

Investigation of the Neural Mechanisms Underlying Cognitive Effort Avoidance

by

Z. Ceyda Sayalı

B.A., Koc University, 2011

A dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in Cognitive, Linguistic, and Psychological Sciences at Brown University

Providence, Rhode Island

May, 2019

© Copyright 2018 by Z. Ceyda Sayalı

This dissertation by Z. Ceyda Sayalı is accepted in its present form
by the Department of Cognitive, Linguistic, and Psychological Sciences as satisfying the
dissertation requirement for the degree of Doctor of Philosophy.

Date _____
_____ David Badre, Advisor

Recommended to the Graduate Council

Date _____
_____ Michael Frank, Reader

Date _____
_____ Amitai Shenhav, Reader

Approved by the Graduate Council

Date _____
_____ Andrew G. Campbell, Dean of the Graduate School

Zeynep Ceyda Sayali
Cognitive, Linguistic & Psychological Sciences, Brown University, Providence, RI
Phone: (401) 256-1792
e-mail: ceyda_sayali@brown.edu

Education

2013 - present Ph.D., Psychology
Brown University, Providence, RI

2006 - 2011 B.A. Double-Major: Psychology & Philosophy
Koç University, Istanbul, Turkey

Honors /Awards

2016 T32 Institutional Training Grant recipient in the Department of Neuroscience

2015 Friedman Family Lecture Forum Award

2013 Sheridan Center's Principles and Practice in Reflective Mentorship Program Award

2011 Vehbi Koç Scholar

2006 – 2010 Tübitak Scholar

Publications

Sayali, C., Badre, D. (2018). Neural Systems of Cognitive Effort Avoidance. *Neuropsychologia*.

Sayali, C., Uslu, E., Menceloğlu, M., Canbeyli, R., Balci, F. (2018). Effects of Acute Physical Exercise on Interval Timing. *Timing and Time Perception*.

Kilic, A., Sayali, Z.C., Oztekin, I. (2016). Aging slows retrieval of temporal information from working memory. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*.

Çoşkun, F., Berkay, D., Sayali, Z.C., Balci, F. (2015). Sequential Temporal Discrimination in Humans and Mice. *International Journal of Comparative Psychology*.

Çoşkun, F., Sayali, Z. C., Gürbüz, E., & Balci, F. (2015). Optimal time discrimination. *The Quarterly Journal of Experimental Psychology*, 68(2), 381-401.

Azeri, S., Sayali, C. (2012). Sign, Concept and Imagination: Vygotsky on Art as Conceptual Cognitive Activity. In Bourdier, V., Bréhaut, K., Dunphy-Blomfield, J., Frath, P. & Hilgert, E., (Eds.), *Res per nomen III : Référence, conscience et sujet énonciateur*. Reims: Éditions et Presses universitaires de Reims.

Posters & Presentations

Sayali, C., and Badre, D. (2018). Investigating the cost of cognitive effort. Poster presented at the 2018 Cognitive Neuroscience Society Conference.

Sayali, C., Badre, D. (2017). Investigating the cost of cognitive effort. Poster presented at the 2017 Annual Meeting of the Society for Neuroscience.

Sayali, C., Ciullo, B., and Badre, D. (2016). Neural systems supporting demand avoidance. Poster presented at the 2016 Annual Meeting of the Society for Neuroscience.

- Sayali, C., Ciullo, B., and Badre, D. (2015). Neural mechanisms of demand avoidance. Poster presented at the 2015 Annual Meeting of the Society for Neuroscience.
- Azeri, S., Sayali, Z., C. (2012). Sign, Concept and Imagination: Work of Art as a Conceptual Organ. *Article presented at the Res per nomen III*, University of Reims Champagne-Ardenne, Reims.
- Sayali, C., Özden, A., Demiray, B., Gülgöz, S. (2011). Second Reminiscence Bump in Turkish Old Adults. *Poster presented at the International Association for Cross-Cultural Psychology*, Istanbul.
- Sayali, C., Cebi, M., Laçin, E., Canbeyli, R. (2011). Interactive Effects of Water Maze Learning and Behavioral Despair in Male Wistar Rats. *Poster presented at the Society for Neuroscience*, Washington.
- Yavuz, M., Nasir, S., Sayali, C. Z., Batum, S., Özdamar, Ö., Moğol, S., Şen, H., Çorapçı, F. Yağmurlu, B. & Aksan, N. (2010). Interplay of emotional reactivity and maternal affect on Internalizing symptoms. *Poster Presented at the International Conference on Infant Studies*, Baltimore, MD.

Memberships

2018	Cognitive Neuroscience Society
2015 - 2017	Society for Neuroscience (SFN)

Teaching

2014 Fall	Cognitive Psychology: Prof. Steve Sloman
2014 Fall	Behavioral Psychology: Prof. Russell Church
2015 Spring	Cognitive Neuroscience: Dr. David Badre
2015 Fall	Clinical Psychology: Prof. Jack Wright
2016 Spring	Social Psychology: Prof. Bertram Malle

Work Experience

09/2011 – 05/2013	Lab Manager & Research Assistant of Dr. Fuat Balcı and Dr. İlke Öztekin Timing, Decision Making and Memory Lab, Koç University
-------------------	---

Research Experience

09/2013 – present	Grauduate student, Badre Lab, Brown University
09/2011 – 05/2013	Research Assistant of Dr. Fuat Balcı and Dr. İlke Öztekin Timing, Decision Making and Memory Lab, Koç University
09/2012 – 06/2013	Independent Project at Koç University, Philosophy Department <i>Principal Investigator</i> : Prof. Zeynep Direk <i>Project Title</i> : Levinasian altruism and neuro-ethics
07/2010 – 06/2012	Research Assistant at Koç University, Cognition Lab <i>Principle Investigator</i> : Prof. Sami Gülgöz <i>Project Title</i> : Second Reminiscence Bump in the Autobiographical Memories of Turkish elderly
01/2010 – 05/2011	Research Assistant at PsyBio Lab, Boğaziçi University <i>Principle Investigator</i> : Prof. Reşit Canbeyli <i>Project Title</i> : Water Maze Learning and Behavioral Despair in Male Wistar Rats

- 10/2010 – 02/2011 Independent Project at Koç University, Philosophy Department
Principle Investigator: Prof. Siyaveş Azeri
Project Title: Sign, concept, imagination: Art as conceptual cognitive activity
- 09/2008 – 05/2009 Independent Study, Developmental Psychology Lab, Koç University
Principle Investigator: Prof. Nazan Aksan
Project Title: Longitudinal study of mother-child relationship and child temperament on developing self-regulation.

Internship Experience

- 2008 Davranış Bilimleri Enstitüsü, Istanbul, Turkey
2008 Persona Psychological Counseling Center, Istanbul, Turkey

Mentoring

Brown University: Undergraduate Directed Research & Seniors Honors Thesis

- Louise Stolz
- Jordan Rubin-McGregor

Brown University: Undergraduate Directed Research

- Nada Hamzah
- Ria Vaidya

Brown University: *Women in Mind Mentorship Program*

- Elizabeth Lanzillo
- Carolina Santiago
- Emmaline Suchland
- Idara Ndon

Certifications

- 2015 Independent Functional Magnetic Resonance Imaging (fMRI) scanner
2014 Transcranial Magnetic Stimulation (TMS) training at Berenson-Allen Center for Noninvasive Brain Stimulation, Harvard Medical School

University Service

- 2017 Organized *Psychedelics in Science and Medicine Conference*, Brown University
2016 – 2017 Member of Stand Up for Graduate Students (SUGSE)
2015 Organized a science communication seminar *Coffee Break Science*
2013 – 2014 Started a mentorship program for women undergraduate students in STEM fields (*Women in Mind*)

ACKNOWLEDGEMENTS

This dissertation is dedicated to all the strong women in my life who taught me that effort is invaluable on the path to learning.

