gamma function, or as finite impulse response (FIR; “tent functions”), and in some instances amplitudes of gamma and FIR functions were parametrically modulated using the `-stim_times_AM2` argument to `3dDeconvolve`. The specific GLM modeling approaches associated with different analyses are described in the Results section.

A.5 Behavioral Results

A.5.1 N-back Behavioral Results

As shown in Figure A.2, just like in the second session, performance remained high, but was monotonically decreasing with N-back load. Also, like the second sessions, participants in the high AUC and low AUC groups did not differ in performance. An ANOVA reveals that N-back d' varied linearly by Load ($F_{1,140} = 131.4, p < 0.01$) and there was a trending difference by Group ($F_{1,140} = 3.23, p = 0.07$), but no interaction $F_{1,140}$

= 0.77, $p = 0.38$). It is worth noting that a Group difference of higher performance for the Low AUC group would contradict the hypothesis that individuals discount more because they perform more poorly. However, the group difference, while marginally significant across loads, was not significant at any single load (all p 's ≥ 0.17).

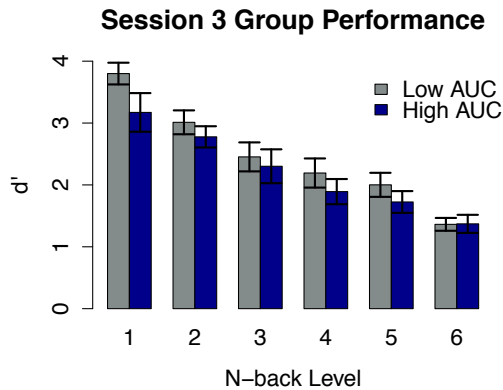


Figure A.2 Group performance d' by load and by AUC group for Session 3.

As shown in Figure A.3, the “lure rate” (percentage of false alarm lure trials), does vary with load ($F_{5,137} = 5.85$, $p < 0.01$ in Session 3), appearing to increase up to $N = 3$, and then remain constant after. Just as for Session 2, post-hoc pairwise comparisons reveal that the only reliable differences are between a higher lure rate for levels $N \geq 2$ as compared with $N = 1$ ($p < 0.01$). Importantly, however, as was shown for d' , there was no reliable group difference in lure rates in Session 3 ($p = 0.15$), supporting that participants in both groups rely on familiarity to a similar extent.




Figure A.3 Lure rates across all loads and both groups for Session 3.

Response time (RT – as measured by the central-tendency parameter μ RT from ex-Gaussian analyses of each participants’ distribution of response times on correct trials) is statistically indistinguishable across Groups in Session 3 ($F_{1,137} = 1.91, p = 0.17$). This is also true when using the tail of each groups’ RT distribution: τ RT, ($F_{1,137} = 1.20, p = 0.27$, respectively). Unlike Session 2, μ RT was found to vary with load in Session 3 ($F_{5,137} = 2.32, p < 0.05$). This tendency reflects an apparent RT slowing with load up to the 3-back, but asymptotically, remaining flat at higher loads. The only reliable pairwise differences in Session 3 RT values are between N = 1 and other task loads (all p 's < 0.01). All other task loads are statistically indistinguishable.




Figure A.4 μ RT across loads and groups for Session 3.

Participants, on average, performed better on N-back re-do trials (after COGED decision-making) than they did in the prior practice in Session 2. In Session 3, however, average performance did decrease slightly with a d' drop of -0.16 units in the low AUC group and -0.24 in the high AUC group, however, performance remained high across all levels that were repeated. To verify this assertion, performance was compared against the d' cutoffs that were originally used to define inclusion into the study. Across both sessions and groups, participants performed re-do trials above these cutoffs in 43 cases, and fell below cutoffs in only five cases (in Session 2, two high AUC participants on the 5-back $d' = 0.92$ and $d' = 0.68$, and in Session 3, one high AUC participants on the 3-back $d' = 1.17$, and one on the 6-back $d' = 0.94$, and one low AUC participant on the 5-back $d' = 0.47$). Importantly, as with Session 2, high and low AUC groups did not differ in terms of performance differences between prior exposure and re-do trials ($p = 0.70$).

A.5.2 COGED Decision-Making Behavior

As shown in Figure A.5, discounting was reliable and monotonic, as it was for Session 2, such that estimated subjective value reliably decreased with load across all levels of the N-back task in Session 3. Formally, Load x Amount ANOVAs reveal reliable effects of Load ($F_{4,137} = 31.6$, $p < 0.01$), but no effects of Amount ($F_{2,137} = 1.75$, $p = 0.18$).




Figure A.5 Indifference points across N-back task levels and base offer amounts for COGED decision-making in Session 3.

A.5.3 COGED In-Scanner Decision-Making

In scanner, participants made a series of choices between variable offers for the 1-back, and each of the task-base amount pairs, where the variable offers were set according to a proximity parameter controlling whether the offer for the easier (1-back) task is above or below, and proximity to, subjective indifference between offers over 150 trials (see Appendix A Methods for full details). Indifference was established by a standard COGED procedure, immediately prior to scanning in Session 3. Unexpectedly, three participants showed no discounting during indifference point determination, only selecting the more demanding option for larger reward. Because this pattern of choice behavior is uninterpretable, these three participants were removed, leaving a total of 21 participants from further analyses of choice behavior and brain imaging analyses for Session 3.

Bias settings were established in piloting to yield two kinds of effects: first, proximity parameters should bias participants towards the low or high demand option depending on whether it modulates the offer for the 1-back task above, or below

indifference, respectively. Second, biasing should make decision-making easy or difficult, depending on whether the offer for the 1-back is far from, or close to, estimated indifference. As shown in Figure A.6, participants showed a consistent and robust effect of bias on choice towards or away from the easy offer. When the proximity parameter was negative (the offer for the 1-back was below indifference) participants were reliably more likely to choose the more demanding option at every parameter value (all p 's < 0.01). When it was positive, participants were reliably more likely to choose the 1-back (all p 's < 0.01). Though bias manipulations were generally effective, when broken down further, particular task N-back levels and bias settings were indistinguishable from 50% choice frequency. Namely, participants' choices on the 2- and 3-back were indistinguishable from chance at proximity parameter values of +0.4 (both p 's ≥ 0.09) and +0.6 (both p 's ≥ 0.07). The same was true of participants' choices on the 6-back at a parameter value of -0.2 ($p = 0.09$). All other bias and load settings were distinguishable from chance (all p 's < 0.05).




Figure A.6 Choice frequency by task load and by choice biasing condition during effort-based decision trials.

While a robust effect of bias was anticipated, N-back task level was not anticipated to affect choice frequency. Surprisingly, however, it did. Formally, a repeated measures Level x Bias ANOVA revealed a linear effect of both Level ($F_{1,20} = 24.5$, $p < 0.01$) and of Bias ($F_{1,20} = 341$, $p < 0.01$), but no interaction between the two factors ($F_{1,20} = 2.38$, $p = 0.14$). As clearly shown in Figure A.6, the linear effect of Level is that participants are less likely to choose the more demanding option, for a given proximity parameter setting, as task level increases. For example, at a parameter setting of +0.6, participants were reliably more likely to choose the hard task if it was the 2-back, than if that hard task was the 6-back ($p < 0.01$), despite the fact that the bias is exactly 60% above the estimated indifference point in both cases. This result was not anticipated and it suggests either a systematic decision-making feature that is not adequately controlled for in the COGED paradigm, confounding the initial estimation of indifference points, or a post-COGED shift in decision strategies. For example, participants may begin to assign relatively more weight to differential reward amounts (rather than differential task demands) when deciding about lower levels of the harder task (the 2- and 3-back versus 5- and 6-back). As a consequence, the biasing effect was not as strong for some task levels and proximity parameter settings (e.g. +0.6 for the 2-back) as anticipated. Nevertheless, over the full range of parameter settings, bias has the desired directional effects on choice frequencies (excepting the cases noted above).

Proximity parameter settings were designed to control not only the direction of the bias, but also the strength and therefore the difficulty of the decision. Smaller magnitude parameter settings are closer to estimated indifference and constitute putatively more difficult choices. At the limit, decision-making is trivial, reducing to

simply identifying the 1-back option (when the proximity parameter is +1.0, and the offer amounts are equivalent and therefore irrelevant), or selecting the larger amount when the amount differential is maximal (the parameter value is -1.0, and \$0 are offered for the 1-back). One measure of choice difficulty is choice frequency itself. At the limits just described, participants selected the 1-back 95.3% of the time, on average, when the offer amounts were equivalent, and selected the 1-back only 4.8% of the time when \$0 were offered, supporting that decision-making was trivial at these limits. Choice frequencies were much less biased at other proximity parameter settings.

Another measure of choice difficulty is choice RT. When trials are broken down by bias manipulation, and those on which participants' choice conforms to the bias manipulation (pro-bias), and those contradicting it (anti-bias), two distinct patterns emerge. First, among pro-bias choices, catch trials remain, on average, faster than regular trials ($p < 0.01$). There is also a numeric pattern by which those bias settings closest to indifference (-0.1 and 0.4) are slightly slower than those further from indifference (-0.2 and 0.6). For example, the average median RT for pro-bias choices at a proximity parameter setting of 0.4 is trending ($p = 0.06$) slower than pro-bias choices at a parameter setting of -0.2. Slower decision-making as bias approximates zero also supports that the bias manipulation successfully modulated choice difficulty. Another pattern, shown in Figure A.7, is that choices participants made in opposition to the bias manipulation, on a given trial, are reliably slower ($p < 0.01$), on average, than those choices conforming to the bias manipulation. Finally, as shown in the Supplement, the bias manipulation had large effects on RT, while other choice trial parameters (load and amount) did not.

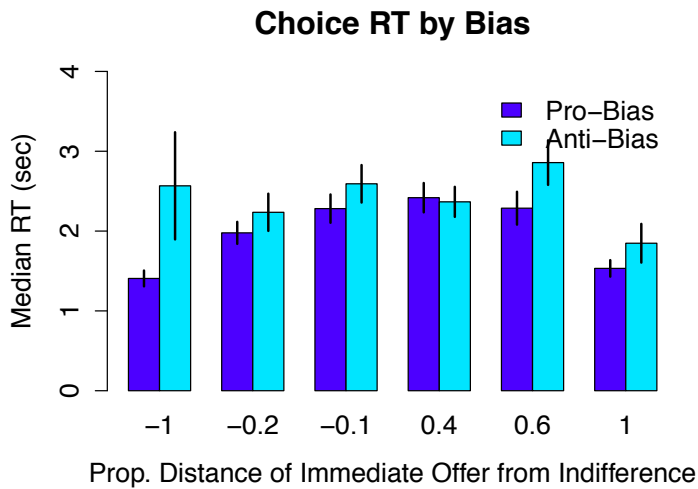


Figure A.7 Choice RTs during effort-based decision-making as a function of pro- versus anti-bias choices and biasing condition.

A.5.4 TEMPD In-Scanner Decision-Making

In addition to COGED trials, participants were also offered 150 TEMPD (intertemporal) choices. As shown in Figure A.8, choice patterns again reflect the intended influence of delay. That is, participants discounted the value of monetary rewards at all delays relative to payment today. Moreover, indifference points decreased with increasing delay showing that, as expected, participants find increasing delays to be increasingly costly. Also, as was observed with COGED trials, the data suggest a small trend of increasing SV with larger amounts, but again the anticipated increase in SV with amounts was not reliable. Formally, Delay x Amount ANOVAs reveal reliable effects of Delay ($F_{4,308} = 37.4, p < 0.01$), but no effect of Amount ($F_{2,308} = 0.37, p = 0.76$). Consequently, subsequent TEMPD analyses collapse across amounts, except where indicated.




Figure A.8 TEMPD SV trials for each delay and amount.

As with in-scanner COGED decisions, in-scanner TEMPD choices were based on offers parametrically adjusted to bias participants towards the higher (delayed) or lower (immediate) offer, and also by offers that are close to, or far from indifference, based on the proximity parameter value. As shown in Figure A.9, participants show a consistent and robust effect of bias on choice towards or away from the immediate offer. When the proximity parameter was negative (the immediate offer was below indifference) participants were reliably more likely to choose the delayed reward at every parameter value (all p 's < 0.01). When it was positive, participants were reliably more likely to choose the immediate reward (all p 's < 0.01). Though bias manipulations were generally effective, when broken down further, particular task delays and bias settings were indistinguishable from 50% choice frequency. Namely, participants' choices for delays of 2 weeks and 1 month were indistinguishable from chance at proximity parameter values of +0.4 (both p 's \geq 0.20) and +0.6 ($p = 0.10$ and $p = 0.06$, respectively). All other bias and delay settings were distinguishable from chance (all p 's < 0.05).




Figure A.9 Choice frequency by delay and by choice biasing condition during intertemporal decision trials.

While a robust effect of bias was anticipated, delay (like N-back level) was not anticipated to affect choice frequency. A repeated measures Delay x Bias ANOVA revealed a trending linear effect of both Delay ($F_{1,20} = 4.07, p = 0.06$) and a reliable effect of Bias ($F_{1,20} = 1290, p < 0.01$), but no interaction between the two factors ($F_{1,20} = 2.57, p = 0.12$). As clearly shown in Figure A.9, the linear effect of delay occurred because participants were less likely to choose the delayed option, for a given proximity parameter setting, as delay increased. For example, at a parameter setting of +0.6, participants were more likely to choose the delayed reward if it was only delayed by 2 weeks, than if it were delayed by a year (at trend-level $p = 0.06$). Again, it is important to note that this effect was unexpected because the bias is exactly 60% above the estimated indifference point in both cases. Nevertheless, over the full range of parameter settings, bias has the desired directional effects on choice frequencies (excepting the cases noted above).

Again, there is evidence that bias modulated decision difficulty as intended, in that offers closer to estimated indifference appeared to index more difficult decisions. On

catch trials, participant choice patterns were nearly uniform (the proximity parameter value is -1.0, and \$0 is offered immediately, or the parameter is +1.0 and \$10, \$15, or \$20 is offered immediately). At these limits, participants selected the immediate reward 99.6% of the time, on average, when the offer amounts were equivalent, and never selected the immediate offer when \$0 were offered, supporting that decision-making was trivial in catch trials. Choice frequencies were much less strongly biased at other proximity parameter settings.

Also, as with COGED, there is clear evidence that the bias manipulation influenced choice difficulty. When choices are broken down into those conforming to the bias manipulation (pro-bias), and those contradicting it (anti-bias), two distinct patterns emerge. First, among the pro-bias choices, catch trials remain, on average, much faster than regular trials. Also, mirroring COGED choices, there is also a numerical pattern (that is not statistically significant) by which decisions at those bias settings closest to indifference (-0.1 and 0.4) are slightly slower than those further from indifference (-0.2 and 0.6). Another pattern, shown in Figure A.10, is that choices participants made in opposition to the bias manipulation, on a given trial, are reliably slower ($p < 0.01$), on average, than those choices conforming to the bias manipulation.




Figure A.10 Choice RTs during intertemporal decision-making as a function of pro- versus anti-bias choices and biasing condition.

A.5.5 COGED Versus TEMPD

As with COGED, an AUC index can also be computed for temporal discounting (AUC_{TMP}). As shown in Figure A.11, there is a strong trend of higher COGED AUC predicting higher TEMPD AUC. Including the full dataset (all 25 participants), the effect (shown by dashed line) is not statistically significant ($p = 0.12$). However, prior studies have suggested a relationship between the two variables, whereby steep effort discounting participants (low COGED AUC) are very likely to be steep delay discounting participants (low TEMPD AUC), but shallow effort discounting participants show the full range of steep to shallow delay discounting. The current dataset is largely consistent with this pattern with the exception of the one participant showing both very steep effort discounting and very shallow delay discounting ($AUC_{3S} = 0.31$, and $AUC_{TMP} = 0.87$, respectively).




Figure A.11 Comparison of delay- and effort-based discounting AUC.

A.6 Neuroimaging Results: Brain Regions Engaged in Cost-Benefit Decision-Making

Two kinds of questions can help ascertain which parts of the brain are engaged during economic decision-making, in integrating the costs and benefits of choice options. The first asks which regions are specifically engaged by the cognitive processes associated with offer comparison, over and above those regions that are involved in offer perception, or response mapping and motor planning and execution. The second asks where activity encodes critical choice dimensions of costs and of benefits, or both (e.g., in terms of SV). This chapter addresses both questions in turn.

A.6.1 Regions Engaged by Close Offer Comparison

A first question asks which regions of the brain are engaged by offer comparison, over and above perceptuo-motor processes. The answer to this question can be addressed by comparing patterns of activity under difficult decision-making (when offers are close in subjective value) to trivial identification and indication of a much-preferred alternative. As described in the methods and verified by behavioral analyses, the trial structure includes decisions in which the offers are approximately close in terms of participants'

subjective values (indifference points) between easier, smaller rewards, and more effortful, larger rewards. These close proximity choices are more likely to promote involved weighing of costs and benefits. The trial structure also includes decisions in which the offers are very far apart such that one offer is clearly superior to the other, rendering the decision trivial. These trials were designed both as catch trials, to prevent participants from engaging in quick and simple rules (e.g. always selecting the larger offer), and also to provide a contrast with regular (close proximity trials). The contrast between regular and catch trials provides a means to identify regions that are engaged by the presumably cognitively elaborated process of comparing decision options in terms of costs and benefits, over and above the simpler processes of offer perception and motor execution for response selection.

As described in the methods, in the 2nd fMRI session, participants made a series of decisions in the scanner, between monetary offers contingent on either re-doing N-back tasks of a specific level (effortful rewards) or on waiting a specific time (delayed rewards). Each trial began with one offer presented in isolation, hereafter referred to as the “valuation period”, during which participants were instructed to consider how much they like that offer. The valuation period was designed so that choice dimensions of costs and benefits could be unambiguously represented (for a single offer). This valuation period always featured the larger, more effortful (delayed) reward, and lasted 6 seconds, before the smaller, less effortful (immediate) reward was also presented and the participant made their choice. The time between the second offer onset and the participant response is hereafter referred to as the “decision period”.