Table of Contents

CHAPTER 1: GENERAL INTRODUCTION	1
HYPOTHESIS SPACE OF COGNITIVE EFFORT AVOIDANCE	7
RESOURCE COST	7
PROCESSING COST	15
POTENTIAL PROBLEMS AND SOLUTIONS REGARDING THE TESTING OF COST ACCOUNTS IN THE LITERATURE	23
1) DIFFERENCES IN TASK CHARACTERISTICS	23
2) <i>DIFFERENCES IN TASK INSTRUCTIONS</i>	27
3) <i>DIFFERENCES IN THE TIMING OF THE MEASURES</i>	33
CURRENT RESEARCH PROGRAM	34
 CHAPTER 2: TESTING THE NEURAL SYSTEMS UNDERLYING COGNITIVE EFFORT AVOIDANCE DURING EFFORT SELECTIONS	 37
INTRODUCTION	37
METHODS	39
RESULTS	54
DISCUSSION	85
SUPPLEMENTARY TEXT	97
 CHAPTER 3: TESTING THE NEURAL SYSTEMS UNDERLYING COGNITIVE EFFORT AVOIDANCE DURING EFFORT LEARNING	 100
INTRODUCTION	100
METHODS	102
RESULTS	114
DISCUSSION	125
 CHAPTER 4: TESTING THE ROLE OF DOPAMINE STATUS DURING EFFORT LEARNING AND SELECTIONS IN PARKINSON'S DISEASE PATIENTS	 134
INTRODUCTION	134
METHODS	137
RESULTS	145
DISCUSSION	154
 CHAPTER 5: GENERAL DISCUSSION	 163
<i>TASK CHARACTERISTICS</i>	163
<i>TASK INSTRUCTIONS</i>	165
<i>TIMING OF MEASURES</i>	167
<i>PROCESSING COST</i>	172
<i>RESOURCE COST</i>	174
<i>FUTURE DIRECTIONS</i>	173
REFERENCES	185

LIST OF TABLES

TABLE 1.1.CLUSTER EXTENT (K) THRESHOLDS	50
TABLE 1.2.DESRIPTIVE STATISTICS	62
TABLE 1.3.DEMAND AVOIDANCE ACROSS BLOCKS	63
TABLE 1.4.CURVE FITS FOR INDEPENDENT ROIS	66
TABLE 1.5.WHOLE BRAIN ANALYSIS RESULTS	69
TABLE 1.6.CLUSTER COORDINATES FOR POSITIVE PCs.....	74
TABLE 1.7.CLUSTER COORDINATES FOR NEGATIVE PCS.....	75
TABLE 1.8. PRINCIPAL COMPONENT CLUSTER CORRELATIONS WITH DEMAND AVOIDANCE BEHAVIOR	82

LIST OF FIGURES

FIGURE 1.1. SCHEMATIC OF THE PARAMETRIC DEMAND SELECTION TASK. (A) EXAMPLE TRIALS OF THE CATEGORIZATION TASK WITH A DIGIT PRINTED ON A COLORED CIRCLE. THE COLOR CUED WHICH TASK TO PERFORM. THE RESPONSE MAPPINGS SHOWN BELOW THE DIGITS ARE FOR THE READER’S REFERENCE AND DID NOT ACTUALLY APPEAR ON THE SCREEN. (B) AN EXAMPLE OF SIX DECK CUES WITH THEIR ASSOCIATED EFFORT LEVEL PAIRING, BASED ON THE PROBABILITY OF TASK SWITCHING IN 13 TRIALS FROM THAT DECK. (C) SCHEMATIC OF BLOCK EVENTS DURING THE LEARNING PHASE. FIRST, A SYMBOL ICON IS PRESENTED FOR 1 SEC, AFTER WHICH THE PARTICIPANT COMPLETES 13 CATEGORIZATION TRIALS WITH THE ASSOCIATED CUE TILED ON THE BACKGROUND. (D) SCHEMATIC OF BLOCK EVENTS DURING THE TEST PHASE. IN THE “SELECTION” EPOCH, THE PARTICIPANT CHOOSES BETWEEN TWO SYMBOL ICONS THAT ARE ASSOCIATED WITH TWO DIFFERENT EFFORT LEVELS, WITH A DEADLINE OF 3 SEC. IN THE “EXECUTION” EPOCH, THE PARTICIPANT EXECUTES THE EFFORT LEVEL ASSOCIATED WITH THE SELECTED OPTION, WHILE THE SELECTED CUE IS TILED ON THE BACKGROUND.

40
56
57
59
67
70
71
72
73
79
115

FIGURE 12.1. DEMAND AVOIDANCE ACROSS PARTICIPANTS. 1A) OVERALL PROBABILITY OF SELECTING THE EASIER TASK. ALL PARTICIPANTS SELECTED THE EASIER TASK MORE THAN 50% OF THE TIME. 1B) PROBABILITY OF SELECTING AN EFFORT LEVEL ACROSS PARTICIPANTS. THE PROBABILITY OF SELECTING A TASK SIGNIFICANTLY DECREASED WITH INCREASING EFFORT LEVEL.

115
119
121
137
138
149
151
153

Chapter 1: General Introduction

Demand avoidance is operationalized as the preference for the easier task when given the option of what task to perform. This phenomenon can be measured in the laboratory using a paradigm called the ‘Demand Selection Task’ (DST). In these tasks, participants choose between two options that require different degrees of cognitive effort, usually manipulated by: the probability of task-switching; the number of items to maintain in memory; or, the frequency of conflicting stimuli. Despite individual differences, recent work (Kool et al., 2010; McGuire & Botvinick, 2010; Gold et al., 2015; Schouppe et al., 2014b) has established that most participants prefer the less effortful task, indicating that effort is aversive. The aversion towards effort has been coined by Hull (1943) as the ‘law of less work’ in psychology, and effort expenditure has been considered as ‘intrinsic disutility’ by economists (Baker, 1988).

The cost associated with effort has been a subject of economic theories in the form of labor-decisions for many years, since labor and value are exchanged bi-directionally. As labor is transformed into consumable goods, wages compensate for the amount of labor an individual is willing to supply. For example, influential microeconomics theories, such as Labor Supply Theory, predict the amount of labor an individual worker will supply in exchange for a certain wage rate (Becker, 1965; Heckman, 1974). According to these theories, labor is defined as the time and the effort required for obtaining the wage. This theory suggests that the worker aims to optimize the trade-off between the time he/she puts into labor to receive income and the time for leisurely activities, in order to maximize their utility function. The underlying assumption for this optimal balance between wages, leisure, and labor is that wages and/or leisure compensate for the cost associated with labor.

Although the intrinsic disutility associated with effort has long been introduced into economic models of decision-making, it was not until Conlisk (1987) proposed the cost associated with reasoning as ‘optimization cost’, prior to which most models had assumed that there was no cost to cognitive effort, i.e. decision-making. While physical effort has been the predominant focus in the field of psychological sciences, rather than mental effort costs, *Kool and Botvinick (2014) recently* investigated the explicit allocation of time on effortful tasks in a labor/leisure decision-task, showing that compensated wage decrease led to an increase in time allocation to the low demand task. Therefore, allocating time to leisurely activity was part of a cost-benefit calculation that depended on independent contributions of up-front and earned payments in a cognitive effort task.

The idea that cognitive effort enters into cost-benefit calculations as a ‘cost’ signal is further supported by effort discounting paradigms (Libedinsky 2013; Westbrook et al., 2013; Massar et al., 2015). In these tasks, parametrically titrated effort levels are paired with monetary reward for the participant to choose between. It was shown that subjective value of a task is reduced when the task to achieve the reward is high in effort. Furthermore, the change in the subjective value is due to individual differences in behavioral and biological measures of motivation and arousal. For example, apathy was observed to be associated with an increased discounting of monetary reward by effort, suggesting that the individual’s motivational trait influences cost/benefit calculations regarding effort (Hartmann et al., 2014). In a similar vein, subjective scores of willingness to expend effort, as measured by the Need for Cognition inventory, negatively correlates with effort discounting (Westbrook et al., 2013), suggesting that the immediate cost associated with effort is modulated by individual differences in pursuing mentally effortful activities in real life settings, such as problem solving or thinking.

The cost associated with effort has been further evidenced by changes in reward-related brain activity in neuroscientific experiments. It was argued that NAcc dopamine (DA) energizes the system to overcome the cost associated with physical effort to achieve reward, while leaving the reward value of the outcome untouched (Berridge et al., 1998; Evans et al., 2006; Salamone et al., 2007). In the same vein, DA depletion in NAcc leads to the avoidance of the high effort/high reward option, but does not affect the preference for the high reward option, when there is no physical effort required to achieve it (Salamone et al., 1994; Floresco et al., 2008). Furthermore, DA-depleted animals increase their chow-intake, in order to compensate for foregoing the more rewarding food (Font et al., 2008), suggesting that DA influences the cost/benefit calculations regarding physical effort. However, the role of NAcc is not specific to physical effort. Although the directionality of its effect has been recently challenged (Schouppe et al., 2014a), evidence from fMRI suggests that NAcc activity also scales with mental effort (Botvinick et al., 2009). For example, as NAcc showed increased activation to monetary cues rather than non-monetary cues, it showed a relative decrease in activation following the performance of high-effort tasks (Botvinick et al., 2009). This decrease in activation is interpreted as evidence for cognitive effort being registered as a loss in the brain.

Thus, there is growing interest in understanding the process by which organisms seek to minimize effort. These studies have aimed to confirm and measure the bias against cognitive effort in its relation to reward, motivation, and individual differences. Many, seemingly different accounts of demand avoidance agree on the idea that task demands are incorporated in the evaluation of task value, and are registered as something costly.

Only a few contemporary theories have explicitly defined what makes a cognitive task effortful. Furthermore, there have been no empirical designs to date that have directly compared

these theoretical accounts. The current review paper aims to fill this gap in the literature by identifying the hypothesis space for the cost of effort, and pitting the theoretical accounts against each other based on the existing evidence and premises each holds for the unit cost of effort.

First, we categorize the literature on the cost of cognitive effort into two broad categories: ‘Resource cost’ and ‘Processing cost’ accounts. In this section, we define the theories that make each account, and provide supporting and challenging evidence from the literature. An investigation of the literature suggests that the discrepancy between the findings might stem from differences in task characteristics. Next, we define these differences with examples from the literature and offer solutions for hypothesis testing, guidelines of which we utilize in our research program.

Due to research on cognitive demand avoidance still being in its infancy, there have been many varied attempts to explain the mechanisms underlying this behavior. These accounts, although distinct, are not mutually exclusive. This is mostly because the source of costs arising from cognitive effort is not as clear. For example, in a cognitively demanding task, such as in a two-digit arithmetic summation, one needs to transfer a digit to the next column, while actively maintaining the representation of the partial sum and inhibiting it from interfering with the carry operation. This means that an arithmetic summation task will require increased cognitive control, in order to resolve the increased conflict between task-representations, the cost of which will be partially observable in task performance, through indicators, such as response time and accuracy. The reason for why we avoid a two-digit arithmetic summation, lies, at least partially, in the internally monitored experience of control and increased response conflict, in addition to the averseness of impoverished performance measures, such as increased likelihood of error and increased time-spent-on-task. Finally, our meta-cognitive beliefs regarding the difficulty of solving math problems could also influence our effort judgements, regardless of whether the experience of this problem would be difficult or not. And, not surprisingly, each of these constructs separately has been proposed as the underlying cause of effort avoidance.

For the purpose of simplicity, we group the literature into two broad categories: “Resource Cost” and “Processing Cost” accounts. We define these accounts on the general basis of process-specificity. For example, what is shared between all Resource Cost accounts is that the mechanism that underlies effort avoidance is not process-specific, and that the cost of effort relies on a limited resource. In other words, regardless of the type of mental process the task employs, a change in defined task-parameters such as error rates, time-on-task, or average reward rate yield the same avoidance outcome. Thus, here we make the assumption that

‘Resource Cost’ accounts are unified under the characteristic that the source of effort costs is based on a depletable resource, and not the mechanism they assume to underlie effort avoidance. In that regard, processing accounts are different than Resource accounts by assuming that the process the task recruits underlies the cost of effort. ‘Processing cost’ accounts assume that the cost of effort is the recruited cognitive processes during task performance. In the following section, we provide supporting and challenging evidence for each cost account, and discuss potential ways for hypothesis testing. As some hypotheses are informed by the same task parameter/variable as another, the isolation of evidence into separate hypotheses is not always possible in current research programs, either because boundaries between hypotheses are not rigid or the current task designs are not well-adjusted to control for confounding variables. We discuss the latter in the following sections.

Hypothesis Space of Cognitive Effort Avoidance

Resource Cost

Resource accounts assume that the effort consumes a resource that gets depleted by various variables. The resource can be 1) a biological resource, 2) a motivational resource, or 3) performance-based.

The limited nature of effort exertion can be illustrated best in a phenomenon called ‘ego depletion’. Ego depletion is a concept that refers to the diminishing of mental resources after acts of self-control (Baumeister et al., 1994; Muraven et al., 1998). Self-regulatory behavior is an example of executive functioning, which, when compared to automatic behaviors, might be costly. These studies showed that when cognitive tasks that require self-control, such as affect regulation, preceded the execution of another cognitive task, participants showed decreases in performance in the following task (Hagger et al., 2010). The effect size across 83 studies that tested for ego-depletion was medium-to-large, and that the effects depended on task-characteristics such as task simplicity, the inclusion of interim periods between tasks, and participants’ beliefs about future control recruitment (although see Vadillo et al., 2016 for potential publication bias regarding these studies). The results were argued to support the idea that controlled behavior relies on a limited resource, and is depleted by use.

Some biological resource accounts of effort argue that the resource underlying this mental depletion is metabolic. Baumeister (1998) proposed that self-regulation requires energy that can be depleted by demanding tasks. This model suggests that controlled behavior leads to the consumption of glucose, the most important cellular energy source for bodily organs, including the brain (Gailliot et al., 2007). Although this account can be easily challenged by the recruitment of perceptual processes that increase glucose consumption more than the recruitment

of controlled processes, despite not leading to depletion effects (Kurzban, 2013). Moreover, it was later showed that the mere placement of a glucose-rich liquid in the mouth of the participant reversed the ego-depletion effects. This finding suggests that exertion of effort does not necessarily deplete limited biological resources, but perhaps change the allocation of motivational resources from the difficult task to the easier, or more rewarding task. For example, studies have shown that cognitive effort exertion (compared to experience of boredom) increased reward sensitivity, indicating that attentional resources are re-allocated following effort exertion (although a recent study showed that boredom also increases reward sensitivity; Milyavskaya et al., 2018). On the other hand, it could be that the reward magnitude retrospectively influences the perception of the difficulty of the task, and thus leads to an increased sense of entitlement following larger rewards. It was shown that when reward magnitude is mostly contingent upon the difficulty of the task, it alters judgments about how much effort was invested prior to the reward receipt (Pooremaeili et al., 2015).

A recent alternative biological resource model has been proposed by Holroyd (2016). However, unlike other biological models, this theory argues that it is not the fuel of control that is depleted, but rather it is the accumulation of toxic byproducts, such as Amyloid- β , that fuel consumption produces, which leads to demand avoidance. Accordingly, acts of control result in increased connectivity across distant brain regions, which, in turn, produce increased Amyloid- β production in neural tissue, a biological waste that is associated with neurodegenerative disorders such as Alzheimer's disease. Thus, demand avoidance is an adaptive response to minimize the long-term cost of accumulating waste products in the brain. The underlying assumption is that organisms have evolved an estimate of this long-term cost, which influences cost/benefit decisions regarding the performance of an immediate action.

At least one piece of supporting evidence for this account involves the manipulation of overall energy levels in relation to demand avoidance. *Libedisky et al. (2012)* investigated the influence of sleep deprivation on delay- and effort- discounting rates. In the relevant experiment, where both tasks were administered, participants were recruited for two sleep conditions on different days: sleep deprivation and restful wakefulness. The results showed that sleep deprivation lead to effort discounting but not delay discounting. These results are consistent with both ego-depletion models, and the waste-accumulation model (Holroyd, 2016). Specifically, Holroyd argues that sleep appears to replenish depleted resources, because the accumulated Amyloid- β are cleared from in between neural tissue during sleeping. In order to avoid the potential risk of over-accumulation, sleep-deprived participants could be more conservative about their resources.