Activity arising during the decision period (accounting for hemodynamic lag) was contrasted across regular and catch trials to determine which regions of the brain are engaged by cost-benefit decision-making. Accordingly, GLMs were built for both delayed and effortful rewards, coding valuation window onsets and a separate decision window onset for each of the two trial types: regular and catch. The GLMs also included an onset for the infrequent display of the menu of cost options that may be presented prior to any given trial (text lists of the possible N-back levels or delays). For these analyses (unlike later ones), stick functions coding each event onset were convolved with a canonical (gamma) hemodynamic response function. Regression weights for regular and catch trial onsets were then compared via t-test at the group level.

Prior studies have shown particular engagement during difficult decision making of some of the same working memory, attention, and control networks tracking cognitive load during the N-back as investigated in the previous section. For example, the IPS in the DorAtt network has been implicated in cost-benefit evidence accumulation when costs and benefits are well matched (Basten et al., 2010). Similarly, the dACC, part of the Sal network, is particularly active during selection among competing offers that are close in value (Shenhav, Straccia, Cohen, & Botvinick, 2014). Thus, the contrast of regular and catch trials, for both delay and effort-based decision trials, is investigated in the same five networks shown to encode load in the previous chapter. The significance of this contrast is evaluated with t-tests conducted on the average response function regression weights across networks (or alternatively in particular nodes within those networks). A complementary whole-brain, voxel-wise analysis is provided in the Supplement.

As shown in Table A.1, many of the networks shown to encode load during the N-back also show a reliable contrast with greater activity on regular over catch trials in both delay and effort-based decision trials. The only network not showing a network-averaged effect among those tested is the DMN.

Network	Effort-Based		Delay-Based	
	t-statistic	p-value	t-statistic	p-value
DorAtt	4.81	<0.01	2.16	0.04
CO	4.13	<0.01	1.75	0.09
Sal	2.96	<0.01	1.78	0.09
DMN	-1.19	0.25	-0.22	0.83
FP	3.23	<0.01	2.29	0.03

Table A.1 t-test of hemodynamic response function regression weights on regular versus catch trials in both effort-based and delay-based decision trials. Shading indicates $p < 0.01$ (dark), $p < 0.05$ (medium), and $p < 0.10$ (light).

These results are mostly consistent with prior literature. The IPS, in particular, was previously implicated in integrating costs and benefits during economic decision-making (Basten et al. 2010) in such a way that activity should be higher for decisions in which costs and benefits are more closely balanced. Hence it might also be more active closer to indifference, and thus for regular versus catch trials. In fact, nodes most closely approximating the loci identified by Basten et al. show particularly strong contrasts of regular and catch trials during effort-based decision-making. These are DorAtt node #259, 11 mm from Basten et al. center of mass ($x = -33, y = -46, z = 47; t = 7.17; p < 0.01$) and DorAtt node #260, 5 mm from Basten et al. center of mass ($x = -27, y = -71, z = 37; t = 5.62; p < 0.01$). The dACC has been shown to positively encoding decision difficulty during foraging tasks (Shenhav et al., 2014). The node closest to the peak close-choice contrast in the Shenhav et al. study did not, itself show a reliable contrast between regular and catch trials: FP node #202 ($x = -3, y = 26, z = 44; 9$ mm away; $t = 1.40; p = 0.17$). However, a nearby CO network node #213, showed a robust regular

versus catch trial contrast ($x = -1, y = 15, z = 44$; 18 mm away; $t = 6.30$; $p < 0.01$). This result confirms the hypothesis that part of the dACC support decision-making and the prediction that it is particularly engaged by close offer comparison. Note that node #213 was also the focus of subsequent investigation. In particular, this dACC node was also investigated for encoding subjective value, as described in a subsequent section. Finally, the lateral PFC has been implicated in promoting patient over delayed intertemporal choices in a TMS study, indicating that the lateral PFC also supports careful comparison of costs and benefits (Figner et al., 2010). Two lateral PFC nodes (left IFG $x = -47, y = 11, z = 23$; $t = 4.99$; $p < 0.01$ and right IFG $x = 47, y = 10, z = 33$; $t = 6.62$; $p < 0.01$) also show stronger activity for regular versus catch trials providing confirmatory evidence that these lateral PFC nodes also support close offer comparisons.

The convergence of signals from both domains of delay and effort shown in Table A.1 imply a domain-general set of networks engaged by close offer comparisons. It is important to note, however, that smaller t-statistics indicate weaker effects (though not reliably in any of the networks tested: largest domain contrast $p = 0.17$) on delay relative to effort-based trials. Moreover, a complementary voxel-wise approach, described in Appendix D, shows greatly attenuated signal in the delay relative to the effort-based decision trials. The pattern of results implies that signal-to-noise is relatively diminished in the delay-based decision-making dataset. Diminished signal-to-noise may have to do with relatively diminished engagement in decision-making during the delay trials. This may have to do with the fixed task structure in which delay decision-making always came last. Given that participants would have completed refresher N-back practice for all six N-back levels, 225 effort-based decision trials, and 75 delay-based decision trials before

ever making in-scanner delay-based decisions, it is reasonable to suspect that participants were fatigued and not engaged in effortful decision-making during those trials.

Alternatively, fatigue may have led them to use simplifying heuristics, rendering regular trials more like catch trials. Indeed, response times for regular trial, in-scanner delay-based decisions (median 1.88 s) were reliably faster than regular trial, in-scanner effort-based decisions (median 1.94 s; $p < 0.01$).

Based on the current results, it seems clear that follow-up studies aimed at comparing effort- and delay-based decision-making should employ appropriate counterbalancing. It is worth noting, however, that effort-based decision-making was the focus of this investigation and the fixed decision-task order was an intentional feature designed to specifically enhance the effort-based decision-making results. Because effort-based decision-making is the focus of the current study, subsequent results presented in Appendix A will be exclusively focused on analyses related to this component of the session. However, for the sake of completeness, the delay-based decision-making results are reported in Appendix D.

A.6.2 Regions Encoding Choice and Bias During Decision-Making

In addition to the basic contrast presented above of regular and catch trials, decision-making trials also varied parametrically on a trial-by-trial basis in terms of choice bias. Specifically, the second offer (always for performing the low effort 1-back) was designed to bias participant choice by being either slightly above or slightly below the indifference point with respect to the first offer. Hence, activity pursuant to the decision window may reflect decision biasing. Moreover, because participants commit to a decision during this time, it is also possible to probe for activity encoding the actual

choice (selection of hard versus easy option). Decision window activity encoding both biasing and choice outcome are investigated in this section.

In this analysis, to enhance sensitivity to subtler trial distinctions of bias and choice, and to gain more flexibility to analyze differential responses at particular time points, rather than a mean response across that window, trial epochs were modeled as a series of impulse response functions (also known as “tent” functions in AFNI). Trial epochs are defined as 13 time points from trial onset to 24 seconds after (one time point every $TR = 2$ seconds). Due to the rapid pacing of decision trials (one every 13 seconds), trial analysis epochs thus overlap 2 successive trials. Collinearity of time points are avoided however, by the fact that there is a semi-random order biasing (and corresponding choice) conditions from one trial to the next. Thus, there is trial-wise jittering by trial type that permits resolution of response profiles for one choice / bias condition compared to another.

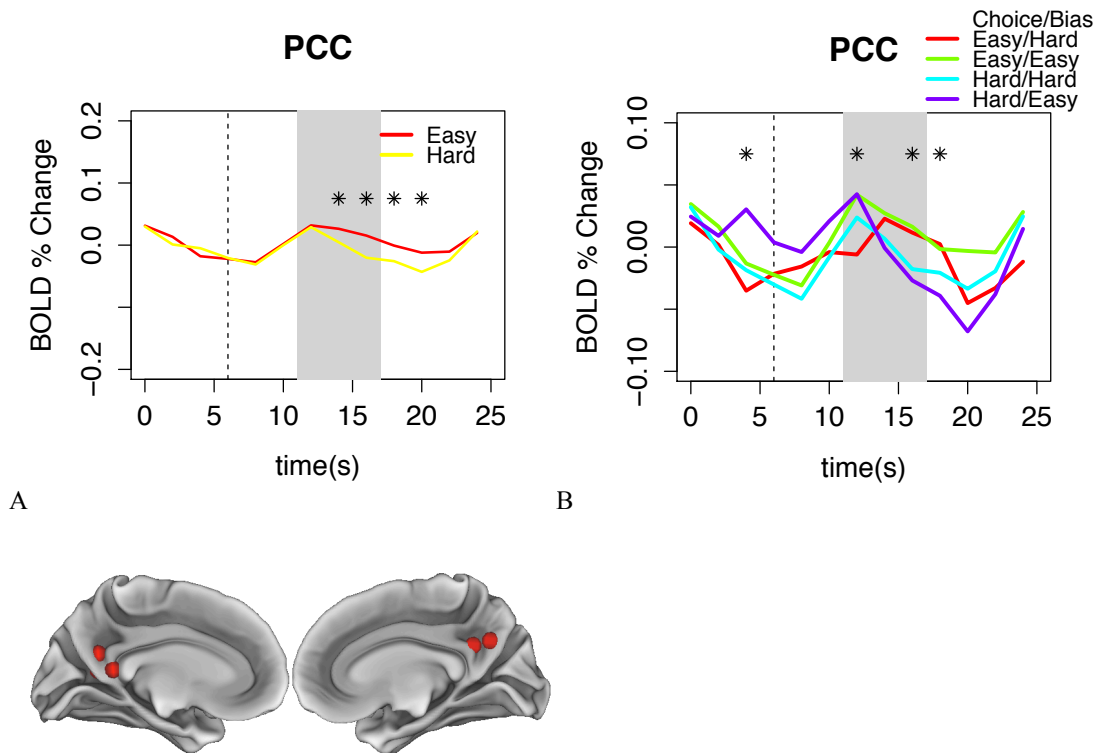
An impulse response model also affords the flexibility of focusing on the time points of interest pursuant to the decision period: 12 to 16 seconds (6—10 seconds after second offer onset). Thus, GLMs were constructed including 13 impulse response regressors for each 2 second interval from 0 to 24 seconds after trial onset (the trial epoch). In the analysis of choice encoding, separate regressors were modeled for hard choices and for easy choices. Then, a series of paired t-tests were used to contrast the time points of interest corresponding to either set. For example, a paired t-test would contrast the modeled response at 12 seconds on a trial in which participants selected the hard choice versus 12 seconds when they selected the easy choice.

Interestingly, a complementary voxel-wise analysis of time points of in the decision window (described in Appendix D) revealed a large cluster in which activity in the PCC that was less active on trials in which participants chose the hard option, relative to those in which they chose the easy option.

Given that the PCC is an a priori region of interest for value encoding (Bartra et al., 2013), one hypothesis explaining the negative activation for a hard over an easy choice is that it reflects a post-decision outcome valuation of greater future costs incurred by the decision (a form of regret). This hypothesis would predict negative activation for hard over easy choices observed in the PCC, which has been otherwise implicated in representing the subjective value of choice outcomes, and in monitoring choice outcomes and the need to alter behavior (Pearson, Heilbronner, Barack, Hayden, & Platt, 2011). Alternatively, rather than reflecting anticipated outcome costs, the negative activation might reflect recognition that the participant made a selection that contradicted their own value function. That is, that they made a mistake and the PCC reflects error magnitude. In that case, the PCC cluster would be most strongly negative in cases when the participant chose against the bias: e.g. they decided against an offer biasing them toward the less costly (easier option) and yet they selected the more costly (harder option).

The caudal portion of PCC reflecting the negative contrast, along with another left lateral cluster also identified for showing a negative contrast at the same time point (see Appendix), map closely to a functionally coupled network that otherwise co-activates with the medial PFC wall, in particular (Leech, Braga, & Sharp, 2012). The set of nodes (Nodes #88—92) overlapping the cluster defined by the contrast of hard versus easy choice outcomes were used to test these hypotheses. Nodes #88—92 were examined

more closely for encoding choice and bias across the decision epoch by extracting their entire time series, after regressing out motion and slow polynomial regressors (to account for scanner drift), and averaging over all trials (and all 5 nodes) according to whether the trial was one that the participant selected the easy or hard option (Figure A.12A), and whether the second (easy option) offer biased them towards the hard or easy option (Figure A.12B).



C

Figure A.12 Averaged time series in the PCC (Nodes #88—92) for trials in which the participant selected the hard or easy option, and whether they were biased towards the hard or easier option. The grey shaded region indicates the time points of greatest interest for the decision period: 6—10 seconds after decision period onset. A) Time series are averaged according to whether the hard or easy option was selected, and * indicates a reliable difference between hard and easy outcomes at $p < 0.05$. B) Time series are averaged across all trials based on both biasing and choice and * indicates a reliable difference between hard choice / easy bias (purple) trials and easy choice / hard bias (red) trials at $p < 0.05$. C) Shows the location of the caudal PCC nodes.

As shown in Figure A.12A, activation peaks in the PCC at 12 seconds and is well matched whether participants select the hard or easy option. After this point, the function

drops reliably faster for the trials in which the hard option was selected. This result suggests that the PCC is indeed processing post-decision outcomes and encoding the greater costs of the harder option outcome. When trials are averaged according to both choice and offer biasing, there is no time point in which the trials on which participants violated their own assumed value function (red and purple trials) are distinct from when they followed their value function (cyan and green). Together, these results suggest that the PCC is not encoding “errors” in the sense of choice outcomes that do not align with internal value functions. Instead, they support that the PCC is encoding anticipated costs after a choice has been committed. Interestingly, there is reliably greater activity at response peak (12 seconds; Figure A.12B) for trials in which participants selected the hard trial, but were biased by the easy offer relative to trials in which they choose the easy offer but were biased towards hard (purple > red trials), and this is also true just prior to the valuation window. That is, activity is higher at response peak when participants select the harder option, and particularly when must overcome an offer bias not to. This result suggests that the PCC also carries information about intrinsic motivation in decision window activity.

Note that Nodes #88—92 were all slightly posterior to loci identified by Bartra et al. (2013) for positively encoding SV during offer valuation and decision outcome receipt (24-33 mm away). However, a node approximating the Bartra loci (Node # 133 – discussed subsequently for SV encoding) did not show reliable choice or bias effects during the decision-window and so is not discussed here. Beyond the Bartra loci, vmPFC loci have also been otherwise implicated in encoding SV during decision-making, and do

show effects of choice and bias during the decision window. These are discussed in the next section.

A.6.3 Regions Encoding Value Dimensions

While examination of the decision window revealed a network of regions involved in cost-benefit decision-making over and above perceptuo-motor processes, and regions (particularly the PCC) encoding choice outcome costs, investigation of the valuation window, in which a single (high demand) option is considered in isolation, permits unambiguous resolution of brain regions encoding key decision-making cost and benefit variables. Note that regions encoding choice variables may not vary in mean signal between regular and catch trials, but nevertheless encode dimensions relevant to choice in both regular and catch trials.

A key choice dimension is SV. In the human neuroeconomics literature, a standard approach to identifying regions encoding SV is to conduct a parametric analysis of regions tracking trial-by-trial variations in SV, e.g. (Kable & Glimcher, 2007; Peters & Buchel, 2010; Pine, Shiner, Seymour, & Dolan, 2010). Here, SV is defined as a quantity that maps the objective choice dimensions (e.g. amount, delay, risk, and effort) into a subjective (individual-specific) quantity that describes each individual's decision-making behavior. This section starts by testing for the encoding of the SV of the first (pure valuation) offer by how the amplitude of the hemodynamic response is modulated by the SV of the first offer on every trial. In the current study, SV is defined according to the participants' indifference points for each amount-load pairing (relative to the 1-back offer), taken from the (out-of-scanner) sessions.

As before, to avoid constraining restrictions on the shape of the response function, and maximize sensitivity to trial-by-trial variations in SV, valuation (and decision) processes are modeled using the impulse response “tent” functions, which provides an estimate of signal at each time point in a trial epoch. Unassumed impulse response models are consistent with the approach taken elsewhere in the neuroeconomic literature, e.g. (Kable & Glimcher, 2007). From a practical standpoint, they are also the best choice given the design of the present decision-making paradigm. Rapid pacing and fixed trial onsets makes canonical hemodynamic response functions a poor choice for modeling valuation as a distinct event (valuation periods are a fixed, 6 second interval, followed immediately by the decision window). Though variable response times provide some natural jittering between the offset of the decision window, and the onset of the following valuation window, there is considerable overlap between valuation and decision processes. Namely, the trailing edge of hemodynamics pursuant to response execution overlaps ramping activity pursuant to valuation processes. Hence, mean variation in valuation processes, as modeled by a canonical HRF, can be obscured by decreasing activity from response execution processes on the previous trial. The tent function used to model each trial covers 24 s trial epochs (20 s beyond the decision window onset); thus, trials were modeled as 13 time points, spaced at 2 seconds each.

Despite the fact that trial epochs overlap multiple trials (2, to be precise), randomized trial ordering effectively jitters trial types (as with choice and bias conditions) enabling clear resolution of the degree to which tent functions are modulated by key decision variables. Task-amount pairs are assigned random order, hence key variables like subjective value (SV), amount, or cost (delay or N-back load), vary trial-to-

trial in a pseudo-random fashion. Thus, when modeling trial epochs as parametrically varying in amplitude by key variables, time points are sufficiently jittered for precise estimation of that amplitude modulation.

Figure A.13A shows the mean response function for an example set of DMN nodes used in the previous N-back analysis, while Figure A.13B shows how the amplitude of the hemodynamic response function is modulated by SV at each time point for these nodes. Note that SV is defined by participant-specific discounting functions in the third session, for all task-amount pairs. For example, if a participant is found to be indifferent between an offer of \$2 for the 2-back and \$1.73 for the 1-back, then the raw (non-normalized) SV of the offer of \$2 for the 2-back, for that participant, is \$1.73. Raw SV amounts were thus individually-defined and mean-centered for each participant, across all 15 reward amount – task load pairs (i.e., 3 reward amount x 5 load levels) to form the parametric predictor values.

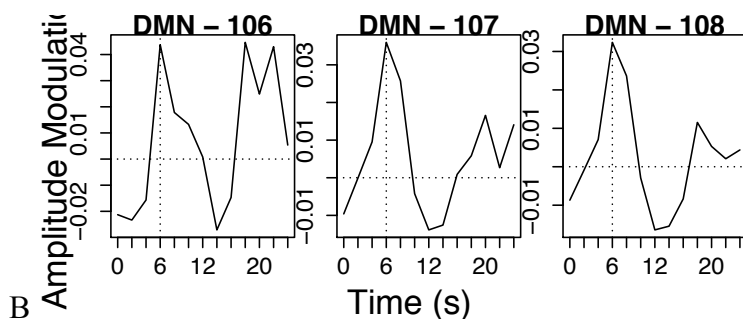



Figure A.13 A) Mean and B) SV amplitude modulation functions for decision trial epoch in three example DMN nodes, estimated as a set of impulse responses spaced every 2.0 s TR. Solid lines provide the response functions for effort-based, and dashed lines the response functions for delay-based decision trials. Valuation windows begin at 0 s, and conclude with the start of the decision window at 6 s.

As can be clearly seen from the three example nodes in Fig A.13A, the mean response is dominated, in the decision epoch, by response-related processes, peaking ~16-18 seconds after the trial has begun. This dominant function is also reflected in the trailing edge that lapses into the valuation window (0-6 seconds) in the three example nodes. Not all nodes show this kind of function. Indeed, mean response functions take a range of shapes across nodes and networks. All nodes are provided in the Supplement, for reference.

Critically, however, while the mean function reflects the trailing edge of the response-related peak through the valuation window, there are very prominent peaks of amplitude modulation by SV pursuant to the valuation window (6-8 seconds) in all three nodes. Hence, the pseudo-randomized trial order (with respect to task-amount pairs) permits clean resolution of amplitude modulation effects apart from the mean response function. It is important to note that these amplitude modulation effects reflect the response to pure valuation of a single offer, accounting for hemodynamic lag, and are unconfounded by the second offer onset since they peak at the same time of the second offer onset (and immediately after).

A number of a priori loci have been identified for SV encoding during economic decision-making. In the next section these nodes are tested to determine whether they also encoded SV during decision-making about cognitive effort. Although the main analysis is focused on a priori loci, a complementary whole brain analysis is provided in the Supplement.

A.6.3.1 A Priori Nodes Encoding SV

Two recent meta-analyses implicate a set of nodes that belong to a core valuation network: first, Levy and Glimcher (2012) report two key nodes in the vmPFC identified through a meta-analysis of consisting of 13 valuation studies of primary and monetary rewards. Second, Barta, McGuire, and Kable (2013), report a wider network of 9 regions, including the vmPFC, but also the ACC, the dACC, the PCC, the brainstem, and the striatum, from a broader meta-analysis of general SV encoding, that included studies of hypothetical and real, primary and monetary rewards that were delayed, probabilistic, effortful, or in the certain punishment domain. Additionally, they suggest that the dACC and anterior insula (AI) may most strongly encode the negative, or cost dimension, of SV. While these meta-analyses have focused on general SV encoding, only one study to date has investigated the encoding of SV of cognitively effortful rewards. Massar et al (2015) report 7 regions that may specifically encode SV during valuation of rewards contingent on cognitive effort in particular. The locations of a priori nodes of interest are shown in Figure A.14, color-coded by study of origin. Note that some a priori nodes, namely the VS nodes, were defined independently of the 264 node set used elsewhere in this dissertation. This is critical given that the 264 node set does not map well onto subcortical structures like the VS.




Figure A.14 A priori nodes of interest for encoding SV. Nodes are color-coded by their study of origin. Green and orange nodes are from the meta-analyses of Levy and Glimcher (2012) and Bartra, McGuire, and Kable (2013), respectively. Blue nodes are from the study of cognitive effort by Massar et al. (2015).

To investigate all SV encoding in all a priori nodes, 12 mm nodes centered on all loci of interest were tested for significant amplitude modulation 6 to 8 seconds after trial (valuation period) onset. This window was selected because it corresponds to the onset of the decision period, and should reflect pure valuation of the first offer alone (always the harder task for more money), and also because it is consistent with peak SV encoding in prior studies (Kable & Glimcher, 2007). Table A.2 shows the results of t-tests for reliable amplitude modulation by first offer SV in all 18 a priori nodes.

Anatomical Description	(x,y,z)	t-stat	p-value
Levy and Glimcher (2012)			
282 – l vmPFC	(4,35,-12)	4.65	<0.01
281 – r vmPFC	(-7,38,-11)	4.04	<0.01
Bartra, McGuire, and Kable (2013)			
265 – l striatum	(-12,12,-6)	3.53	<0.01
266 – r striatum	(12,10,-6)	2.85	<0.01
267 – vmPFC	(2,46,-8)	4.55	<0.01
268 – l AI	(-30,22,-6)	2.30	0.03
269 – r AI	(32,30,-6)	2.44	0.02
133 – PCC	(-2,-35,31)	3.05	<0.01
271 – Brainstem	(-2,-22,-12)	1.78	0.09
215 – ACC	(0,30,27)	4.15	<0.01
213 – dACC/pre-SMA	(-1,15,44)	3.06	<0.01
Massar, Libedinsky, Weiyan, Huettel, and Chee (2015)			
274 – r supramarg. gyr.	(33,-52,32)	0.33	0.74
275 – l cingulate	(-24,-49,36)	0.97	0.34
276 – l inf. temp. gyr.	(-58,-35,-22)	3.10	<0.01
277 – l IFG	(-43,53,-4)	1.48	0.15
278 – l IPL	(-30,-43,43)	1.49	0.14
279 – l IPL	(-41,-55,46)	0.34	0.74
280 – l postcentral gyr.	(34,16,-26)	3.21	<0.01

Table A.2 t-tests for reliable, trial-wise, parametric amplitude modulation by first offer SV in a priori nodes. Shading indicates $p < 0.01$ (dark), $p < 0.05$ (medium), and $p < 0.10$ (light).

SV is reliably encoded as positive amplitude modulation in all a priori nodes identified for general SV encoding from the two meta-analyses. Most of the nodes reported in the single cognitive effort study (Massar et al., 2015), by contrast, did not show reliable trial-wise amplitude modulation by SV, with the exception of Node 276 and 280 in the inferior temporal gyrus and left post-central gyrus, respectively.

A.6.3.2 Amount and Load Encoding

For a region to encode SV, it must reflect both benefit and cost magnitude. In our effort-based decision trials, that corresponds to encoding both the offer amount and the task load in the amount-load pair. It is possible, however, that a priori nodes reflect one dimension or another, only. To determine whether the nodes indeed encode both, time series were extracted from each node. As before, motion effects were regressed out first

prior to analysis. The resulting node-averaged time series were then separated into trials and aggregated by trial type: that is, by amount and by N-back task load.

A series of regressions were conducted to test whether these 18 a priori nodes encode SV of cognitively effortful rewards, in particular, and the independent dimensions of task load and reward amount. Two types of regressions were conducted for each node, each predicting activity averaged across time points 6 and 8 seconds: 1) the first has predictors of N-back load and offer amount (of the first offer) centered and 2) the second has the ratio of load / amount centered. Alternately separating out and combining the two key dimensions of choice allows for tests of whether putative SV nodes show a combined representation of dimensions, and also whether they show independent representation of dimensions, respectively. Again, variable intercept models are used to account for the nesting of trials within participants.

$$\text{BOLD}_{\text{lag}3-4,i} = B1_{j[i]} + B2_i (\text{Amount} / \text{Load}) + \varepsilon_i \quad (\text{A.1})$$

or

$$\text{BOLD}_{\text{lag}3-4,i} = B1_{j[i]} + B3_i \text{Amount} + B4_i \text{Load} + \varepsilon_i \quad (\text{A.2})$$

$$B1_j = \gamma_{1,0j} + \eta_j \quad (\text{A.3})$$

As shown in Table A.3, every node identified for positive SV encoding, in the two meta-analyses, either significantly or at trend-level, encodes the amount-to-load ratio positively. Also, all of the core valuation nodes in the vmPFC and VS show independent encoding of both amount and load, as expected. This result constitutes a critical and novel finding, as it is the first time that the core valuation network has been shown to encode cognitive load as a cost dimension during offer valuation. It also supports the

hypothesized cognitive effort discounting patterns observed in the VS during reward receipt in a prior study (Botvinick et al., 2009).

Anatomical Description	Amount B3×10 ⁻² (<i>p</i> -value)	Load B4×10 ⁻² (<i>p</i> -value)	Amt. / Load B2×10 ⁻² (<i>p</i> -value)
Levy and Glimcher (2012)			
282 – l vmPFC	2.42 (<0.01)	-1.40 (<0.01)	1.83 (<0.01)
281 – r vmPFC	2.06 (<0.01)	-1.53 (<0.01)	1.82 (<0.01)
Bartra, McGuire, and Kable (2013)			
265 – l striatum	2.83 (<0.01)	-1.19 (<0.01)	1.67 (<0.01)
266 – r striatum	3.06 (<0.01)	-1.12 (0.02)	1.82 (0.01)
267 – vmPFC	2.76 (0.01)	-2.12 (<0.01)	2.44 (<0.01)
268 – l AI	0.43 (0.39)	-0.81 (<0.01)	0.79 (0.06)
269 – r AI	0.46 (0.26)	-0.56 (0.02)	0.58 (0.08)
133 – PCC	1.02 (0.29)	-1.46 (<0.01)	2.28 (0.02)
271 – Brainstem	0.05 (0.91)	-1.08 (0.05)	1.91 (0.10)
215 – ACC	4.42 (0.04)	-1.39 (0.01)	1.84 (0.02)
213 – dACC/pre-SMA	0.02 (0.35)	-1.56 (<0.01)	1.54 (0.06)
Massar, Libedinsky, Weiyan, Huettel, and Chee (2015)			
274 – r supramarg. gyr.	0.05 (0.91)	-0.47 (0.07)	0.63 (0.10)
275 – l cingulate	0.10 (0.78)	-0.27 (0.17)	0.44 (0.13)
276 – l inf. temp. gyr.	0.72 (0.17)	-1.10 (<0.01)	1.25 (<0.01)
277 – l IFG	0.57 (0.55)	-1.23 (0.03)	1.62 (0.04)
278 – l IPL	0.31 (0.54)	-0.49 (0.10)	0.58 (0.19)
279 – l IPL	0.16 (0.85)	-0.72 (0.15)	0.82 (0.26)
280 – l postcentral gyr.	1.94 (0.04)	-0.93 (0.08)	0.87 (0.25)

Table A.3 Amount, load, and amount/load ratio as predictors of activity 6 and 8 seconds after trial onset, and also N-back activity as a predictor of SV in a priori SV nodes. Shading reflects $p < 0.01$ (dark), $p < 0.05$ (medium), and $p < 0.10$ (light).

Interestingly, Table A.3 also identifies a number of nodes, like the PCC, left AI and dACC, that reliably encode load but not amount. This result suggests that SV, *per se*, is not encoded in the PCC, AI or dACC, but instead that cost information is selectively encoded, even though cost correlates with SV. By contrast, both the striatum and vmPFC nodes show reliable encoding of both dimensions. Also, while these results largely recapitulate the effects observed in the meta-analyses, they largely fail to replicate encoding of SV in loci identified in the Massar et al. (2015) study, with a few exceptions. Namely, the left IFG and inferior temporal gyrus both show encoding of load, and,

thereby the encoding of (correlated) SV. However, none of the loci encode both dimensions of load and amount.

A visual representation of the independent encoding of both load and amount dimensions in 4 vmPFC and VS nodes is presented in Figure A.15. As before, these figures were made by averaging error time series, after regressing out motion and slow polynomial regressors, across all trials of the same first offer amount or load. Also, linear effects of load and amount are tested at each time point across all trials. Note the temporal specificity of linear load and amount effects are restricted primarily to the time window of interest – 6 to 8 seconds after valuation period onset.

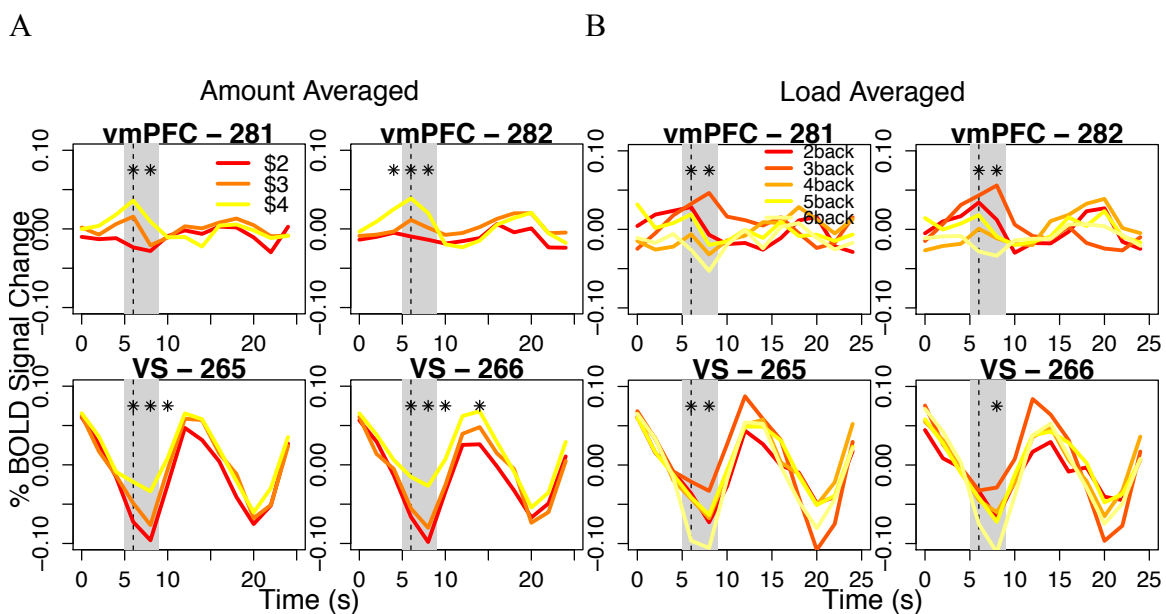


Figure A.15 Timecourse plots showing effects of A) amount, and B) load on activity in key vmPFC and VS nodes. Grey shading indicates the valuation window (6—8 sec after valuation period onset). * Indicates a reliable linear effect (at $p < 0.05$) of amount or load at each time point, according to variable-intercept, multi-level models with trials nested within individuals.

As mentioned above, a number of nodes encode load (negatively) during valuation and hence their trial-wise amplitude modulation would correlate with trial-wise first offer SV, but on closer analysis were found to not also encode amount. Notably, this includes the dACC node of interest. This result is notable in part because the dACC has

been hypothesized to select the value of cognitive control task sets based on their expected value (benefits minus costs) (Shenhav et al., 2013), and the present data only supports the cost side of the equation. Of course there are many reasons why the dACC might not report benefits and costs at the level of fMRI BOLD signal. A single unit recording study of monkey ACC neurons has shown simultaneous positive and negative encoding of offer value (which may thus cancel out at the level of local field potentials) (Kennerley et al., 2011). Thus these data do not rule out benefit encoding in the dACC, but instead do provide support for cost encoding during valuation.

The encoding of load in the dACC during valuation is also interesting because the same node also showed robust contrast of regular versus catch trials. As shown in Figure A.16, averaged trial epoch timecourses from this cluster aggregated, alternately, by trial type or by anticipated cognitive load, show greater decision-window activity for regular versus catch trials and greater valuation-window activity for lower versus higher load. That is, it simultaneously encodes decision load positively and anticipated N-back load negatively. Note that while this same node shows reliable effects of load during offer valuation, it does not, as indicated above, show reliable effects of reward amount at any time point.




Figure A.16 Decision-window and valuation-window activity in the a priori dACC node also identified by the contrast of catch and regular decision trials. A) Time series averaged by catch and regular decision trials. B) Time series averaged by anticipated load of first offer. C) Location of dACC node. * Indicates a reliable linear effect, at $p < 0.05$, of load on activity at the corresponding time point in a multi-level model with trials nested within participants.

A.6.3.3 A Priori SV Nodes Predict Choice

An important question is whether SV encoding is determinative of choice, or merely correlative. While the vmPFC and VS regions clearly encode both choice dimensions pursuant to valuation of the first offer, a further question is whether activity in a priori regions is also predictive of the subsequent choice that the participant will make, regardless of the second offer. This question was examined by coding trials according to two factors: 1) the subsequent choice made on that trial, the high-load, harder option (the option presented during the valuation period) or the low-load, easier option (presented during the decision period); and 2) the bias induced on that trial (to choose harder or easier option), based on whether the low-load offer is below or above

the estimated indifference point relative to the high-load offer. A priori vmPFC nodes, in particular, show a reliable increase in activity, pursuant to first offer valuation on trials in which the participant ultimately chooses the harder option when the trial is designed to bias them towards the easier option. In other words, valuation-related activity in the vmPFC anticipates whether a participant will elect the high demand offer, and even more so when that means overcoming a bias not to (similar to what was observed for the PCC cluster during the decision window). The mean response across choice and bias conditions in the three a priori vmPFC nodes (267, 281, and 282) is shown in Figure A.17 below.




Figure A.17 Averaged time series in three vmPFC nodes (267, 281, and 282) encoding SV, for trials grouped by whether participants chose the hard (first) or easy (second) offer and also by whether the easy offer was designed to bias participants to choose it or the hard offer. Green and orange nodes are from the meta-analyses of Levy and Glimcher (2012) and Bartra, McGuire, and Kable (2013), respectively. * Indicates a difference (at $p < 0.05$) by time point of hard choice / easy bias (purple) trials, and easy choice / hard bias (red) trials.