While the above biological accounts address a global resource limit, the capacity limit could also be localized to the cognitive representations the task employ (Brzezicka et al., 2013). It has been suggested that control-dependent processes are computationally costly, because the role of control is in resolution of competing representations (Feng et al., 2014). Processes that rely on overlapping representations and processing pathways are subjected to cross-talk within shared resources. This cross-talk deteriorates behavioral performance, and leads to interference effects. As the complexity of the task increases, cross-talk increases within the local module (i.e. phonological or visual representations) that the task is performed under. The role of cognitive control is in the resolution of this cross-talk; however, the cost is not experienced at the level of a general control mechanism, but instead, at the level of local overlapping task representations. Thus, the cost of control is the resolution of the local conflict that arises from multitasking. Musslick (2016) showed that there is a tradeoff between developing shared task representations

and generalization. Accordingly, shared representations facilitate learning at the expense of diminishing performance during multitasking. Thus, a local bottleneck of shared representations might pressure the system to limit control recruitment. The readout of this local resource cost would then be apparent in diminished task performance, and increased neural activity in regions that are associated with conflict monitoring, all of which would be predicted to increase the likelihood to protect depleting resources, or in other words, effort avoidance.

Alternatively, motivational resource cost accounts such as ‘Opportunity cost’ accounts suggest that the cost of effort is the value of the foregone opportunity for being on task. A general framework that accounts for this trade-off is proposed by Kurzban (2013). Kurzban argues that the cost of effort is measured in terms of the foregone opportunity of what has not been attended to. The alternative task and its value depend on the task environment. For example, the presence of a smart phone in the study room will reduce the expected value of the task by increasing its opportunity-cost. Therefore, it is predicted that the presence of a smart phone will increase the likelihood of demand avoidance. Since the possibility of multitasking is lower during more effortful tasks, easier tasks might be selected because they allow for utilizing alternative tasks during time spent on task. Therefore, the cost of an effortful task is the subjective value of the alternative task. This account naturally incorporates the influence of time-on-task, as the number or the value of alternative tasks can increase with prolonged task engagement. Overall it can be predicted that if the cost of control is the value of the foregone task, then demand avoidance will positively correlate with the value of the alternative option. The value of the alternative task will scale with prolonged time on effortful task, where multitasking is not possible. In other words, if the opportunity cost of a task is low, the model predicts that resources are fully allocated to the task at hand. However, if the opportunity cost is

high, then the processing capacity is forced to divide between the current task and the alternative task in order to maximize reward, hence increasing the sense of effort.

However, one important question regards how these decisions concerning opportunity costs are made, since it is not always possible to know the exact value of the alternative task. Boureau and colleagues (2015) argue that the average reward rate and the controllability of the environment determine effort expenditure. The value of a task, therefore, is a function of both the average long-term reward rate of the foregone task over the duration the task is performed, and the advantage of executing control in a given environment over making random decisions. Thus, demand avoidance should increase with increasing reward-rate of the foregone opportunity and decrease with the value of control.

Niv (2007) argues that the average rate of reward (reward-per-unit-time) is what drives avoidance and approach decisions. The agent aims to maximize reward-per-unit-time. In order to do so, the agent might respond faster, although this might mean increasing energetic cost, because the cost of not receiving reward during the moment of waiting is higher. According to this account, the optimal response latency is inversely proportional to the reward rate of the task. If the reward rate is high, fast responses are optimal. However, if the reward rate is low, then response latency should be longer. In a context of high rate of reward, all actions including those that are not necessarily predictive of reward must be faster since the opportunity cost of time is higher. Therefore, the response vigor is closely tied to the expectation of reward in a given context. As reward-based learning has been shown to be recorded in terms of phasic dopamine bursts, tonic dopamine keeps track of the average rate of reward-per-unit-time. By reducing the average rate of reward, effort must reduce tonic dopamine levels in the system. Thus, the cost of effort can be tracked by tonic dopamine levels. A continuation of this theory suggests that if the

average-reward-rate is high, the likelihood of engaging in time-costly, and error-prone effortful tasks will decrease.

Therefore, this model (Cools, 2015) predicts that increased dopaminergic tone would lead to greater opportunity cost of time, thus greater effort avoidance and greater response vigor. For example, a test of this prediction could involve measuring demand avoidance rates in Parkinson's Disease (PD) patients. PD results in the death of DA neurons in the substantia nigra, a structure situated in the nigrostriatal pathway, so the mesolimbic 'reward' system remains relatively intact. Thus, DA medication in PD patients overstimulates the mesolimbic DA system, and leads to established alterations in reward-based decision-making. PD patients, who are on-medication would be expected to have greater opportunity-cost of time due to greater average reward-rate of time as signaled by greater tonic DA levels, and thus avoid effortful tasks more than healthy adults and participants off-medication.

However, the investigation of opportunity cost accounts is particularly difficult due to the value of latent variables, especially the value of the easier task can be conflated with latent reward factors such as the increased capacity to multitask during the easier task (Kurzban, 2013; Feng et al., 2014). This is because during easier task, fewer resources are occupied by the experimental task and attention can be divided between the experimental task and other rewarding tasks such as daydreaming.

Another possibility is that successful execution of the more effortful task is intrinsically rewarding, inflating the value of the effortful task. Lewis (1965) suggested that effort is related to the motivation level of the participant. In fact, he argued that the anticipation of effort increases the value of the end-state. The secondary reinforcer associated with high effort is preferred, compared to the reinforcer associated with low effort. Although he refrained from

making a definitive claim regarding why that would be the case, it is possible that increased arousal during effortful trials might increase the value of the subsequent stimulus. This idea is consistent with current empirical evidence.

A recent study showed that in the absence of monetary reward, voluntary selection of the more effortful task led to increased activity in the dorsal striatum (caudate) (Schouppe et al., 2014a). This result is interpreted as evidence for the intrinsic motivation during the selection of effort. Most real-life scenarios presuppose that the degree of effort is a good predictor of the attainable reward. Therefore, these results could also suggest that selection of the more effortful task(s) is accompanied by increased reward anticipation, or a heightened feeling of entitlement to reward.

Finally, the performance cost accounts of effort would argue that impoverished task-performance during the performance of difficult tasks yield an intrinsic cost associated with task performance as well. Effortful tasks naturally come with the cost of impoverished performance such as increased RT and error-rates, compared to easier tasks. As such performance measures can be adopted as a proxy to measure effort demands, the diminished performance itself might also be costly. There is an intrinsic relation between time and accuracy in many experimental tasks, known as the 'speed-accuracy trade-off'. More effortful tasks are associated with longer processing times to compensate for decrements in accuracy. They are also associated with increased error likelihood, specifically when the task also requires quick responses. Thus, effort cost is often confounded with delay cost and error likelihood. If the cost of control is time-on-task, then the avoided or discounted effortful task will be accompanied by longer task durations compared to the less effortful task. If the cost of control is error likelihood, then the avoided or

discounted effortful task will be accompanied by increased error rates compared to the less effortful task.

However, performance costs often index other cost measures. Increased time-on-task and error rates should also increase the duration of the control signal, which could be costly on its own or stand for the increased energy consumption of the brain. Additionally, increased time-on-task inevitably correlates with the frequency of alternative tasks one could instead attend to, and thus can also represent an opportunity-cost. Since performance costs usually correlate with effort manipulations, most studies attempted to control for their effects with effort avoidance.

One important challenge in testing for error avoidance is that although it is possible to control for error rates, it is not completely possible to control for the perception of increased error likelihood during the anticipation of more demanding tasks. It has been shown that medial prefrontal cortex neurons increase their activity in proportion to the error likelihood of a context during the cue period (Brown and Braver, 2005). This result suggests that the cue associated with increased effort could signal increased error likelihood, and thus be avoided, regardless of whether the average error rate is equated across tasks. However, this hypothesis was later challenged and the findings of the original study could not be replicated (Nieuwenhuis et al., 2007). It was shown that although medial prefrontal cortex shows increased activity in response to increased error rates, this effect was not generalizable to the cue period that predicted contexts with increased error likelihood. Instead, the activity change in this brain region was sensitive to a culmination of factors that included task difficulty, response conflict, negative feedback and the experience of errors.

Processing cost

Engagement in control processing is determined by the relevance of an action in a given context. The relevance of the action is determined as a function of its likelihood in reaching the reward state, which is a function of the costs and benefits of task performance. Thus, motivation/reward, and performance-monitoring, play crucial roles in the neural representation of controlled behavior.