Because only one option is available, and hence participants cannot make a choice during the valuation window, the fact that valuation-related activity anticipates subsequent choice implicates the vmPFC in calculating a pre-decision quantity of SV, rather than passively reflecting the value of the chosen offer. It is also important to note

that because this analysis collapses across all amount-load pairs, the orderly result implies that the extent to which participants value a given task-amount pair varies somewhat independently across trials, and this independent variation plays a role in subsequent choice. One interpretation of this pattern is that on some trials the participant may encode the high-effort option as being more subjectively valuable than even their own mean SV for that amount / load combination, based on endogenous or stochastic factors. On those trials, the data suggest that the participant will be more likely to select that high-effort option, presumably reflecting its higher relative valuation compared to the low-effort option. As described above, this effect mirrors the distinction made in the decision-window (at 12 seconds and at 4 seconds, prior to the valuation window) in the caudal PCC, implying that the intrinsic motivational state information is shared between the two valuation regions. Also, notably, both the caudal PCC and the vmPFC nodes appear to show significantly lower activity in the decision window on trials in which the hard task was ultimately selected.

Note that the valuation-window effect of greater activity for hard choice / easy bias trials (purple; cf. Figure A.17) relative to easy choice / hard bias trials (red) is also observed during the valuation window in four of the nodes encoding load during valuation. At 6 seconds, there is a reliably greater activity in hard choice / easy bias trials in the left AI (Node #268), and at 4 seconds there is reliably greater activity in the PCC (Node #133), the right supramarginal gyrus (Node #274), and (at trend-level $p = 0.07$) in the dACC (Node #213). These results implicate these nodes in encoding both load and intrinsic motivation to engage with a demanding task, along with the vmPFC and caudal PCC.

A.6.3.4 Subjectivity in SV Encoding

Beyond reward amount and load, a third dimension of SV is subjectivity itself. Subjectivity refers to the result that there are idiosyncratic preferences (i.e., high vs. low effort discounting) by which individuals vary in their valuation of an offer over and above the particular amount and load combination presented on that trial. In general, participants who display steep effort discounting must find either the offered rewards to be less valuable, the cognitive loads to be more subjectively costly, or both. The net result is quantified as reduced SV and AUC. In this section, a priori nodes are tested for evidence of subjectivity.

One way to investigate subjectivity, given that mean activity is positively modulated by SV in the a priori nodes, is to test whether mean activity in the SV nodes varies by individual differences in discounting. A series of t-tests reveals that in none of the a priori nodes encoding SV does mean activity across 6 and 8 seconds vary reliably by AUC_{3S} (all p 's ≥ 0.22). However, the pattern of activity in all three vmPFC nodes reveals a clear ordinal pattern. As shown in Figure A.18, the mean time course deflection is more positive for shallow ($AUC_{3S} > 0.8$) versus medium ($0.6 < AUC_{3S} < 0.8$) versus steep ($AUC_{3S} < 0.6$) discounters, across all three vmPFC nodes. Moreover, this pattern corresponds to a reliable linear effect of AUC_{3S} on activity at a later time point (12 sec after decision period onset) whereby shallower effort discounters have a more positive deflection. Although the later time point reflects activity pursuant to decision-making rather than pure valuation, the sum of the evidence strongly suggests that subjectivity is playing a role in value encoding in the vmPFC.




Figure A.18 Averaged time series from the three a priori vmPFC nodes (267, 281, and 282) pursuant to the valuation period on effort-based trials, for steep, medium, and shallow discounters. Participants are grouped according to their AUC_{3S} . Green and orange nodes are from the meta-analyses of Levy and Glimcher (2012) and Bartra, McGuire, and Kable (2013), respectively. * Indicates a reliable linear effect, at $p < 0.05$, of AUC_{3S} on mean activity at a given time point.

A similar, but more rigorous analysis asks whether, controlling for objective factors of load or amount, subjectivity, as captured by AUC_{3S} , explains additional variance in activity in the vmPFC cluster. If AUC_{3S} were to explain variance beyond that explained by task and load, it would constitute strong evidence for the encoding not just of objective, but *subjective* value in this region.

A multi-level, multiple regression provides a formal test, in which trial-wise average regression weights (across time points 6 and 8 seconds) are explained by fixed effect predictors of task load and amount and, simultaneously, with random effects of participant-level AUC_{3S} as a subject-level predictor. Here the question is whether there are subject-level effects of (average) discounting beyond those fixed effects of amount and load on activity in the a priori nodes.

$$BOLD_i = B1_{j[i]} + B2_i \text{ Amount} + B3_i \text{ Load} + \varepsilon_i \quad (\text{A.4})$$

$$B1_j = \gamma_{1,0j} + \gamma_{1,1j} AUC_{3Sj} + \eta_j \quad (\text{A.5})$$

While amount and load show effects reflecting those provided above in Table A.2, AUC_{3S} explains no additional variance in mean signal (all p 's ≥ 0.24), controlling for these variables. Though this analysis does not support the full hypothesis of subjective value coding in a priori nodes, the overall small sample size and restricted variance in discounting constitute a low-power, insensitive dataset for addressing the question of subjectivity satisfactorily. A rigorous test for subjective coding beyond objective choice dimensions during effort-based decision-making requires a larger sample with greater variance in effort discounting.

While the previous analysis does not support the full subjectivity hypothesis, another way to test for subjectivity and the relationship with amount and load encoding is to test whether inter-individual amount and load effects, during offer valuation, vary by discounting themselves. To test this, a series of linear models of effects of amount and load on valuation period activity, similar to Eqns. A.2 and A.3, were fit for each subject and a priori node separately. Next, models were fit to test whether AUC_{3S} predicts individual differences in these subject and node-specific slope terms. Interestingly, while AUC_{3S} does not predict load effects in any nodes (all p 's ≥ 0.14), AUC_{3S} does predict the slope of the amount effects in the left VS (Node 265: $B = 0.15$, $p = 0.02$). Note that the relationship holds when excluding the apparent outlier with the very high amount effect slope (excluding the participant: $B = 0.10$, $p = 0.01$). The positive relationship indicates that in shallow effort discounters there is a more positive effect of offer amount on response functions in the VS during valuation. These results support subjective variation

in the response to amount stimuli, if not load stimuli.




Figure A.19 Amount effects in the left VS (Node 265) during valuation positively predicted by participant AUC_{3S}. Inset provides coronal slice with left and right VS nodes in red. Green crosshairs are centered at the left VS node location (radiological convention; MNI x = -12, y = 12, z = 6).

A.6.4 Integrating Effort-Tracking Information into Valuation

A central question of this research is to investigate whether there are brain regions involved in tracking cognitive effort during task engagement that communicate this information to valuation regions at the time of effort-based decision-making. As revealed in the main text, a number of regions show features that make them good candidates for tracking effort for this purpose. Task-positive networks, in particular, showed negative relationships between mean, load-independent signal and discounting. The task-negative DMN did not show reliable individual difference relationships with discounting. However, it did show approximately monotonic decline in activity with load, thus mirroring robust within-individual variation in SV. Hence the pattern of activity in these networks during the N-back may anticipate subsequent dynamics in the valuation period in a priori SV nodes. In this section, the relationship between patterns of activity in the N-back are tested for predicting patterns of activity in a priori SV nodes during offer valuation.

A.6.4.1 Does load-independent N-back activity predict valuation activity?

A first test examines whether individual differences in mean, load-independent N-back activity in the networks of interest (CO, FP, Sal, DorAtt, and DMN) predicts valuation period activity in the a priori SV nodes. Specifically, subject-averaged regression weights (across all loads) were used as a measure of cognitive effort during the N-back, and these subject-averaged regression weights were tested for predicting valuation period activity (6—8 seconds after valuation period onset). In only one a priori SV node (identified by Massar et al.) does load-independent N-back activity in task-positive networks predict (at $p < 0.05$) valuation period activity (Node #279 - l IPL). Activity in this node is predicted positively by mean N-back activity in the FP ($B = 0.90$; $p = 0.03$) and Sal networks ($B = 0.75$; $p = 0.05$). Incidentally, this node is virtually overlapping (Euclidean distance = 2 mm from) a FP node. Thus, the finding indicates that higher mean activity during the N-back in the FP network appears to predict higher valuation period activity in a specific FP node (the left IPL) during valuation. This may indicate a key node for effort tracking, but the sign of the relationship complicates this hypothesis since higher activity during valuation should correspond to higher SV, while higher load-independent activity in the N-back corresponds to lower SV. Moreover, Node 279 did not show reliable relationships with load, amount, or SV during valuation. Thus the evidence for this region caching load information for the purposes of offer valuation is not particularly strong.


By contrast with the task-positive networks, individual differences in load-averaged N-back activity in the DMN appears to reliably correlate with averaged valuation period activity in multiple a priori SV nodes. Table A.4 provides a list of all of the nodes where this relationship is significant ($p < 0.05$) or trending ($p < 0.10$).


Intriguingly, these nodes are primarily overlapping the set of nodes showing load (if not amount) effects during the valuation period. That is, these nodes show negative encoding of load during valuation, and also are positively predicted by load-independent activation in the DMN during the N-back.

Anatomical Description	B (<i>p</i> -value)
Bartra, McGuire, and Kable (2013)	
268 – l AI	0.66 (0.04)
269 – r AI	0.49 (0.04)
271 – Brainstem	0.81 (0.06)
215 – ACC	0.90 (0.06)
213 – dACC/pre-SMA	1.18 (0.06)
Massar, Libedinsky, Weiyang, Huettel, and Chee (2015)	
277 – l IFG	0.63 (0.10)

Table A.4 A priori SV nodes in which valuation period activity is predicted by average N-back β weights in the DMN. Shading reflects $p < 0.01$ (dark), $p < 0.05$ (medium), and $p < 0.10$ (light).

The monotonically decreasing load function in the DMN makes it, in particular, a strong candidate for predicting load effects during valuation. The sign of the effects in Table A.4 indicate that more deactivation in the DMN during the N-back predicts more deactivation in these nodes during offer valuation.





C

Figure A.20 DMN N-back activity predicts valuation period activity in two AI nodes. A) Node 268 – l AI and B) Node 269 – r AI. C) Left and right node locations. Dashed lines give 95% C.I.

A positive relationship between DMN activity during N-back performance, and SV node activity during offer valuation is sensible, given that the DMN during the N-back and the SV nodes during offer valuation both negatively encode load. The SV node-DMN relationship results from the fact that individuals with less N-back deactivation in the DMN also find the N-back less costly, and hence they show greater valuation period activity than those with more N-back deactivation. This interpretation implicates both DMN N-back activity and load encoding of load in these SV nodes during offer valuation in the subjective effort cost encoding. It is also possible that the correlation has nothing to do with valuation. An alternative interpretation implicates a connection between the extent to which individuals' Sal networks (largely overlapping those nodes showing a relationship with DMN N-back activity) deactivate during offer valuation, and the way that their DMN responded to working memory tasks (systematic differences in session signal-to-noise, e.g.). Under the alternative interpretation, the fact that both the DMN and these SV nodes encode load negatively, during the N-back and valuation respectively, is incidental (i.e., it would be expected in any task context, not just during N-back and valuation).

It is important to note that the prior result linking individual differences in DMN N-back activity to valuation period activity is also complicated by the fact that the DMN

was the only network *not* showing an individual difference relationship with discounting. Thus, it is not clear why individual differences in activity should be predictive of valuation activity but not explicit valuation. One potential explanation has to do with the nature of the link between DMN deactivation and subjective effort. It is possible, for example, that the DMN needs to be deactivated to a fixed degree to perform an effortful task appropriately, at a given level of load (i.e., attenuating between-subjects variability), but that some participants find it more subjectively costly to do so. Hence, individual differences in discounting may not be directly reflected in individual differences in N-back DMN deactivation, even if valuation period activity during valuation is related to both.

A.6.4.2 Do DMN load effects during the N-back predict load effects during valuation?

Another analysis that could provide strong evidence implicating DMN deactivation in valuation-related activity is one that focuses directly on the effects of load in both contexts. Given that both SV and DMN N-back activity exhibit negative linear load effects, individual differences in the load-related linear slope of DMN deactivation might also predict linear load slope effects in a priori nodes during valuation.

Importantly, this constitutes a stronger test of integration than the previous one, given that the previous one focused on load-independent (i.e., mean activity) individual differences, whereas this one focuses on *load-dependent* (i.e., the steepness of load slope effects). To test this, separate models for each participant were fit to describe the linear effect of load (i.e., load slope) on DMN N-back activity, and also in the a priori nodes, the linear effect of load slope on valuation period activity (6—8 seconds after trial onset). Next, the N-back load slope effects were used in a model predicting valuation period load

effects. The results, shown in Table A.5, indicate a positive relationship in multiple a priori SV nodes. Critically, this analysis constitutes a stronger test of integration than the prior one, in that the previous analysis focused only on individual differences load-independent (i.e., mean) activity in the DMN, whereas this one focused on *load-dependent* activity – linking individual differences in the way load affects DMN deactivation during N-back task engagement to the way load is encoded when prospectively considering load during valuation. The prior result may be explained by a more uninteresting link between DMN deactivation during both N-back and decision-making tasks. Indeed, vmPFC nodes of interest are part of the DMN itself and some participants may just have stronger DMN deactivation responses across diverse kinds of tasks. However, the results presented in Table A.5 avoid that possibility since it links the result specifically to how a given participant’s a priori SV nodes encode prospective load.

Anatomical Description	B (<i>p</i> -value)
Levy and Glimcher (2012)	
281 – r vmPFC	0.95 (0.03)
Bartra, McGuire, and Kable (2013)	
133 – PCC	1.71 (0.03)
271 – Brainstem	1.42 (0.03)
215 – ACC	1.22 (0.07)
Massar, Libedinsky, Weiyan, Huettel, and Chee (2015)	
279 – l IPL	1.32 (0.09)

Table A.5 A priori SV nodes for which N-back load effects in the DMN predict load effects during the valuation period. Shading reflects $p < 0.01$ (dark), $p < 0.05$ (medium), and $p < 0.10$ (light).

The positive sign in these relationships indicates that individuals with steeper load slope effects in the DMN during the N-back task also have steeper load slope effects in this network of a priori SV nodes during valuation. This result supports the hypothesis that the DMN tracks subjective costs as a function of load during task performance. It also supports the hypothesis that the cost information is integrated into valuation processes via a priori valuation nodes (and the vmPFC, PCC, brainstem, ACC, and IPL in

particular). Finally, it constitutes additional evidence for subjective encoding of effort costs, given that the relationship is based on individual differences. Thus, the finding supports the intriguing hypothesis that deactivation of the DMN tracks subjective cognitive effort and that this deactivation informs effort anticipation during valuation and decision-making.

A.6.4.3. Do nodes encoding load during valuation also encode load during the N-back?

A final link between effort cost valuation and effort tracking during task engagement may be found in the a priori nodes themselves. If nodes encoding load during valuation also encode effort costs during N-back performance, it would support the hypothesis that these valuation nodes themselves cache effort information (i.e., extracted during N-back performance) for subsequent decision-making. To investigate this hypothesis, activity in the a priori nodes shown to encode load during valuation is tested for encoding discounting (SV_{3S}) during the N-back. Formally, the same modeling approach used in the main text (Eqns. 4.5—4.6), is used again to test the effects of participant-centered, load-dependent effects of brain activity (β_{ctr}) and also load-independent (β_{avg}) effects, during the N-back, on subsequent discounting in each of the a priori SV nodes. Mirroring the results of the main text chapter, none of the a priori SV nodes showed load-dependent effects, while many of them showed load-independent effects. Those effects and their p-values are included in Table A.6.

Anatomical Description	Activity effects *10 ⁻¹ (p-value)	
	Load-specific B3	Load-independent γ_1
Levy and Glimcher (2012)		
282 – l vmPFC	1.87 (0.38)	-1.94 (0.54)
281 – r vmPFC	1.57 (0.39)	-1.98 (0.60)
Bartra, McGuire, and Kable (2013)		
265 – l striatum	0.08 (0.98)	-18.40 (<0.01)
266 – r striatum	2.65 (0.45)	-12.76 (<0.01)
267 – vmPFC	1.50 (0.44)	-4.97 (<0.01)
268 – l AI	-3.54 (0.58)	-16.97 (<0.01)
269 – r AI	2.53 (0.76)	-2.06 (0.72)
133 – PCC	1.32 (0.80)	-13.88 (<0.01)
271 – Brainstem	-0.90 (0.85)	-7.69 (0.02)
215 – ACC	0.20 (0.97)	-9.15 (<0.01)
213 – dACC/pre-SMA	0.79 (0.85)	-7.30 (<0.01)
Massar, Libedinsky, Weiyang, Huettel, and Chee (2015)		
274 – r supramarg. gyr.	0.51 (0.97)	-17.15 (<0.01)
276 – l inf. temp. gyr.	0.63 (0.88)	2.47 (0.62)
277 – l IFG	-1.80 (0.63)	-8.25 (0.02)
280 – l postcentral gyr.	1.58 (0.65)	1.50 (0.63)

Table A.6 Load-specific and load-independent N-back activity. Shading reflects $p < 0.01$ (dark), $p < 0.05$ (medium), and $p < 0.10$ (light).

Strikingly, nearly all of the SV nodes anticipated by the meta-analysis of Bartra, McGuire, and Kable (2013) show both negative effects of load during offer valuation and also negative encoding of averaged N-back activity on discounting. That is, during N-back task performance, these nodes show sensitivity to individual differences, in that they exhibit greater activity in those individuals that are steep cognitive effort discounters (i.e., low AUC individuals). Though it is striking these many of these nodes encode load negatively during valuation and higher subjective effort positively (as load-independent activation) during the N-back, the opposing signs complicate interpretation. In other words, if these regions signaled high cognitive effort during the N-back (greater activity in individuals that subjectively experience greater cognitive effort), then one might also expect a positive encoding of cognitive effort during valuation (i.e., greater activation on

decision trials associated with a high-load / high-effort option). But instead all of these regions negatively encode cognitive effort during valuation (i.e., indexing it as a cost or discount factor that reduces the value of the option), with activity reducing as the signaled effort associated with that option goes up. Although the exact reasons for the sign changes in the effects observed in these nodes is unclear, the encoding of both types of effects does raise the possibility that any one (or all) of these a priori SV nodes may play critical roles in tracking effort costs during engagement with a demanding task for the purposes of integrating that cost information during effort-based decision-making.

A.6 Discussion

In this chapter, two broad questions were asked about brain activity during effort-based offer valuation and decision-making: 1) which regions are involved in active decision-making beyond non-decision processing of response execution and 2) which regions are involved in encoding key choice dimensions during valuation.

During decision-making, task-positive showed stronger activity on close offer comparisons (on regular versus catch trials) confirming their role in supporting demanding offer comparison and decision-making in addition to their role in supporting demanding N-back performance. Nodes of interest in these networks, especially including the dACC and IPS, showed particularly strong effects of choice difficulty. Encoding of choice difficulty was considerably more robust for effort-based than delay-based decision trials. This likely reflected the accumulated fatigue that was much greater for delay-based trials, which always came at the end of a long protocol. Unfortunately, this also precludes rigorous comparison of delay and effort-based decision-making

activity patterns, which although not the primary focus, was a secondary goal of the study.

In addition to choice difficulty, the encoding of choice and biasing conditions was investigated during the decision window. Most notably, the caudal PCC was found to play a key role in encoding choice of hard task > easy task approximately 8-12 seconds after both offers were made available. On closer inspection this reflected timecourses which were initially indistinguishable on hard choice versus easy choices, but then dropped significantly faster for hard choice versus easy choice trials later in the trial epoch. One interpretation is that this region encodes a form of “outcome regret” that associated with committing to higher costs. When time series were further broken down by biasing condition, a more complex pattern was revealed, indicating this region peaked in activity at 6 seconds after offer period onset and was greatest when participants chose the hard task relative to the easy task, particularly when overcoming an offer bias towards the easy task. Hence, the caudal PCC, which is otherwise implicated in valuation processes, appears to carry information both about intrinsic motivational state, and also post-decision cost encoding. Note that this region was close to loci identified in the meta-analysis of Barta et al. (2013) for SV encoding during valuation and during choice outcome receipt. Node #133, best approximating the loci of Barta et al., did not encode choice and bias like the more caudal PCC cluster identified in this dataset. However, it did encode task load during the valuation window (6 to 8 seconds) and thus negatively correlated with trial-wise variation in first-offer SV. Moreover, it was one of the nodes in which N-back load effects in the DMN predicted load effects during valuation.

The encoding of key decision variables was investigated in a series of a priori nodes heretofore implicated in encoding either domain-general SV or effort-based SV. Consistent with meta-analyses implicating the vmPFC and VS in representing SV during decision-making, nodes in both regions positively encoded the SV of effortful rewards, by amplitude modulation of the hemodynamic response, 6 to 8 seconds after participants were presented with a single offer. Furthermore, averaged time courses from the VS and the vmPFC revealed encoding of both reward amount (positively) and task load (negatively), providing evidence that the VS and vmPFC encode anticipated cognitive load as a cost during effort-based offer valuation. To my knowledge, this finding is the first unambiguous demonstration that the VS and vmPFC incorporate information related to cognitive effort as a unique cost factor that discounts the subjective value of an offer during decision-making. Furthermore, single offer valuation-related activity in this region was predictive of subsequent choice in a manner suggesting a causal role in decision-making and overcoming offer biases. Prior to this study, there has been indirect evidence suggesting that the VS might encode the cost of cognitive effort. Botvinick et al. (2009) showed that the VS was less active when processing a reward cue immediately after disengaging from a highly demanding cognitive task versus a less demanding task. While this may have reflected, as the authors claim, “cognitive effort discounting” in the VS, it may have also reflected some other (non-cognitive) response in the striatum following vigorous activity immediately prior. Similarly, Schmidt et al. (2012) have shown that VS activity scaled with reward and also performance, during task engagement, that suggested it might encode cognitive motivation. However, this is also indirect in the sense that it may have only encoded reward and not cost information. Hence, the direct encoding of

cognitive cost during pure valuation is a an important and novel finding, implicating the VS in encoding a critical dimension of cognitive motivation.

While the vmPFC and VS showed reliable encoding of both choice dimensions of reward and amount, most other a priori nodes encoded SV, but only because they negatively encoded load and not amount (with the exception of the more ventral/anterior ACC node – #215 – which encoded both dimensions). This result was interesting for a number of reasons including that load encoding was so widespread and robust across a wide range of nodes, and also because reward magnitude has elsewhere been shown to have such widespread and robust effects (Vickery et al., 2011). The widespread encoding of load was also interesting because it suggested a number of loci for the integration of effort costs during effort-based valuation and decision-making.

Although the vmPFC showed encoding of both load and amount dimensions, it showed only somewhat limited evidence for *subjective* encoding of these dimensions. Although there were ordinal differences in activity of this region that sensibly tracked individual differences in effort discounters (higher activity in shallow discounters compared to steep discounters), the only reliable statistical evidence that discounting predicted vmPFC activity obtained late in the time course, during the decision rather than the valuation window. This pattern of results suggests that limited sample size and likely limited power to detect individual differences at this level of analysis.

Intriguingly, though evidence of subjectivity was limited in the vmPFC, there is evidence of subjectivity elsewhere. In particular, discounting (AUC_{3S}) predicted individual differences in the slope of the linear amount effect in the left VS, positively, such that shallower discounters had stronger amount effects. Additionally, novel but

tentative evidence was observed for the subjective encoding of effort costs in other valuation nodes. Specifically, in the subset of nodes that encoded load but not reward amount (e.g., AI, dACC), mean valuation period activity was positively correlated with mean N-back activity in the DMN. One exciting possibility is that the brain tracks effort by the degree of DMN deactivation during task engagement (which deepens with increasing load), and this information is integrated into valuation processes via SV nodes (which also show decreasing activation with increasing load) during offer valuation. Again, the fact that mean DMN activity load effects did not vary with individual differences in cognitive effort discounting limits stronger inferences about evidence for subjectivity in cost encoding. On the other hand, another particularly intriguing finding was that individual differences in the linear slope of load effects (i.e., increasing deactivation with load) in the DMN during the N-back predicted the linear slope of load effects during valuation (again increasing deactivation with load) in multiple a priori SV nodes. This result supports the hypothesis that subjective effort is tracked by DMN deactivation and is integrated via cost encoding in the PCC, ACC, vmPFC, IPL, and brainstem during effort-based valuation and decision-making.

Chief limitations of this experimental design include a sample size that was somewhat small for a focus on individual differences, and may not have been sufficiently powerful enough to detect subtle patterns reflecting subjectivity effects in amount and load encoding. Another limitation was that the delay-based decision trials always occurred at the end of a long experimental session, and so may have been confounded with increased fatigue and disengagement during this part of the session. Evidence of disengagement included faster response times during delay relative to effort-based

decision trials, and a key consequence is that choice dimensions were weakly encoded. The fixed experimental order was designed intentionally to strengthen results on effort-based decision trials, which were the primary focus of the experiment. However, differential signal-to-noise confounded cross-domain comparisons of effort and delay. Proper comparisons can be achieved in future studies with counterbalanced order. Finally, the rapid pacing and fixed trial-to-trial interval were not optimal for detecting parametric variation by trial parameters. Jittering by trial type enabled resolution of key parameters of interest, but future designs that incorporate greater variability in inter-trial intervals would potentially yield greater resolution and effect estimation. These modifications may make decoding of subjectivity easier to detect. Additional, more general issues that could be addressed in follow-up studies are discussed next in the final chapter.

A.7 General Discussion

A.7.1 Decision-making Behavior

Discounting procedures like COGED are assumed to yield subjective indifference points that quantify effort costs, and the present results support this assumption. They typically proceed by stepwise titration of offers until a point of subjective equivalence is reached, i.e., when participants are indifferent between higher demand task for more money and lower demand task for less money. The titration procedure is a common one for finding indifference points in the more extensive risk and delay based discounting literatures. And yet, few studies have verified that indifference points established in this way are stable and meaningful. Here, indifference point stability was confirmed by the pattern of choices observed in the final imaging session. During that session, offers for

the 1-back were designed to be slightly (or largely) above or below indifference points measured for each participant and each N-back level, earlier in the same session. Choice frequency plots demonstrated that, across levels, the participants were reliably more likely to choose the 1-back when the offer was higher than indifference, and more likely to choose the harder task when the offer was lower than indifference. Indifference point precision was confirmed by response time analyses. These results showed that decisions made about offers very far from indifference (when one option is clearly superior to another) were relatively fast, while decision times slowed as offers were made in closer proximity to indifference. This pattern of results was as strong, if not stronger, than that observed in the delay discounting data, for which procedures are much more well-established (Green & Myerson, 2004). As such, the particular discounting procedure and assumptions about indifference point stability and precision were validated generally. More immediately, strong inferences are supported regarding precise and reliable measurement of subjective costs in performing various levels of the N-back task and across various participants.

The lack of a correlation between COGED and TEMPD AUC in this dataset was not anticipated. However, the general trend of the data was consistent with prior observations (Westbrook et al., 2013). In the prior study, discounting of the two domains was positively related because steep effort discounting participants were exclusively steep delay discounters, while shallow effort discounters showed a range of delay discounting. The current dataset follows the same general pattern, with the exception of a single participant who was a steep effort discounters but also a shallow delay discounters. Given the small sample size, it is difficult to draw any strong conclusions about

individual differences. Nevertheless, the current data are not fundamentally inconsistent with prior observations.

A.7.2 Brain Regions Engaged in Effort-Based Decision-Making

A vast neuroeconomics literature investigating decision-making regarding diverse cost dimensions including delay, risk, and physical effort has implicated a canonical and purportedly domain-general network of regions that encode SV or are involved in offer comparison. Regions representing SV include the vmPFC, VS, and PCC, while the dACC, dlPFC, and IPS have been particularly implicated in decision-making on difficult choice trials, i.e., when offers are close in SV (Bartra et al., 2013; Basten et al., 2010; Kahnt, Heinzle, Park, & Haynes, 2011; Levy & Glimcher, 2012; Shenhav et al., 2014). As described, only a handful of studies have examined cognitive effort value encoding directly (Botvinick et al., 2009; Massar et al., 2015; T. Otto, Zijlstra, & Goebel, 2014; Schmidt et al., 2012; Schouppe et al., 2014; Vassena et al., 2014). Studies in the domain of cognitive effort have implicated a network of regions encoding cognitive motivation (as during effort anticipation or performance) that is very similar to the network implicated in decision-making generally, particularly including the ACC and the striatum. This is consistent with the aforementioned role hypothesized for the dACC in value-based regulation of cognitive control (Shenhav et al., 2013), and for the ventral striatum in general value-learning about states and actions, reviewed in (Niv, 2009). So far, however, only one study has examined activity encoding choice dimensions during cognitive effort-based decision-making (Massar et al., 2015). Hence, there is very little data about regions supporting this critical class of decisions.

Two sorts of questions were asked of the present dataset: 1) what is the brain doing during active offer comparison apart from more basic perceptuo-motor processes associated with indicating a trivially better option? and 2) where does the brain encode choice dimensions during effort-based decision-making? The design of the in-scanner decision trials permits asking both questions cleanly.

Scanner decision trials included offers that were close to and far from each participants' indifference points for every task load. At the limit, when offers are very far from indifference, decision-making amounts to trivially identifying the obviously superior offer. These catch trials provide a good contrast against those trials in which offers are close to indifference and offers must be compared carefully. Catch trials engage, and therefore provide a good control for those perceptuo-motor "non-decision" processes like response mapping and execution. A contrast of catch and regular decision trials has revealed multiple loci including the dACC, the IPS, and the dlPFC, all of which have otherwise been implicated in comparison of close offers (Basten et al., 2010; Pine et al., 2009; Shenhav et al., 2014). This result validates that the network of regions observed for supporting close decisions in other domains extends to decisions about cognitive effort as well. It also validates a key assumption of the COGED paradigm: that it precisely estimates subjective indifference so that offers close to indifference are actually close to indifference. Interestingly, a dACC node that showed a robust regular versus catch trial contrast, and was thus more active for difficult decision trials, was also shown to encode anticipated cognitive load negatively during offer valuation. This implicates dACC in concurrently tracking cognitive effort during decision-making in both a prospective manner (in terms of the effort associated with the choice) and as it unfolds,

during the decision-making process itself. Interestingly, these two forms of demand were encoded in opposite directions and at different stages of the trial: as a load-related deactivation during the valuation window, and as choice-difficulty related activation during the decision window. A key implication for future research is that the dACC appears to encode both ongoing and prospective effort simultaneously. As such, future studies investigating the role of the dACC in effort expenditure and decision-making should be designed to clearly resolve both influences.

In addition to difficulty, choice and bias were also encoded during decision-making. Of particular note, a set of caudal PCC nodes were reliably less active 8—12 seconds after decision window onset on trials in which participants selected the more demanding over the less demanding option. One interpretation is that the caudal PCC encodes greater costs associated with committing to the more demanding option. This interpretation is consistent with the hypothesis that the PCC tracks action outcomes to drive adaptive changes in behavior (Pearson et al., 2011). Prior to this, at 6 seconds after decision window onset, activity peaks in these caudal PCC nodes and is reliably higher on trials in which the hard option is selected overcoming an easy offer bias relative to when the easy option is selected despite that offer being below subjective indifference. This suggests that the caudal PCC, like the vmPFC (and supramarginal gyrus, AI, and dACC) encodes state intrinsic motivation during effort-based valuation and decision-making.

The second kind of question addressed by this paradigm is which brain regions encode choice dimensions. The experimental design includes a valuation period, during which participants have the opportunity to consider a single offer in isolation. This design

contrasts with most other decision-making studies that have considered valuation effects during a time window in which both offers are available. The advantages of the current design is that it allows for a pure analysis of whether various decision variables, such as amount and load, are encoded for this offer, without the associated complication and confounds of distinguishing multiple offers and the associated decision period.

All a priori regions of interest identified for encoding SV in previous meta-analyses, including the vmPFC, VS, AI, PCC, brainstem, and ACC showed positive encoding of SV pursuant to single offer valuation. Interestingly, several of these regions – the dACC, PCC, and AI in particular – showed reliable encoding of load (negatively) but not reward amount. Hence, they evinced positive SV encoding, but only because they encoded load negatively. This result is also interesting given that the AI, in particular, has elsewhere been implicated as part of a “pain matrix” encoding aversive stimuli, and responding robustly, in terms of increased activation, both to punishments and to physical and cognitive effort (Craig, 2002; T. Otto et al., 2014; Prévost et al., 2010; Treadway et al., 2012). The negative sign of load encoding observed in the AI during the current study is thus not straightforwardly predicted by this prior literature. Nevertheless, positive encoding of both rewards and punishments have been observed in the AI (and the striatum as well) (Bartra et al., 2013), supporting that this region plays some as yet unidentified role in incorporating cost and penalty information into cost-benefit valuation processes. The VS, for its part, has been shown to encode the selection of high cognitive load positively (in case a participant freely selects higher load) or negatively (in case the participant is forced to select higher load) suggesting that more complex underlying decision-making processes can yield either sign at the level of local field potentials and

BOLD signal (Schouppe et al., 2014). Thus, it is not straightforward to predict the sign of the effect; the important result is that load is encoded. The most direct interpretation of the current result is that many of these valuation regions are negatively encoding costs and positively encoding benefits towards a single common currency of SV during single offer valuation. In the COGED paradigm, the cognitive load of the N-back associated with an offer factors in as a cost variable that causes participants to discount the value of that offer. As such, load (and concomitant effort costs) should be negatively encoded in an SV region.

Although a number of regions were found to encode load and amount as anticipated, there was less robust evidence that the encoding reflected the subjective dimension of valuation, at least in terms of individual differences in effort discounting. In particular, load effects and amount effects were not found vary with individual differences in discounting rate. One notable exception is that a left VS node was shown to have an amount effect that reliably increased with AUC_{3S} . This does provides some support for the idea that reward amount, if not cognitive load, is reflected in the subjective encoding of value in the VS. It is possible that further subjectivity effects were not detected because the limited sample size provided insufficient power to detect relatively subtle individual differences. Also, subjectivity is likely to be a weaker effect to the extent that state factors introduce variability in subjective value encoding (e.g. with intrinsic motivation, from one trial to the next).

Intriguingly, however, there was some evidence of subjective encoding of load with respect to individual difference covariation between DMN activity during N-back performance and activity in left and right AI during single offer valuation. Specifically,

during the N-back, the DMN negatively and monotonically varied with load, while during valuation, the AI nodes negatively and monotonically varied with load. Critically, these effects correlate (for both nodes) across participants such that those individuals exhibiting stronger load-independent DMN deactivation in the N-back also showed stronger average AI deactivation during offer valuation. This result provides evidence in support of two hypotheses: 1) that subjective, phenomenal effort reflects the extent to which individuals suppress activity in the DMN while they perform demanding tasks and 2) that such cost information becomes integrated into subjective cost-benefit valuation via the AI.

Even more intriguingly, individual differences in the steepness of load functions in the DMN predicted steepness of load effects during the valuation period in several a priori SV nodes. Again, this result supports the hypothesis that DMN deactivation tracks subjective effort during task engagement, and cost information is incorporated into valuation processes through the vmPFC, PCC, IPL, ACC and brainstem in particular.

In sum, the present data set has provided strong and novel evidence for encoding of anticipated cognitive load as a cost during valuation and decision-making, and preliminary evidence for subjectivity in representations of effort costliness and reward desirability in a core valuation and decision network. In addition, this data has yielded multiple lines of evidence suggesting that effort costs are tracked by DMN deactivation and also the particular loci by which this effort cost information becomes integrated into a valuation process.