Expected Value of Control theory ([EVC] Shenhav et al., 2013) formalizes the recruitment strategies for cognitive control, in which the expectation of the value of control determines its employment. This theory proposes that the current control signal is determined as a function of the probability of occurrence of the outcome states, with immediate reward estimated for performing them and subtracted by the intrinsic cost of executing the control intensity that is required to obtain them. This model accounts for all potential control costs that might arise from metabolic, structural and representational capacity limits during effortful task execution (Shenhav et al., 2017). However, regardless of the underlying capacity limit, the cost of control is assumed to reflect the need for protecting such limited resources and employing the intensity of control necessary to attain reward or avoid punishment. In congruence with this theory, individual differences in the recruitment of cognitive control during task-performance (i.e. switch-costs in a task-switching paradigm) positively correlated with individual differences in demand avoidance rates (Kool et al., 2010; Gold et al., 2015), suggesting that to the extent that an individual recruited controlled processes, they preferred the easier task.

From this account, it can also be derived that if the cost of control is control signal, then cognitive control related brain activity such as in PFC, will signal/accompany demand avoidance behavior. At least one piece of neurological evidence for this account comes from an imaging

experiment. *McGuire and Botvinick (2010)* aimed to investigate the neural basis of cognitive effort decisions by manipulating task-switching frequency. They showed that LPFC, a brain area that is implicated in top-down processing during task switching, positively correlated with the avoidance ratings. In a task-switching paradigm, in order to change tasks, a ‘task-set reconfiguration’ must take place before appropriate task-execution can proceed. Therefore, during this time, one has to shift attention between associations, retrieve goal states that are in line with this new task set, retrieve its rules from memory, and inhibit the rules associated with the previous task set. Thus, task-switching is an example of response conflict, the resolution of which requires cognitive control (Monsell, 2003).

Hence, the relation between response conflict and control recruitment has been argued to provide the basis for avoidance learning. The capacity for conflict detection has been proposed as the mechanism that monitors the increased need for cognitive control (Botvinick et al., 2001). For example, in a Stroop Task, the goal is to name the color of a printed word. In congruent Stroop trials, the color of the ink the word is printed in and its reading match (i.e. ‘green’ printed in green ink). In incongruent Stroop trials, the color and the reading do not match (i.e. ‘green’ printed in red ink). Assuming that there are two representations of a word, namely the color of the ink and the reading of the printed word, incongruent Stroop trials lead to the coactivation of competing representations (the reading ‘green’ and the color, red). The role of cognitive control is in resolving this competition by strengthening the connection between the input and the color representation. Accordingly, increased error rates, as well as increased task demands during an incongruent trial, could both signal the need for controlled engagement to the current task and the increasing negative value of the task. Botvinick and colleagues argue that the conflict monitoring function of ACC allows this brain area to detect the degree of required control to

resolve conflict, while simultaneously informing the BG of the updated value of the task via the registration of conflict as a cost. Similarly, EVC (Shenhav et al., 2017) also predicts that ACC incorporates the cost associated with control, whether the required demand is instantiated by conflict, surprise or error likelihood.

In that respect, it is important to test the cost hypothesis during the presentation of reward receipt following the execution of effort, where the history of action could influence the evaluation of the outcome. *Botvinick et al. (2009)* investigated the activity in the reward circuitry of the brain following effort execution. In this experiment, the behavioral task consisted of low and high demand tasks that were categorized according to the proportion of trials that required task switching. The results showed that effortful blocks were accompanied by greater dACC and a reduced Nacc activity. There was a significant inverse relationship between ACC activation during task performance and NAcc activation during the reward cue, indicating that ACC tracks task difficulty, probably as indexed by the degree of conflict, which is processed as a cost in calculating reward in NAcc.

Accordingly, it has been claimed that ACC inhibits dopamine release during conflict processing (Botvinick et al., 2004), possibly via its connections with NAcc (Botvinick et al., 2009). Through experience, the reward value of an action is updated via reward prediction errors (Schultz et al., 1997; Schultz & Dickinson, 2000; Dayan et al., 2002). The reduction in the probability of reward, in turn suppresses the activity of the dopaminergic neurons in VTA, leading to a negative prediction error. The change in the dopaminergic activity alters the plasticity of the associated action representations. Those representations of actions in PFC that consistently lead to negative prediction errors in VTA get weakened, while those that lead to positive prediction errors are strengthened (Hyman et al., 2006; Shen et al., 2008). Thus, the

organism learns to avoid those actions that lead to negative prediction errors, and maintain goal-relevant actions in order to maximize reward. The transient pause in dopamine bursts experienced during control demanding, conflicting events, leads to avoidance learning via dopaminergic frontostriatal loops, and consequently to the habit of demand avoidance.

In line with this model, Cavanagh and colleagues (2014) have shown that cues that are associated with high effort induce increased coupling between punishment related mPFC theta power and cue avoidance, and decreased coupling between reward-related theta power and cue selection. Frontal theta power has been associated both with the need for an increase in controlled processes as well as prediction errors that could inform action selection (Cavanagh and Frank, 2014). These results suggest that high-effort context cues might drive demand avoidance behavior via inducing a negative reward prediction errors. The subsequent decision to select from the effortful task despite the negative value of effort might require increased top-down motivation to overcome the adaptive response to avoid effort.

Further supporting evidence for the role of dopamine in effort allocation comes from physical effort studies. It was shown that Parkinson's disease patients, who suffer from striatal dopaminergic loss, increase willingness to expend physical effort when on-medication. Parkinson's Disease medications have been shown to overstimulate the mesolimbic dopaminergic reward pathway. This suggests that increased dopaminergic tone yields reduced sensitivity to the cost associated with effort. These results can be mapped onto the cognitive effort literature under the assumption that control recruitment during cognitive effort exertion registers as a cost. However, as described in the previous section, Opportunity Cost theories of dopamine makes the opposite prediction, where increased DA tone will lead to greater average rate of reward, and thus lead to greater opportunity-cost of time during effort decisions. As

cognitive effort avoidance in PD patients awaits empirical testing, it promises to be an important avenue for contrasting Opportunity Cost and Processing Cost hypotheses of cognitive effort.

In the literature, the Processing Cost accounts has been equally challenged and supported. The disagreement between the results of the studies stems mostly from: (1) a lack of consistent neural evidence for brain areas that are associated with conflict monitoring; and, (2) a lack of control for the cost associated with performance that do not solely signal the required control intensity.

Performance Cost accounts specifically pose a challenge for fMRI studies that discuss the role of ACC in conflict resolution. As well as response conflict during effortful tasks, uncontrolled ERs could also be represented by increased ACC activity following effort exertion. A study by Brown and Braver (2005) contrasted the effects of response conflict and error-likelihood through computational neural modeling. The model incorporated the role of mid-brain dopaminergic signals in reinforcement learning, whereby errors lead to a suppression of phasic dopaminergic activity. The neural model was presented with two contexts that were associated with low- and high- error-likelihood. Following the presentation of the context cue, the model solved for a task that either involved response conflict or not. Correct trials in high error-likelihood conditions led to increased ACC activity compared to correct low error-likelihood conditions, even when the response conflict across conditions was equated. Those contexts that were associated with low error-likelihood showed increased dopaminergic suppression and increased ACC activity following errors, consistent with the teaching role of dopaminergic activity in reward-prediction errors and avoidance learning (Frank et al., 2004; Hazy et al., 2010). This result, although Nieuwenhuis and colleagues (2007) could not replicate it later, raises concerns about whether increased ACC activity in response to high-effort trials might signal the

anticipation or the experience of high error-likelihood contexts, rather than the degree of response conflict.

Additionally, the Psychological Entropy model (Hirsh et al., 2012) argued that entropy increases with time, which could raise a potential problem, specifically, for the effortful option. Entropy related uncertainty is also argued to activate ACC, which could be due to the effects of prolonged time-on-task during effortful tasks. The model also emphasizes that it is not necessarily the response conflict that increases entropy, but rather it could be the individual's own estimation of whether conflict resolution is possible, given limited resources. Therefore, it is possible that the individual's performance costs share an interaction with the level of response conflict during an effortful task. This would suggest that entropy related anxiety would be higher for those individuals who perceive the successful effortful task performance to be unattainable, due to prior bias about the task, overall fatigue levels, or state anxiety.