A.7.3 Future Directions

Regarding decision-making, an obvious future direction is to examine functional connectivity among the many regions of interest. To the extent that the AI nodes track costs that become integrated with benefits in the vmPFC, for example, we should expect not only trial wise correlation with offer SV, but also greater functional connectivity between these regions during the valuation period. During the decision period, to the extent that the IPS is involved in cost-benefit comparison, it may also show greater functional connectivity with the vmPFC on regular trials relative to catch trials. Also, aside from the intriguing individual differences correlation between the DMN during the N-back and the nodes including the AI, PCC, ACC, and vmPFC during offer valuation, functional connectivity analyses could be used to examine the linkages between specific regions implicated in tracking effort, like particular DMN nodes, and regions integrating effort cost information like the AI or vmPFC. Dynamic causal modeling may also prove particularly informative about how cost information is conveyed to valuation regions during valuation, as one would predict directed transfer from a tracking region to a valuation region under that condition.

Appendix B

B.1 Response times for effort-based decision trials

As described in Appendix A, COGED decision difficulty was successfully manipulated by modulating the bias parameter. This was demonstrated by examining both choice probabilities and RTs. As shown in Figure B.1, median RTs were faster at the limits, and slower for smaller magnitude proximity parameter settings (note that these same data were presented, broken out by biasing condition in Appendix A, but here collapse across that variable for comparison purposes). This supports that proximity parameter settings of -1.0 and 1.0 are treated as trivial decisions (referred to as “catch trials”) and those closer to indifference as more difficult, as intended.




Figure B.1 Median effort-based decision RTs by proximity parameter.

The large effect in median RT values by proximity parameter values contrasts with the very small effect of other decision trial parameters: base amount, and task level. As shown in Figure B.2, the effects of either of these parameter variations are relatively small in terms of median RTs. Moreover, there are no pairwise differences in median RTs among base amount, or among task level parameter settings.

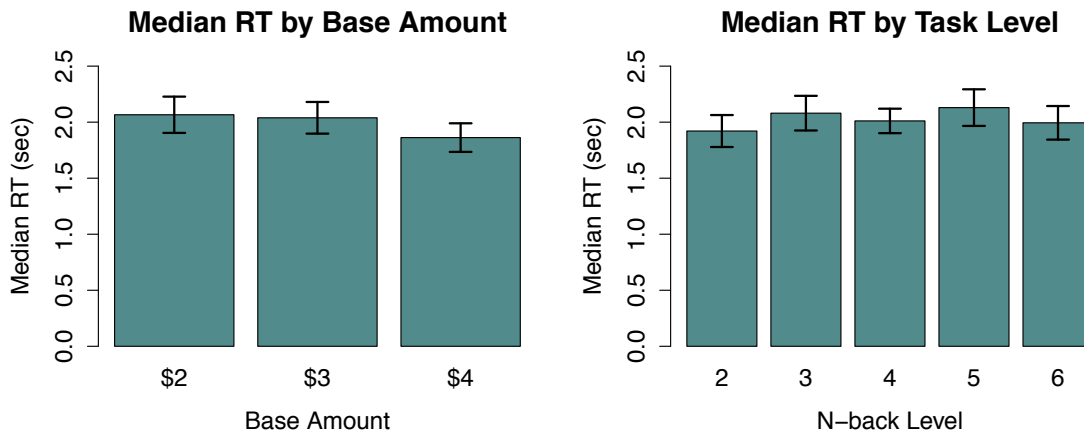


Figure B.2 Median effort-based decision RTs by base amount and N-back Task Level.

B.2 Response times for delay-based decision trials

RT patterns also provide evidence that bias influenced TEMPD decision difficulty as anticipated. As shown in Figure B.3 (again these are the same data presented in Appendix A, but here collapsed across bias condition), median RTs were faster at the limits, and slower for smaller magnitude proximity parameter settings. The large difference in median response times reported in the main text supports that the participants robustly treat decision trials at the limits differently than they do decision trials closer to indifference.

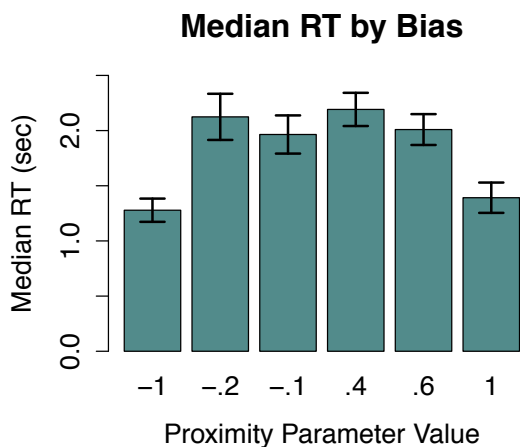



Figure B.3 Median delay-based decision RTs by proximity parameter value.

Again, as with COGED, the large effect in median RT values by bias settings contrasts with the very small effect of other decision trial parameters: base amount, and delay. As shown in Figure B.4, the effects of either of these parameter variations are relatively small in terms of median RTs. Moreover, there are no pairwise differences in median RTs among base amount, or among task level parameter settings (all p 's ≥ 0.26).



B.4 Median delay-based decision trial RTs by base amount and delay.


Appendix C

C.1 Regions tracking effort during the N-back


C.1.1 Whole-brain voxel-wise load functions

One of the questions explored in the main text is how brain regions vary as a function of N-back load. An analysis of a priori task-positive networks and the DMN revealed a set of linear and quadratic functions in each network. To complement these results, voxel-wise whole-brain t-tests were conducted using linear contrasts of block regression weights to test for linear effects of load, and polynomial contrasts of block regression weights to test for quadratic (inverted-U) effects of load. These t-tests were then thresholded at each voxel at $p < 0.005$, and cluster corrected to $p < 0.01$ (cluster extent ≥ 97 voxels). Given the unexpected dip, in most task-positive networks, at the 5-back, linear and quadratic effects were tested both over the range $N = 1\text{---}5$ and $N = 1\text{---}6$. Figure C.1 gives both of these results for the whole brain.

A – Linear, $N = 1\text{---}5$




B – Inverted-U, $N = 1\text{---}5$



C – Linear, $N = 1\text{---}6$

D – Inverted-U, $N = 1\text{---}6$



C.1 t-tests of linear (A,C) and negative quadratic (inverted-U; B,D) across load levels $N = 1—5$ (A,B) and $N = 1—6$ (C,D). Linear and quadratic effects are predominantly negative across both ranges.

As shown in Figure C.1A and C.1C, linear effects are predominantly negative and also largely within the DMN – in particular, the medial PFC wall, the mid-cingulate cortex and the anterior lateral PFC show negative effects of load, as expected for the DMN. As shown in Figure C.1B and C.1D, the quadratic tests reveal strong inverted-U (negative quadratic) load functions in the dlPFC, vlPFC, dACC/preSMA, IPL and IPS. These results, therefore, confirm (negative) linear effects in the DMN, and inverted-U load functions in the FP, Sal, and DorAtt networks, in particular.

C.1.2 Network load functions

By investigating the change in regression weights across loads in each of the 13 intrinsic networks of the brain, as identified by Power et al. (2011), various patterns emerged, including flat, monotonic, and inverted-U functions. A formal test of these functions for each network was accomplished by multi-level models with variable intercepts as described by equations 4.1—4 in the main text. The result of those multi-level models for all 13 networks are provided in the following table. Note that because of the unexpected jump in activity at $N = 6$, relative to $N = 5$, linear and quadratic models are fit to both ranges $N = 1—5$ and $N = 1—6$, for exploratory purposes.

Network	Linear effects		Quadratic effects	
	B2*10 ⁻³	p-value	B3*10 ⁻³	p-value
Across Loads N = 1—5				
Unc	-4.5	0.11	-1.5	0.46
SmtSn	-5.3	0.05	0.7	0.64
CO	-4.1	0.05	-3.0	0.05
Aud	-5.6	0.03	1.3	0.35
DMN	-6.7	0.01	2.4	0.23
MmRtr	-2.5	0.39	-1.7	0.42
VntAtt	-6.6	<0.01	-1.4	0.31
Vis	-7.1	<0.01	0.2	0.92
FP	1.7	0.45	-6.6	<0.01
Sal	-1.7	0.40	-4.2	0.02
SubC	-2.0	0.45	-3.5	0.08
Crblr	-4.8	0.08	-2.6	0.10
DorAtt	-2.2	0.34	-4.5	0.01
Across Loads N = 1—6				
Unc	-2.3	0.23	0.8	0.50
SmtSn	-2.9	0.11	1.8	0.13
CO	-2.7	0.11	-0.2	0.82
Aud	-2.8	0.12	2.2	0.03
DMN	-4.7	0.01	2.1	0.10
MmRtr	-1.6	0.38	0.01	0.96
VntAtt	-3.9	0.02	1.2	0.16
Vis	-3.9	0.03	2.1	0.06
FP	0.9	0.62	-3.0	<0.01
Sal	-1.3	0.42	-1.3	0.23
SubC	-1.8	0.40	-1.2	0.31
Crblr	-2.9	0.12	0.2	0.88
DorAtt	-2.2	0.09	-1.7	0.27

Table C.1 Linear and quadratic fixed effects of load in all 13 networks for N = 1—5 and N = 1—6.

Shading: light grey for $p < 0.10$, medium for $p < 0.05$, and dark for $p < 0.01$. Network labels are abbreviated: Unc for Uncertain, SmtSn for Somatosensory, CO for Cingulo-Opercular, Aud for Auditory, DMN for Default Mode Network, MmRtr for Memory retrieval, VntAtt for Ventral Attention, Vis for Visual, FP for Frontal-Parietal Task Control, Sal for Salience, SubC for Sub-cortical, Crblr for Cerebellar, and DorAtt for Dorsal Attention.

A full set of boxplots across load levels for all networks is provided in Figure C.2.

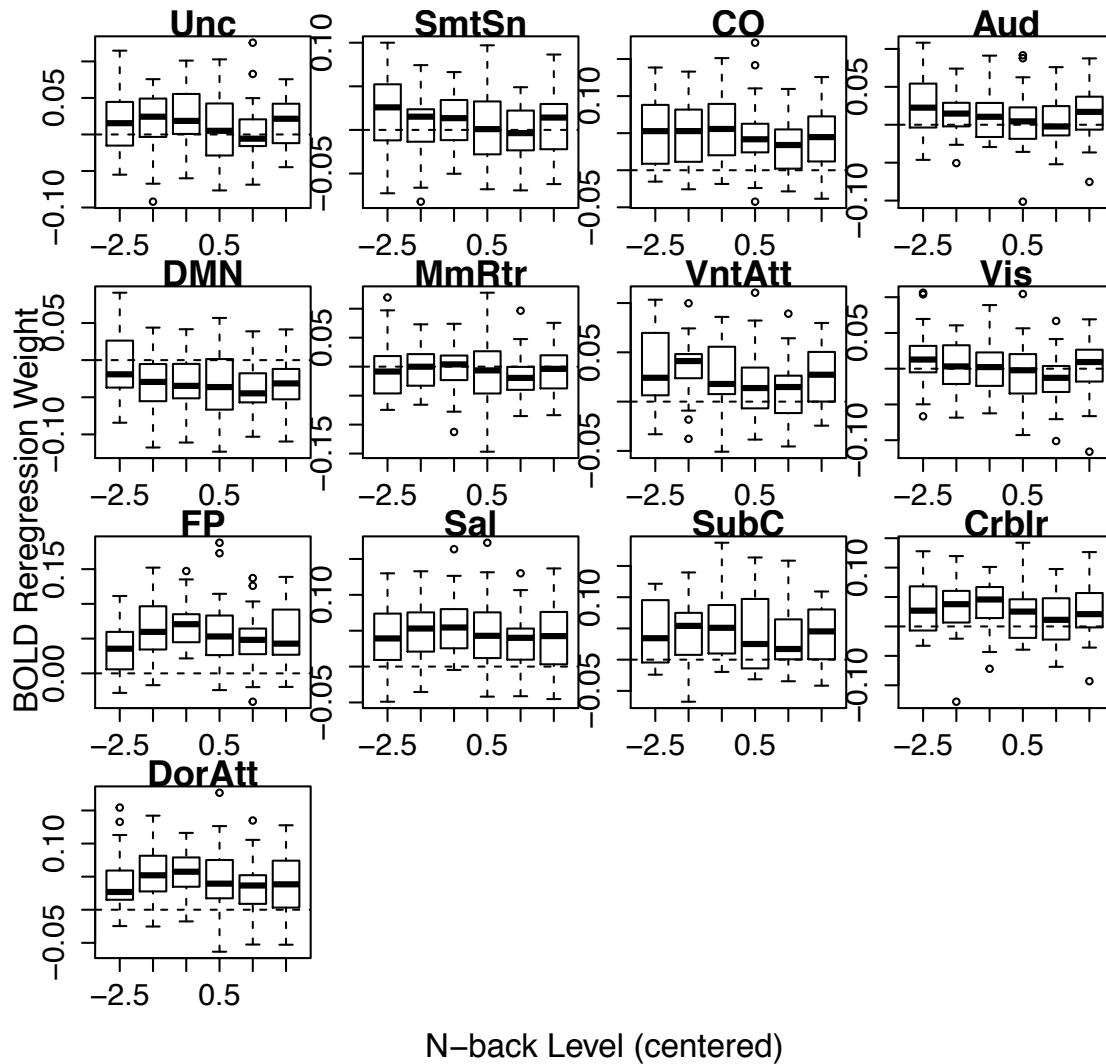


Figure C.2 Mean recruitment within each of 13 networks across N-back levels $N = 1-6$.

C.1.3 What is the relationship between discounting and activity at each load and in each network?

As described in the main text, separate multi-level models of N-back regression weights were computed for each network with predictors of AUC_{3S} and dummy codes indicating Load (1-back as the baseline for contrast – thus making no assumptions about the shape of the load function), along with the interaction of AUC_{3S} and each load level, all nested within participants. The general form of each multi-level model is given by:

$$\beta_i = B1_{j[i]} + B2_{j[i]} 2back + B3_{j[i]} 3back + B4_{j[i]} 4back + B5_{j[i]} 5back + B6_{j[i]} 6back + \varepsilon_i \quad (C.1)$$

$$B1_j = \gamma_{1,0j} + \gamma_{1,1j} AUC_{3Sj} + \eta_{1j} \quad (C.2)$$

$$B2-6_j = \gamma_{x,0j} + \eta_{xj} \quad (C.3)$$

where β is the network-averaged regression weight being predicted by an intercept term $B1$, and $B2$ — $B6$ referring to dummy codes for the 2-back—6-back, respectively. Note that all predictors are subscripted to indicate the nesting of load level i , within participant j . Note also that AUC_{3S} is included as a predictor at the participant level of the model for the intercept, but not as a predictor of Load effects (no cross-level interactions between Load and AUC_{3S}). A full model in which AUC_{3S} was included as a predictor of dummy-coded Load effects was found to explain insufficient variance to justify the additional degrees of freedom in nested model comparisons, in all networks. Thus, the simpler models are presented here. This also means that model fits do not support cross-level Load x AUC_{3S} interactions. The results for all 13 models (one for each network) are provided in Table C.2 below.

Network	Fixed Effects Estimates*10 ⁻² (p-values):					
	AUC _{3S}	2-back	3-back	4-back	5-back	6-back
	$\gamma_{1,1}$	$\gamma_{2,0}$	$\gamma_{3,0}$	$\gamma_{4,0}$	$\gamma_{5,0}$	$\gamma_{6,0}$
Unc	-4.8 (0.09)	-0.60 (0.58)	0.24 (0.82)	-0.14 (0.23)	-0.18 (0.10)	-0.57 (0.60)
SmtSn	-5.3 (0.02)	-1.7 (0.07)	-1.1 (0.22)	-2.0 (0.05)	-2.5 (<0.0)	-1.4 (0.15)
CO	-8.5 (<0.01)	-0.04 (0.97)	0.42 (0.70)	-0.37 (0.75)	-1.9 (0.07)	-0.59 (0.57)
Aud	-3.4 (0.14)	-1.4 (0.10)	-1.5 (0.11)	-2.2 (0.03)	-2.4 (0.01)	-1.2 (0.19)
DMN	-1.0 (0.68)	-2.2 (0.04)	-2.1 (0.04)	-2.9 (0.02)	-3.0 (<0.0)	-2.7 (<0.01)
MmRtr	-7.4 (<0.01)	-0.04 (0.97)	0.07 (0.95)	-0.13 (0.92)	-1.2 (0.23)	-0.39 (0.72)
VntAtt	-4.6 (0.04)	0.03 (0.97)	-0.59 (0.53)	-1.8 (0.05)	-2.4 (<0.0)	-0.99 (0.25)
Vis	-6.9 (<0.01)	-1.4 (0.13)	-1.3 (0.20)	-2.3 (0.03)	-3.1 (<0.0)	-1.5 (0.17)
FP	-7.2 (<0.01)	2.5 (0.02)	3.4 (<0.0)	2.1 (0.07)	1.0 (0.32)	1.8 (0.11)
Sal	-8.6 (<0.01)	0.80 (0.48)	1.5 (0.20)	0.57 (0.64)	-0.74 (0.51)	0.17 (0.86)
SubC	-7.9 (<0.01)	0.33 (0.80)	1.3 (0.36)	0.18 (0.90)	-0.91 (0.49)	-0.28 (0.80)
Crblr	-12. (<0.01)	-0.10 (0.93)	0.28 (0.81)	0.84 (0.44)	-2.0 (0.06)	-0.70 (0.56)
DorAtt	-9.7 (<0.01)	1.6 (0.11)	1.6 (0.07)	0.44 (0.70)	-0.53 (0.59)	-0.00 (0.99)

Table C.2 Fixed main effects (and p-values) of multi-level model of N-back regression weights predicted by AUC and Load. Shading in the table indicates significance level with light grey for $p < 0.10$, medium for $p < 0.05$, and dark for $p < 0.01$.

While nearly all networks show some evidence of negatively encoding AUC_{3S}, it is possible that such diffuse encoding reflects the outsized influence of a single outlier. Closer inspection of the AUC_{3S} effects reveals the outsized influence of a single participant (SU1514, AUC_{3S} = 0.31) driving negative associations between discounting and BOLD signal in multiple networks. An exemplary figure highlighting the influence of this particular participant for the VntAtt network is given in Figure C.3.




Figure C.3 Network-level individual differences activity across VntAtt nodes by AUC_{3S}, suggests that data from SU1514 is driving the negative relationships in this network. The solid line gives the linear regression, and dashed lines give the 95% CI.

Upon removal of the influential participant, and subsequent re-analysis, multiple regions no longer show reliable AUC_{3S} effects. In fact, only two networks show AUC_{3S} effects after removal: the DorAtt and Crblr network, as shown in Table C.3. All other AUC_{3S} effects are no longer reliable (all p 's ≥ 0.14). As noted in the main text, this result does not mean that COGED does not relate to activity in a wider set of networks. Indeed, there is little sampling below AUC_{3S} < 0.5, and SU1514 may accurately reflect the tendency of this (sparsely sampled) underlying population of steep discounters. Moreover, when considering data from all task-positive networks together, the aggregate set of networks has a reliable effect of β_{avg} weights predicting SV_{3S} ($\gamma_{11} = -6.2 \times 10^{-1}$; $p < 0.01$), even excluding SU1514. In that case, there is a wider encoding of discounting across task-positive networks. This result does, however, suggest additional caution in inferring COGED effects across a wider range of networks. Conversely, it indicates that the COGED relationships observed in the DorAtt network are particularly robust.

Network	Main Effects Estimates*10 ⁻² (p-values):					
	AUC _{3S}	2-back	3-back	4-back	5-back	6-back
	$\gamma_{1,1}$	$\gamma_{2,0}$	$\gamma_{3,0}$	$\gamma_{4,0}$	$\gamma_{5,0}$	$\gamma_{6,0}$
Crblr	-9.7 (<0.01)	3.0 (<0.01)	3.5 (<0.01)	1.9 (0.05)	0.1 (0.33)	2.3 (0.04)
DorAtt	-6.8 (0.03)	5.4 (<0.01)	5.5 (<0.01)	3.9 (<0.01)	3.1 (<0.01)	3.7 (<0.01)

Table C.3 Fixed main effects (and p-values) of multi-level model of N-back regression weights predicted by AUC and Load, excluding data from SU1514. Only networks with a significant AUC_{3S} effect are shown, for brevity. Shading in the table indicates significance level with light grey for $p < 0.10$, medium for $p < 0.05$, and dark for $p < 0.01$.

A visual representation of the AUC_{3S} effects in the DorAtt and Crblr networks, by contrast, are provided in Figure C.4. Note that there is a clear effect of AUC_{3S} on BOLD signal such that shallower discounters recruit the Crblr and DorAtt nodes less vigorously.

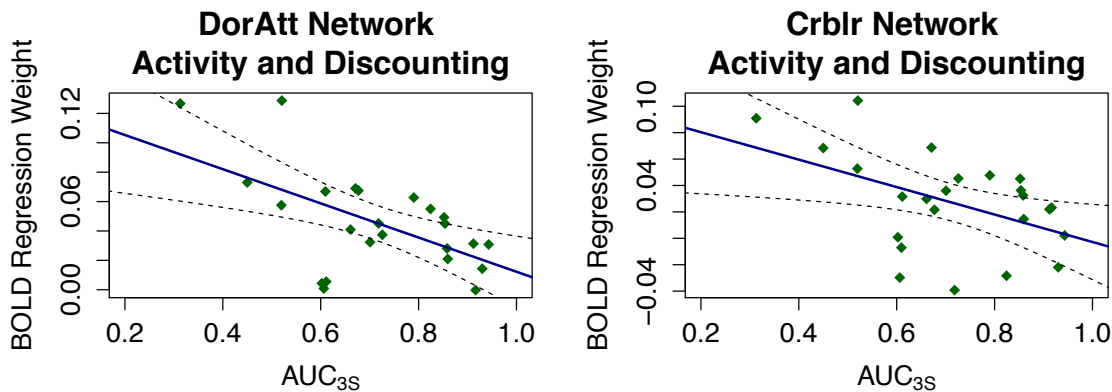


Figure C.4 Network-level individual differences activity across DorAtt and Crblr networks by load-specific A) AUC_{3S}, averaged across task load levels or B) SV_{3S} with task loads modeled separately, showing a clear effect of discounting on activity. Solid lines give linear regressions, and dashed lines give the 95% CI.

C.1.4 Does the relationship between activity and discounting hold when controlling for performance differences?

Activity-discounting relationships support that a network's activity level relates to individual differences in subjective cognitive effort. However, as described in the main text, it is important to show that variance among individual differences in brain response to demanding tasks explained by COGED is not better explained by individual differences

in task performance metrics. To test for this, multiple regression models were fit for each level $N = 2$ — 6 , for those networks showing AUC_{3S} effects, including not only AUC_{3S} , but also measures of response time (ex-Gaussian μRT) and performance (d') at each level included as covariates. Note that ex-Gaussian τRT was also used as an alternative predictor to μRT and fit approximately as well, but the results for the AUC_{3S} effect were the same either way. Also note that exclusion of SU1514 resulted in slightly larger p values for most regions and nodes, but did not fundamentally alter the pattern of results.

$$\beta = B1 + B2 AUC_{3S} + B3 \mu RT + B4 d' \quad (C.4)$$

Table C.4 gives the effect, $B2$, describing the relationship between AUC_{3S} and the network-averaged regression weight for each level of the N-back, controlling for performance measures. As shown, networks particularly including the CO, MmRtr, FP, Sal, Crblr, and DorAtt show relationships to discounting, independent of individual differences in task performance. This supports that these networks track effort per se, and not simply delay or probabilistic discounting associated with N-back task performance. It is important to note, though, that for many regions, inclusion of performance measures reduces the effects to trend-level, with the exception of the MmRtr and DorAtt networks. It is also notable that, again, the most reliable effects are restricted to the 4-back, as shown for the SV analysis above. This suggests that the strongest diagnosticity regarding individual differences in subjective effort might be under the 4-back load.

AUC _{3S} Effect, B2*10 ⁻² (p-value), Controlling for Performance, by Load:						
Network	1-back	2-back	3-back	4-back	5-back	6-back
Unc	-7.6 (0.20)	-4.9 (0.36)	0.7 (0.89)	-13.6 (0.02)	-3.8 (0.41)	-2.0 (0.67)
CO	-10.4 (0.08)	-7.2 (0.18)	-5.0 (0.38)	-12.0 (0.05)	-4.8 (0.33)	-7.2 (0.18)
MmRtr	-7.9 (0.20)	-9.9 (0.03)	-6.4 (0.24)	-21.5 (<0.01)	-11.1 (0.03)	-8.9 (0.09)
Vis	-9.4 (0.07)	-4.5 (0.29)	2.4 (0.62)	-0.10 (0.08)	-3.8 (0.35)	-2.1 (0.70)
FP	-9.7 (0.06)	-10.0 (0.07)	-1.1 (0.82)	-14.8 (0.02)	-6.1 (0.25)	-9.9 (0.12)
Sal	-9.5 (0.11)	-9.7 (0.09)	-4.1 (0.48)	-15.3 (0.01)	-6.9 (0.19)	-9.4 (0.12)
SubC	-9.1 (0.15)	-5.8 (0.36)	-3.9 (0.60)	-18.7 (0.01)	-7.8 (0.18)	-7.6 (0.21)
Crblr	-13.5 (0.01)	-9.4 (0.15)	-3.0 (0.59)	-17.3 (<0.01)	-10.0 (0.06)	-8.5 (0.17)
DorAtt	-9.9 (0.07)	-12.7 (0.01)	-8.3 (0.04)	-18.3 (<0.01)	-11.1 (0.04)	-10.0 (0.09)

Table C.4 AUC_{3S} effects (and p-values) on individual differences in averaged recruitment of each N-back level, for each network, controlling for individual differences in performance. Shading in the table indicates significance level with light grey for $p < 0.10$, medium for $p < 0.05$, and dark for $p < 0.01$.

C.1.5 Do non-linear load functions interact with AUC to predict activity?

Hypothesized Load x AUC_{3S} interactions were not observed in N-back BOLD data. Such interactions would strongly implicate a region in tracking effort during task engagement. Although no interactions have been observed by previous analyses that assume linear load effects, potential non-linear effects of load may harbor Load x AUC_{3S} interactions that were not observed in prior analyses because they only considered load as a linear predictor of recruitment. To investigate this possibility, the random effects (one per subject) from the models fit for Eqns. 4.1—4.4 were further tested for their relationship with AUC_{3S}. Specifically, multiple regressions were fit to test whether linear or quadratic random effects predict AUC_{3S}, for each network showing either a linear or quadratic fixed effect of load.

$$AUC_{3S_j} = B1 + B2 \text{ Linear Random Effect}_j + B3 \text{ Quadratic Random Effect}_j \quad (C.5)$$

Table C.5 shows the resulting regression weights (corresponding to Eqn. C.5) for those regressions that were significant or trending. As with the nodal analysis above, these results are not corrected for multiple comparisons, and are considered exploratory.

Network	Linear effects		Quadratic effects	
	B2	p-value	B3	p-value
Across Loads N = 1—5				
CO	40.1	0.08	66.5	0.09
FP	36.2	0.07	101.7	0.05
Crblr	-4.9	0.38	98.8	0.02
Across Loads N = 1—6				
Crblr	-52.8	0.03	-168.4	0.02
DorAtt	-113.0	<0.01	-12.9	0.15

Table C.5 Significant relationships between the linear or quadratic random effects and AUC_{3S} . Shading: light grey for $p < 0.10$, medium for $p < 0.05$, and dark for $p < 0.01$.

Across $N = 1—5$, the most reliable effects include that in the FP and Crblr networks, which showed reliable (or trending) inverted-U (negative quadratic) fixed effects, more positive (effectively, less negative), quadratic terms predict higher AUC_{3S} . That is, their inverted-U is shallower than those with lower AUC_{3S} values. A visual depiction of this effect is given in Figure C.5. As can be seen, those with lower AUC_{3S} values (darker red) tend to have sharper inverted-U profiles across $N = 1—5$ than those with higher values (lighter yellow). This pattern of sharper inverted-U functions for those with lower AUC_{3S} values is mirrored by trend-level effects in the CO. The inclusion of $N = 6$, for which there was a noticeable increase in activity at $N = 6$ erased this pattern for the FP network and CO networks, and actually reversed it for the Crblr network for which smaller AUC_{3S} participants show an upright-pattern (consistent with a uptick in activity at $N = 6$).

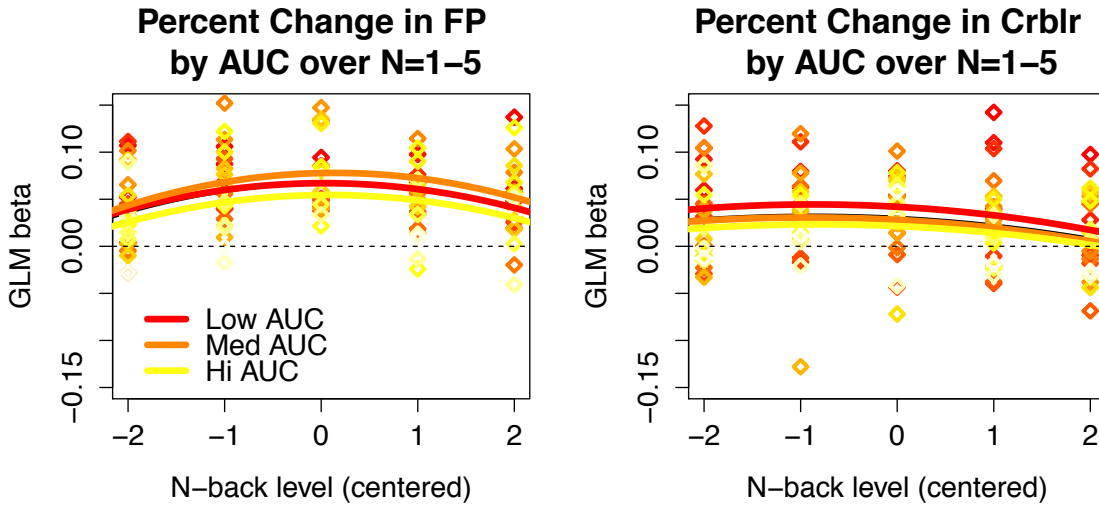


Figure C.5 N-back regression weights across N = 1—5 in the FP and Crblr networks for individual participants as modeled by fixed effects (solid black line) and subject-specific random effects. Fixed and random effects are averaged across three categories of participants, for ease of visualization, according to AUC_{3S} : low ($AUC_{3S} < 0.65$), medium ($0.65 < AUC_{3S} < 0.85$), and high ($AUC_{3S} > 0.85$).

Another reliable effect observed across N = 1—6 is a robust negative linear effect for the DorAtt network. Given that the fixed effect is negative, a negative relationship between linear random effects and AUC_{3S} means that there are stronger negative slopes for those participants with larger AUC_{3S} values (lighter yellow). This effect is depicted in Figure C.6.




Figure C.6 N-back regression weights across $N = 1$ — 6 in the DorAtt and Crblr networks for individual participants as modeled by fixed effects (solid black line) and subject-specific random effects. Fixed and random effects are averaged across three categories of participants, for ease of visualization, according to AUC_{3S} : low ($AUC_{3S} < 0.65$), medium ($0.65 < AUC_{3S} < 0.85$), and high ($AUC_{3S} > 0.85$).

Though these analyses are exploratory, the results suggest that steep and shallow discounters follow different load functions over the range of loads observed here – be it a larger linear decline in the DorAtt network for shallower discounters, or sharper inverted-U functions in the FP and Crblr networks for steeper discounters. A confirmation of any of these patterns in a follow-up study would provide evidence of an effective $AUC \times$ Load interaction, and thereby further implicate the networks in question in tracking cognitive effort.

Appendix D

D.1 Whole-Brain Analyses of Regions Involved in Decision-Making

D.1.1 Voxel-wise Analyses of Regular Versus Catch Trials

As shown in Figure D.1, a canonical working memory and cognitive control network was found to be robustly more active for cognitive effort decision-making in regular compared to catch trials. This network included the bilateral IPS, dACC, pre-SMA, and also the left inferior frontal gyrus and left mid-insula. Other regions more active for regular than catch trials include the bilateral brainstem, lateral occipital lobe, thalamus, and parahippocampal gyrus, and also the left caudal putamen and left cerebellum. Only three regions were less active for regular than catch trials: two clusters in bilateral, ventral inferior frontal gyrus, and one in the left IPL.




Figure D.1 t-test of regions in contrast of regular and catch effort-based decision trials. Clusters shown were voxelwise corrected $p < 0.005$, and cluster thresholded to $p < 0.05$.

In addition to the loci discussed in the main text, the left IPFC has been implicated in working memory processes supporting intertemporal choice, e.g. (Figner et al., 2010); its activity here suggests it supports cost domain-general (both effort and delay-based) decision-making. Left lateral motor cortex engagement reflects the fact that the right

hand was used for responding, and may reflect longer response times and more involved planning when one response must be inhibited relative to another during regular trials. Greater occipital cortex activity may straightforwardly reflect greater visual activity associated with more looking back and forth between two (closer) options.

Also, as shown in the following Figure, the dACC cluster reflecting a robust regular versus catch contrast in effort-based decision trials encompasses the a priori node of interest for encoding SV (Node # 213) described in the main text.




Figure D.2 dACC node overlaid with the medial PFC cluster defined by the contrast of regular versus catch trials.

The same contrast in delay-based decision-making reveals a similar, but much less robust network of regions. In particular, the left motor cortex and left occipital lobe as well as small clusters in bilateral brainstem and IPS were also more active for regular versus catch trials in delay-based decision-making, as shown in Figure D.3. All of these regions mirror effort-based decisions. The contrasts were considerably weaker, however, with obviously smaller cluster sizes, and a peak t-stat of 4.93 for delay-based compared to 9.07 for effort-based decisions. Few other regions show differences between the delay-based decision trial types, with potential exceptions in the vmPFC, posterior medial parietal cortex and, and caudate showing decreased rather than increased activity in regular versus catch trials. While these regions were distinct from those identified for effort-based decisions, the extent and position of these clusters makes them suspect, since

they are lie mostly along edges, ventricles, and in white matter (corpus callosum for the vmPFC cluster).




Figure D.3 t-test of regular versus catch trials during decision window for delay-based decisions. Voxel thresholded at $p < 0.005$, cluster corrected to $p < 0.05$.

D.1.2 Voxel-wise Analysis of Choice and Bias Encoding During the Decision Window

As shown in Figure D.4, the decision to engage in a hard task over an easy task corresponded with lesser activity, 10 seconds in to the decision window, in bilateral PCC and the angular / superior temporal gyrus on the left. This contrast was derived from impulse response functions modeled across the entire decision trial epoch, separately for trials in which the participant selected the hard choice or selected the easy choice. Then, the time point-wise contrasts of the two impulse response functions were tested at each time point in the decision window. The cluster shown below survived correction for the time point 10 seconds in to the decision window.






Figure D.4 Clusters showing lesser activity 10 seconds after decision period onset for trials in which participants select the harder over the easier task. A pair-wise t-test of trial types at that time point was voxel-wise thresholded at $p < 0.005$ and cluster corrected to $p < 0.05$ at the whole brain level. Two clusters obtain A) 227 voxels, cent. of mass: (-1,-51,27), peak t-stat: -4.93 and B) 102 voxels, cent. of mass: (-48,-62,33), peak t-stat = -4.51.

D.1.3 Node-based Analyses of Regions Encoding SV

Using amplitude-modulated tent functions spanning an entire trial epoch, it is possible to test for amplitude modulation by SV across the whole brain. Specifically, after fitting GLMs, 26 predictors of interest are generated (13 mean impulse responses and 13 amplitude modulation responses per trial) corresponding to 13 time points (2 second TRs). Plots of these predictors in all 264 nodes of Power et al. (2011) follows. Note that the plots show the impulse response functions for both effort-based (solid) and delay-based (dashed) decision trials. Also, note that the first set of 264 plots corresponds to the mean response function (Figure D.5), and the second set of 264 plots corresponds to their trial-wise amplitude modulation by SV of the first offer (Figure D.6). As shown in the first set of plots, the impulse response functions for delay and effort-based decision trials are very similar in all nodes, whereas the second set of plots reveal much attenuated amplitude modulation in the delay-based decision trials.