One challenging finding for the Response Conflict account across some studies was that the level of response conflict was not always predictive of the level of executed cognitive control (i.e. the lack of a relation between PFC recruitment and ACC activity in Botvinick et al., 2009; the lack of ACC activity in McGuire & Botvinick, 2010). Response conflict could signal the extent to which control is required to resolve it, as well as the representation of a negative stimulus. This would suggest that response conflict cost could further be defined as: 1) the intrinsic cost of executing control; and 2) the cost of experiencing a negative stimulus.

Botvinick and colleagues (2004) argued that these two components overlap in ACC's conflict-monitoring and avoidance learning functions. However, if these two functions are not necessarily unified, different individuals might associate separate costs to either process. In order to shed light on an individual ability to monitor and adaptively adjust cognitive control

requirements following response conflict, individual control scores could be obtained using independent tasks. Consistent with this concern, Schouppe et al. (2014b) showed that conflict detection and subsequent control adjustments might be decoupled in deriving the cost of effort. For example, incongruent Stroop trials are followed by adaptive increases in cognitive control, which can be detected by behavioral changes in task performance (CSPS effect). However, Schouppe et al. (2014b) showed that CSPC effect did not predict demand avoidance as individual CSPC scores did not correlate with low-conflict preference rates, indicating that there was no relation between the extent to which participants executed cognitive control after response conflict and demand avoidance. Therefore, this study challenges the relation between cognitive control and response conflict in explaining the cost of effort.

Additionally, the negative quality of conflicting events has been recently supported (Dreisbach and Fischer, 2015). It was shown (Dreisbach and Fischer, 2011) that performance during incongruent Stroop trials following negative targets, such as the presentation of a word of negative valence (i.e. ‘murder’), was facilitated compared to incongruent trials that follow positive target events, or congruent trials that follow negative targets. These results indicate that negative affect-regulation and conflict processing might share similar underlying mechanisms that consequently motivate adaptive control strategies.

Finally, the distinction between reactive and proactive modes of cognitive control would inform the differentiation between response conflict and control signal accounts. As response conflict accounts predict a reactive control strategy, in which PFC activity is triggered at the time of the stimulus, control signal account would predict a sustained PFC activity, which onsets at the effort cue. Considering that reactive and proactive control influences the temporal dynamics of PFC recruitment, the cost of control might vary as a function of neural employment.

Resource accounts would predict an increase in metabolic consumption during proactive control engagement, since sustained neural activity naturally employs more resources than transient activity. Further, increasing task conflict leads to a shift from reactive control to proactive control strategies. In other words, highly effortful tasks require proactive control strategies that easier tasks do not. This suggests that empirical designs that test high- versus low-effort conditions might not only manipulate task difficulty, but also the mode of recruited control, the additional cost of which it inflicts is yet unknown.

Additionally, the timing or the absolute recruitment of proactive control might vary across individuals, which might further explain differences in demand avoidance rates. Consistent with this concern, Botvinick et al. (2009) showed an absence of PFC-NAcc relationship following effortful task execution. PFC, which is a brain area indicated for top-down control, did not directly correlate with the extent of discounting. Instead, dACC is argued to be involved in both conflict monitoring and the reactive adjustments in cognitive control intensity (Shenhav et al., 2013). If the role of dACC was in the allocation of cognitive control, then its direct relation to discounting, in the absence of PFC's executive role, suggests that the cost of effort is dissociated from the actual cost of the execution of the specified control. This may possibly be because the demands of the control signal and the executional capacity do not always directly overlap, i.e. both response times and error rates increase during effortful blocks, indicating that the loss in accuracy cannot be compensated with implemented caution in these trials. While there might be limits to how much one can increase control intensity, the specification of the required control signal to achieve optimal performance is tracked and used to discount the value of reward. This is interesting, because it indicates that the monitored conflict is driving the cost/benefit calculations without requiring direct information from the cost of

immediate control execution; perhaps, this is done by calculating the upcoming cost based on the functional form of the cost function defined by past experience or by subjecting the immediate execution signal to subsequent transformations.

Potential problems and solutions regarding the testing of cost accounts in the literature

The lack of conclusive evidence for any cost hypothesis might be due to three reasons: 1) differences in task characteristics; 2) differences in task-instructions; and, 3) differences in the timing of the measures.

1) Differences in task characteristics

a) Type of cognitive task

Among the studies in the existing literature, some adopted a task-switching paradigm to manipulate effort demands, some employed an N-back memory task, some employed backward typing, some adopted a Flanker Task, some adopted a Stroop Task, and some adopted an arithmetic task. As a result of the ongoing debate about the nature of cognitive control, there is no consensus on whether executive processes are unitary or not. One of the first definitions of executive control by Baddeley and Hitch (1974), defined the central executive as a mechanism that controls visual, phonological and episodic components. Although this theory was argued to postulate a homunculus, an overview of cognitive control processes highlighted various functions of executive mechanisms. Further neuroscientific evidence, starting with frontal lobe patients, helped to locate the executive processes in the brain, while potentially studying the relation between these seemingly disparate functions. The view on the modularity of executive functions was mostly supported by the exploration of the brain. It was observed that different

parts of the brain supported different subcomponents of controlled processes. While functions of planning and goal selection were located in dorsal PFC and subcortical structures, attentional control was mostly located in the parietal cortex.

The localization of separate functions is not necessarily enough to support the modular view of executive control, due to high inter-connectivity between these brain areas. However, it underscores the possibility of testing separate components of cognitive control utilizing the employment of different tasks. For example, Stroop Task and Flanker Task conflicts are shown to occur both at the perceptual and response level (De Houwer, 2003; van Veen and Carter, 2005). The unaccounted cost associated with perceptual conflict, makes it difficult to translate these findings to the domain of an arithmetic multiplication task, which does not involve perceptual conflict. Therefore, using separate cognitive tasks to manipulate cognitive effort might lead to translational issues between the results, since it is unknown whether the cost associated with cognitive control is unitary or not.

In the brain, ACC and PFC were shown to be dissociated in response to distinct forms of conflict, namely response and stimulus (Liston et al., 2006), suggesting that according to the form of conflict different brain areas could be employed. The employment of disparate brain areas in response to the same construct of task difficulty poses a critical challenge for the generalizability of the findings. For example, during a hierarchical task where there are multiple task dimensions, the first task dimension requires the association of one dimensional response-task rule mappings, such as the association between one color and one key response. Meanwhile, the second level requires the combination of two different levels, such as the color-response mapping being contingent upon the shape of the stimulus. In this case, different brain regions might be executed at different levels of the task.

Additionally, the strength of this relationship across task dimensionality might interact with the localization of the activated brain region. For example, if both BOLD and RT were mediated by a latent variable, such as difficulty, and each variable had a different functional relation to any of the other variables, a mixed picture might emerge. If the relation between RT and this latent variable is non-monotonic - meaning response times might increase or decrease in response to increasing difficulty - even when RTs are equated across different tasks, they might not be matched in terms of difficulty. Therefore, the relation between difficulty and BOLD activity might be task-specific and region-specific, while the interaction of these contingencies might result in linear or nonlinear effects across parametric levels of the same difficulty construct.

These translational problems between behavioral and neural findings constitute a problem for the determination of the cost of effort, since the cost might vary across different components of cognitive control, as well as the neural mechanism the task employs. In turn, the intrinsic cost or the cost function associated with these cognitive or neural mechanisms might lead to non-linear effects on behavior or neural response.

b) Presence of external rewards

Additionally, the literature so far had inconsistently tested for demand avoidance in the presence or absence of external rewards reward (Botvinick et al., 2009; Schouppe et al., 2014a; Cavanagh et al., 2014; Frobose et al., 2017). However, it was shown that voluntary selection of a harder task might yield an intrinsic value in the absence of external reward (Schouppe et al., 2014a). Specifically, the term ‘intrinsic motivation’ was first used by Harlow (Harlow, 1950) in relation to his observation that rhesus monkeys spontaneously inspected and manipulated complex puzzles that were placed in front of them, in the absence of food, water, or any other external reward. During these manipulations, they improved their task-performance to the point

that they became fully competent at the task. Similar reports across species have shown that this internally guided spontaneous behavior to seek cognitive effort was associated with greater learning, task performance, and positive mood (Ryan & Deci, 2017), a tendency that was undermined by the presence of external rewards (Deci, 1971; Hewett & Conway, 2016; Hidi, 2016). On the other hand, an emerging literature of cognitive effort avoidance postulated that cognitive effort exertion increases negative affect (Kurzban et al., 2013) and yields an intrinsic cost that reduces the reward value of the task (Kool et al., 2013; Shenhav et al., 2013; Westbrook et al., 2013). Cognitive effort discounting procedures (Westbrook et al., 2013; Massar et al., 2015; Apps et al., 2015) consistently show that people demand greater monetary rewards to solve difficult tasks compared to easy tasks, when given the option. These seemingly disparate observations regarding cognitive effort suggest that the reward context under which cognitive effort is exerted might influence the evaluation of effort.

c) Estimation of neural activity

Finally, the fMRI literature on cognitive effort avoidance so far has adopted a subtraction approach. This subtraction approach does not control for a variety of factors besides effort that might categorically differ between high and low effort conditions or other continuous variables that might correlate with effort, but are not collinear with it. For example, in their design, Botvinick et al. (2009) adopted two effort levels. The low-effort task yielded trials that included no task-switching, while the high-effort condition yielded tasks that alternated on every trial. It is possible that the low-effort task is different than that of the high-effort task not only in terms of effort exertion, but also in the relative boredom it might yield due to a lack of control demand in the low-effort condition.

Furthermore, this design does not permit full characterization of the response function

relating effort, as a manipulated task variable, to brain activity. Thus, it was unknown whether the aforementioned networks were exclusively involved in tracking effort costs versus other factors. A parametric manipulation of effort can help mitigate these issues by testing a function that changes in a dose-response fashion over effort levels.

2) Differences in task-instructions

In the existing literature, while some of the studies refrained from instructing the participants about effort manipulations, others explicitly informed them of effort differences between options. Such a difference in the administration of the effort-based decision tasks constitutes a problem in the comparability of the findings. As such, explicit instructions could subject participants to beliefs and biases about effort, which the implicit condition would not.

The importance of task instructions is also underscored in *Westbrook et al. (2013)*. In this task, participants were explicitly informed about the cognitive control manipulation. This opens the possibility that the cost of cognitive effort might be confounded or influenced by beliefs individuals hold about cognitive effort. Through personal communication with the author, it became apparent that the individual COG-ED scores were not correlated to DST scores; the latter was a paradigm administered without explicit instructions about effort. The lack of a significant correlation creates a problem in translating demand avoidance results found by COG-ED, when viewed from the perspective of avoidance learning, and the lack of correlation suggests that these two tasks are not measuring the same construct. This might be because participants in the explicit task instruction condition can monitor effort differences between options easier than the implicit task instruction condition.

For example, Gold et al. (2015) failed to replicate the DST findings and showed that that switch-costs in a DST task-switching paradigm did not correlate with avoidance rates across

individuals. These initial results suggested that demand avoidance could not be explained by control costs. Given these unexpected findings, the researchers decided to alter the task instructions. Usually in DST, task instructions do not involve an explicit statement about how two options differ from one another. Instead, participants start the task by sampling from both options and show increasingly higher demand avoidance rates over the course of trials (Kool et al., 2010), implying that task values are learned and inform decisions in an online fashion. In the second experiment, the instructions were made explicit with regard to the manipulations of cognitive control. The experimenters specifically told the participants that one option involved increased task-switching compared to the other. The rest of the instructions were the same: the participants could freely choose between options. It was hypothesized that the lack of demand avoidance in the patient group was due to a lack of differentiation between the easy and the difficult option. To examine the issue, participants first completed this modified DST. Then, they were instructed to ‘figure out’ the harder option and select from it. In the following task, participants were instructed to ‘figure out’ the easier option. In this experiment, DST choices were similar and above chance for both the patient and the healthy group. Switch-costs were negatively correlated with avoidance rates only in the healthy group, suggesting that increased control explained the cost of effort in healthy participants only. It was important to note that patients showed significant difficulty identifying the harder and the easier options. This suggests that patients could not form an explicit idea about option difficulty.

Moreover, it was shown that general IQ scores positively correlated with demand avoidance rates in both patient and healthy groups. General IQ tests intend to measure adaptation to the environment and of the ability to learn (Sternberg et al., 2005). These results imply that proper utilization of conflict information might be dependent on the awareness of difficulty, in

addition to learning the value of each option, in order to accurately inform cost/benefit calculations regarding effort. Additional multiple linear regression analysis tested for the individual contributions of effort detection and IQ on demand avoidance; while the total effect of both factors was significant, neither survived significance on its own.

These results show that although switch-costs were higher in the patient group, they were not sufficient to modulate demand avoidance behavior. In fact, demand avoidance is only observable when common healthy and diseased populations are explicitly instructed about effort manipulations. This indicates that in the absence of explicit effort monitoring, demand avoidance might not be possible. In fact, demand avoidance is shown to be dependent on explicit effort monitoring, and the capacity to correctly estimate the reward value of an option. Thus, the results of this study argue against the processing cost account and, highlight the possibility that explicit effort monitoring could be mediating the relation between effort avoidance and cognitive control.

The significant value of difficulty monitoring on demand avoidance can be illustrated by the experience of a left medial frontal patient, with lesions including the left ACC (Naccache et al., 2005). The patient was subjected to various Stroop experiments during skin conductance response (SCR) and electroencephalogram (EEG) recordings. In standard Stroop paradigms, the patient showed increased RTs and ERs for incongruent trials compared to congruent trials, like the healthy controls. Additionally, she showed the Gratton effect, which refers to the increase of RTs for congruent trials and a decrease in RT for incongruent trials following incongruent trials. This standard Stroop effect also correlated with event-related potential (ERP) recordings from the right ACC activity, an area that remained intact despite the patient's lesions on the left side. This effect is considered to be an example of increased executive control following conflicting

trials. Therefore, the patient showed both intact executive control and executive monitoring capacity.

Another experiment assessed the subjective feeling of effort following both the congruent and the incongruent trials. Both the healthy controls and the patient were presented with Stroop pairs that differed in difficulty and were asked to verbalize which of the pairs was more difficult. While the normal participants were able to categorize the difficulty correctly, the patient was at chance-level. The authors also explored SCR responses to both congruency types, with the expectation that incongruent trials would result in an elevated SCR response. While the normal participants showed significant fluctuations in SCR responses on a trial-trial basis, with higher SCRs after incongruent trials, the patient's SCRs were flat across trials. Importantly, the patient was able to monitor her own performance in terms of task duration and adjust control requirements accordingly. Thus, the difficulty of the task correctly informed control specification without the explicit knowledge about task difficulty.

The patient's impairment was restricted to evaluating effort. In order to compare her performance on an independent task, she was also given an Iowa Gambling Task, in which she was presented with 4 different decks that resulted in either monetary loss or gain. Although she reported that disadvantageous decks were 'tricky' and that she should not be selecting cards from those decks, she kept picking from the disadvantageous decks during the task. Interestingly, there was no SCR anticipation observed prior to selecting from these decks either; indicating that the SCRs associated with effortful tasks in normal patients could mark the negative value of demanding tasks.

These results indicate that there is dissociation between the explicit knowledge and the behavioral response regarding negative outcomes. This study does not test for demand

avoidance; however, it reports the dissociation between executive control and conscious feeling of effort. The dynamic control adjustments made after difficult trials do not necessarily map onto an explicit knowledge of difficulty. This lack of awareness is paralleled by a lack of preference for monetarily advantageous options, suggesting a common evaluation mechanism underlying judgments regarding effort and the implicit learning of the avoidance of negative outcomes. This explanation questions the validity of Cog-Ed in testing demand avoidance behavior in the absence of explicit beliefs or meta-cognition about difficult conditions. In that respect, the observed results could be due to prior beliefs about how to behave in difficult conditions, instead of the experienced cost of effort during the experiment.