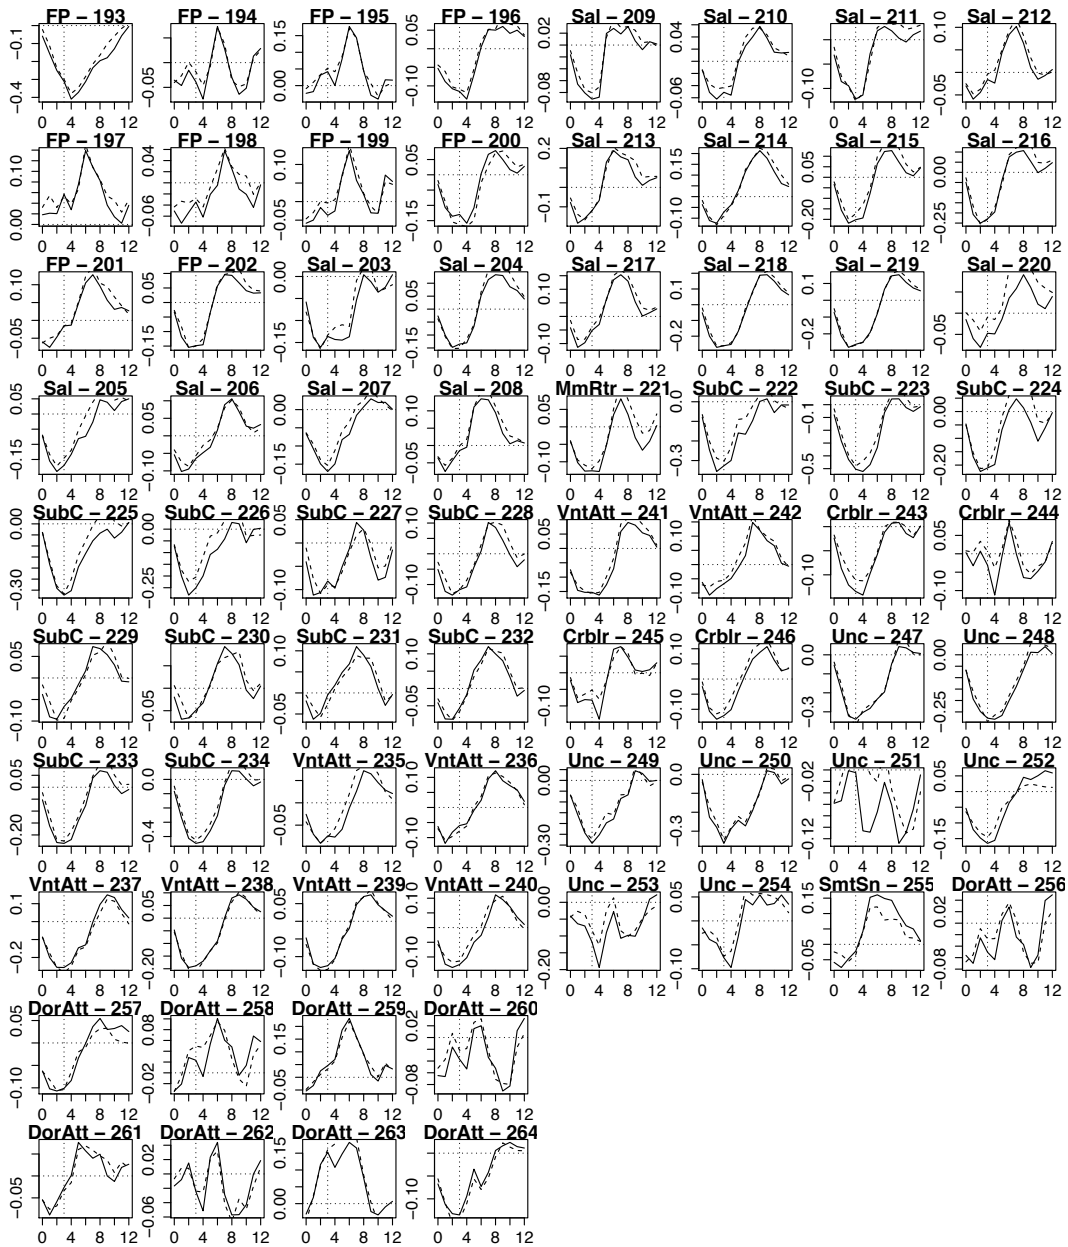

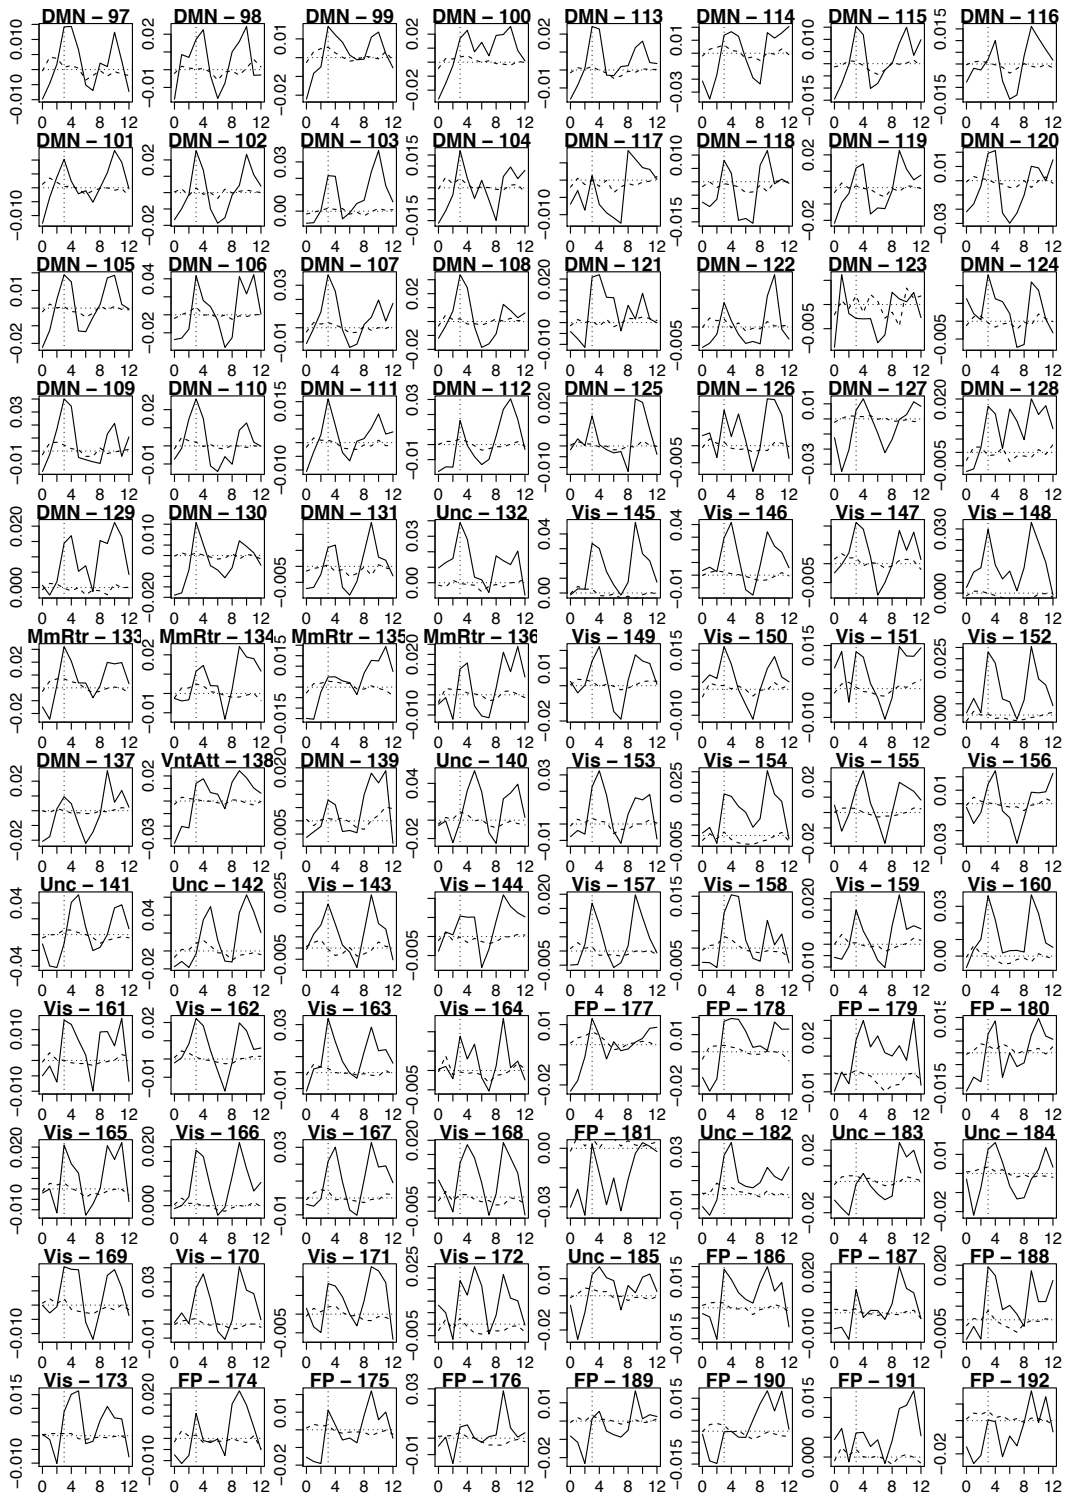


Figure D.5 Mean impulse response function for 13 time points (spanning 24 seconds) across 264 nodes defined by Power et al. (2011). Solid lines reflect effort-based and dashed lines reflect delay-based decision trials. x-axes gives the time point, y-axes give regression weight. Vertical dotted lines show the decision window onset.








Figure D.6 Amplitude modulation of impulse response function for 13 time points (spanning 24 seconds) across 264 nodes defined by Power et al. (2011). Solid lines reflect effort-based and dashed lines reflect delay-based decision trials. x-axes gives the time point, y-axes give regression weight. Vertical dotted lines show the decision window onset.

D.1.4 Voxel-wise Analysis of Regions Encoding SV

D.1.4.1 Regions Encoding SV During Effort-based Valuation

Using amplitude modulated tent functions spanning an entire trial epoch, it is possible to test for amplitude modulation by SV across the whole brain, rather than just in specific nodes as described in the main text. Specifically, after fitting GLMs with 26 predictors of interest (13 mean impulse responses and 13 amplitude modulation responses per trial), a group level contrast between the amplitude modulation response of the valuation window (time points 6 and 8 seconds) is tested against zero. Figure D.7 shows the cluster-corrected result of this contrast for an amplitude modulator of SV.




Figure D.7 t-stat map of amplitude modulation by SV pursuant to valuation. Voxel-wise thresholded at $p < 0.001$, and cluster corrected to $p < 0.05$. The single cluster is 265 voxels, with a peak t-stat of 6.32, and a center of mass: $x = 1.3$, $y = 42.6$, $z = -3.3$ MNI space, LPI convention.

As shown in Figure D.7, the contrast reveals a robust positive deflection in activity, in a single large cluster spanning the vmPFC, as a function of offer SV during effort-based decision making. This cluster overlaps the a priori vmPFC nodes investigated in the main text.

Many of a priori nodes, by contrast, lie outside the vmPFC cluster, so it may seem somewhat surprising that they were not also identified as distinct clusters in the prior whole brain analysis. This may relate to multiple comparisons correction associated with whole brain voxel-wise analyses as the a priori regions did not survive a stringent whole

brain cluster correction. This fact is made clear by a subsequent whole-brain analysis utilizing in an uncorrected $p < 0.05$ threshold, which reveals the dACC, PCC, AI, and brainstem all showing positive amplitude modulation by SV. Hence, the lack of distinct clusters beyond the vmPFC likely reflects a Type II error stemming from low power. The voxel-wise analysis here thus provides additional evidence that SV encoding is particularly strong in this region relative to other regions.

Figure D.8 shows the averaged time series for each of the \$2, \$3, and \$4 offers, and also the averaged time series for each of the 2-back through the 6-back. Note that for the time points of interest (shaded in grey), there is a clear pattern reflecting both modulation by amount such that larger amounts correspond to more positive time series and modulation by load such that larger loads correspond to more negative time series. Hence, this region encodes both dimensions of choice contributing to SV.




Figure D.8 Averaged time series from vmPFC cluster encoding SV pursuant to valuation for each amount and load condition. The grey region corresponds to time points 6 and 8, during which significant amplitude modulation by SV was observed. * Indicates linear effect at $p < 0.05$, by time point, of amount in A and load in B as determined by multi-level model nesting observations within participants.

D.1.4.2 Independent Encoding of Offer Amount and Load During Effort-Based Valuation

As described in the main text, a region encoding SV during effort-based decision—making should independently encode both load and amount. The experimental paradigm included orthogonalized (fully-crossed) task load and amount regressors, allowing this hypothesis to be specifically tested.

Whole brain analyses were conducted by fitting GLMs where amplitude modulation of tent functions was predicted by parametric variation in centered amount, or, alternately, task load rather than SV. Next, group level contrasts of the amplitude modulation predictors at time points 6 and 8 seconds were tested against zero, and the resulting t-maps were thresholded (at voxel-wise $p < 0.001$), and cluster corrected (to $p < 0.05$). Table D.1 reports cluster sizes, locations, and peak t-stat values associated with each of these analyses.

As anticipated by the original analysis that revealed a vmPFC cluster when testing for SV effects which combine both amount and load, a largely overlapping vmPFC region was identified in both the amount and load analyses when conducted independently (Figure D.9). These new analyses thus confirm and corroborate the original analyses in demonstrating that the vmPFC robustly encodes both objective dimensions that are thought to be integrated in SV.




Figure D.9 Group-level t-stat map of activity parametrically modulated by first offer amount (left) and task load (right) at time points 6 and 8 seconds after trial onset. Cluster corrected at $p < 0.05$.

Intriguingly, the results also reveal a novel cluster of activity in the right dorsal anterior PFC that is robustly, and negatively modulated by task load but not by amount. This cluster was not predicted *a priori*. One possibility is that this region, which is also part of the DMN, plays a role in cognitive effort-based decision-making by representing expected cognitive load and conveying this information to the vmPFC where it is integrated with reward magnitude for calculating SV.

Anatomical Description	Size (voxels)	Center of Mass			Peak t-stat
		x	y	z	
<i>Amplitude modulated by offer amount</i>					
vmPFC	257	6.4	40.7	-2.9	5.78
<i>Amplitude modulated by task load</i>					
vmPFC	320	0.0	40.4	-4.4	-6.36
raPFC	183	20.7	40.0	44.9	-6.19

Table D.1 Anatomical description, extent, location, and peak voxel t-stat for amplitude modulation by either task load or amount of first offer at 6 and 8 seconds.

Again, it is possible to test the hypothesis of subjective encoding of effort costs in the regions modulated task load by testing whether the load effects during valuation in these regions correlates with AUC_{3S} . In neither region modulated by load, however, is the relationship between the slope of the load effects and AUC_{3S} significant (vmPFC: $B = -1.64 \times 10^{-2}$, $p = 0.42$; raPFC: $B = -4.33 \times 10^{-2}$, $p = 0.12$). Hence, this analysis does not support subjective encoding of effort in either cluster.

D.1.4.3 Regions involved in encoding SV variables during delay-based valuation

Unlike effort-based decisions described previously, no clusters survived correction for encoding SV during the evaluation window for delay-based decisions. This result was unexpected given the well-established encoding of SV in, among other regions, the very same vmPFC region identified for effort-based decision-making. The

lack of robust SV encoding again suggests diminished signal-to-noise in delay-based decisions. Nevertheless, by using the vmPFC cluster defined for effort-based decision-making as a mask, it is possible to ask whether the same region also encodes choice dimensions during evaluation on delay trials.

As shown in Figure D.10, it appears that there is weak encoding of both amount and delay during the evaluation window of delay-based trials. Interestingly, in both cases, the most pronounced encoding appears to occur earlier than it did with for effort-based decision trials: at 4 seconds after trial onset. This earlier encoding is consistent with the faster reaction times observed for delay trials and the tentative hypothesis that participants shifted decision strategies to a simpler heuristics for these trials.




Figure D.10 Averaged time series of all voxels within the cluster defined as encoding SV during effort-based decision trials, for delay-based decision trials. * Indicates a linear effect, a $p < 0.10$ of amount in A and delay in B, at each time point.

As with effort-based decision trials, there is also evidence that activity in the vmPFC region pursuant to evaluation during delay-based decision trials can also predict subsequent choice, taking into account biasing. Figure D.11 shows that activity is greater, at time points 8 and 10 seconds after trial onset, when participants ultimately select the

delayed reward, and particularly relative to the case when they select the immediate reward but were biased towards choosing the delayed reward.




Figure D.11 Averaged time series in the vmPFC cluster encoding SV on effort trials, for delay trials grouped by whether participants chose the delayed (first) or immediate (second) offer, and also by whether the immediate offer was designed to bias participants to choose it, or choose the delayed offer, with respect to each participants' own subjective indifference points. * Indicates a difference (at $p < 0.05$) by time point of delayed choice / immediate bias trials, and immediate choice / delayed bias trials.

Though the encoding of choice dimensions is weaker for delay than for effort-based decision trials, the similar patterns across trial types supports the broader hypothesis that the vmPFC region identified for encoding SV during effort-based decision trials supports domain-general encoding of SV, across both delayed and effortful rewards. The fact that activity in this region, pursuant to single-offer evaluation, is predictive of subsequent choice, and that prediction interacts with subsequent decision-bias supports the hypothesis that SV, computed in the vmPFC, drives subsequent choice rather than passively encoding a post-evaluative variable.