However, the implications for the implicit nature of DST, regarding demand avoidance, might also be equivocal. In this task, participants are given two options to select from, without prior experience with the association between the context and the task difficulty. Therefore, the participants' job is to determine which one of these contexts is more rewarding, and make a decision in the absence of explicit feedback or incentives. In that respect, DST becomes a reinforcement learning problem in which the agent learns how to map the action to the context, in order to maximize the reward signal. This is evidenced by the increasing probability of selecting the easier task across trials. As the participants start without a preference for any of the contexts, they progressively develop a bias for selecting the easier task through increasing experience with the task. Another characteristic of reinforcement learning is that the agent needs to optimize the trade-off between exploration and exploitation of the options, in order to get a reliable estimate of both options' reward value. This entails that the agent samples from both options enough times and demonstrates a stochastic history of preference on a trial-trial basis. A similar selection profile is evident in DST. Although the participants gradually learn that one of

the options is easier to perform, they continue selecting from the harder option throughout the task, although with a decreasing probability.

The final question, then, concerns the basis on which the reward probabilities of the options are estimated, in the absence of explicit reward. The dopamine reward prediction error hypothesis assumes that the value of an option gets updated via phasic dopamine signals on trial-trial basis with unexpected changes in the value of an option. The strength of the dopaminergic burst scales with the difference between the predicted reward and the actual reward signal. For example, positive reinforcement of an action that originally was predicted to not yield reward will yield a greater positive prediction error than an action that was originally predicted to yield reward. As the reward probabilities of the options are updated via reward prediction errors, the better option is selected to the extent that the difference between the reward probabilities of two options is high. Before the administration of DST, participants practiced a set of task-switching blocks until their within-block accuracy reached 90%. The RTs for the practice phase were not reported; however, it was shown that practice improves task-switching performance across test blocks, specifically in easier task sets (Kray & Eppinger, 2006). While the accuracy for both the easy and the hard options reached around 98% at the end of the experiment, the average RT for the easier deck remained faster than that of the harder option. Assuming that in the absence of explicit feedback, participants could internally generate the reward signal via monitoring of their RTs, the enhanced improvement in the easier task performance across trials could stand for a greater reward prediction error that selectively reinforces the easier option compared to the harder one. In corroboration with Gold and colleagues' (2014) findings, this explanation suggests that the observed avoidance results in DST could also be explained by performance monitoring

as well. This assumption can be tested via modeling the effects of trial-trial reward-prediction errors with changes in performance measures on selecting the easier task.

3) Differences in the timing of the measures

Among the existing few fMRI studies, some of them took measures only following the execution of the effortful task, while the others took measures only during the selection of effort. One of the studies observed that activation to reward in the NAcc decreased with increasing effort to obtain that reward (Botvinick et al., 2009). This was taken as evidence that effort reduces the reward associated with performing a task. Another study observed that selecting an effortful task (i.e., anticipation of upcoming effort) also activates NAcc (Schouppe et al., 2014a). However, in this study, NAcc increased its activation with greater effort, which is at odds with the anticipation of a reduced reward. More importantly, these studies tested effort cost in different experimental phases: either following the experience of effort/reward or during the anticipation of effort.

However, it is possible that these two facets may rely on distinct neural mechanisms or distinct response functions. This concern is warranted in the literature. It was shown that (Sakai et al., 2003) the anticipation of the following memory task yielded activity in brain areas that are both domain-dependent and –independent, in terms of the task requirements. Domain-independent brain activity during task execution was greater in posterior PFC compared to anterior PFC, when compared to its anticipation. While domain-specific posterior areas increased activity during task execution compared to anticipation, anterior PFC did not. This interaction suggests that specific parts of PFC might be exclusively involved in the anticipation of task demands. This idea is corroborated by other studies (Koshino et al., 2011).

It was shown that that medial PFC yields activity during the anticipation period, and the

same region is deactivated during execution. Similar results have been found in the domain of pain perception (Ploghaus et al., 1999), which is considered relevant, given the parallel between cognitive effort and the experience of aversive stimuli, in the literature. It was found that while the anticipation of pain activated regions in the medial frontal cortex, the experience of pain activated caudal ACC. This dissociation among various regions in PFC between the anticipation and execution/experience periods suggests that the differences in brain activity during selection and execution of effort might stem from PFC's role in task preparation. Thus, measuring brain activity during both the experience of effort and its anticipation is required to establish the neural systems underlying effort-based decision-making.

Current research program

The goal of our research program is to test the cost accounts of cognitive effort while addressing the aforementioned problems in the literature. In order to address the potential problem that might arise from differences in task characteristics, we developed our experimental design based on the original DST task. Thus, across all experiments, we adopt a task-switching paradigm to manipulate effort, and assign different effort tasks to separate contexts the participants are asked to choose between. Across all experiments, we avoided the inclusion of external rewards and feedback in order to directly test the intrinsic cost/value associated with cognitive effort.

First, in a behavioral study, we parametrically manipulated effort levels in DST, and show that sensitive demand avoidance measures can be obtained using this parametric DST paradigm (Chapter 2). In all imaging experiments (Chapter 2 & Chapter 3), we use the same parametric DST design that allowed us to observe brain activity change in a dose-dependent fashion. In order to address the potential role of task-instructions in developing effort-bias, we do

not provide explicit instructions regarding difficulty differences between effort options and instead rely on implicit learning mechanisms. We test this implicit learning assumption by utilizing reinforcement learning models in explaining learning the cost associated with effort (Chapter 3) in demand avoidant participants, and show that effort selections can be explained by reinforcement learning mechanisms.

Second, in order to avoid potential problems that might arise from differences in the timing of brain measures, we separate selection and execution epochs of effort in fMRI (Chapter 2), and show that Selection and Execution of effortful tasks recruit distinct brain areas. These results confirm the concern that the discrepancy in the effort literature might be due to acquiring brain measures at different epochs of effort-based decision-making.

Third, one of the least explored areas of research in the cognitive effort avoidance literature is individual differences. While most studies report that majority of participants avoid effort, they consistently show a group of demand seeking participants. In order to understand what brain or behavioral measures underlie differences in effort evaluation, we test a relatively large group of participants in fMRI. In Chapter 2, we show evidence for both demand avoidant and demand seeking individuals, for whom similar resource and processing measures incur opposite values, indicating that a general cost of effort framework might be too simple to explain individual differences in cognitive effort expenditure.

Finally, we test the causal role of reward-related brain activity in modulating effort-based decisions, in order to directly test the hypothesis that cognitive effort signals an intrinsic cost that can be remedied by dopamine administration in the brain (Chapter 4) and show that unlike the energizing role of DA in physical effort, diminished DA leads to greater effort seeking behavior in Parkinsons' Disease patients. As a conclusion, models that favor the intrinsic value of

cognitive effort are incorporated into current models of cognitive effort expenditure in order to explain the observed results across all studies.

Chapter 2: Testing the neural systems underlying cognitive effort avoidance during effort selections

Introduction

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

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

A second question concerns what makes a task cognitively effortful. An influential hypothesis proposes that tasks are effortful to the degree that they recruit cognitive control. Thus, according to the “cost of control” hypothesis, cognitive control might be recruited to the degree

that its associated costs do not exceed its anticipated benefits (Shenhav et al, 2013). Consistent with this hypothesis, people avoid tasks that require more task switching (Kool et al., 2010), greater working memory load (Westbrook et al., 2013), and greater response conflict (Schouppe et al., 2014b), all of which putatively require greater cognitive control.

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

Accordingly, in the present study, we sought to provide a more sensitive test of the brain-behavior relationships predicted by the *cost-of-control* hypothesis. Specifically, we developed a parametric version of the Demand Selection Task (DST) paradigm for fMRI to sensitively test how incremental changes in activity and connectivity levels in FPN, the reward network, and/or other systems related to demand avoidance. Parametric designs allow sampling data from intermediate values of an independent variable. These intermediate values correct the shape of the underlying neural function. Factorial designs, on the other hand, sample from variable extremes and run the risk of undermining the functional relationship between neural activity and

cognitive process. This problem can exaggerate the influence of other factors such as uncontrolled variables in the design. For example, in a design where there are only two effort levels, the change in neural activity across these levels might be interpreted as the effect of the added effort level. However, performing effortful tasks might be less aversive than the boredom experienced during the performance of an underwhelming task. In that case, the low effort task might be accompanied by neural activity associated with boredom. In case the relationship between boredom and task levels is not monotonic, the change in neural activity across task levels would be a mixture of two different processes and interacting neural functions. Such a problem could be remedied by a parametric design which permits the estimation of separate non-monotonic functions. Consequently, prior to model fits of the neural data, the function of interest can be isolated by removing out the effects of other functions.

Further, we addressed how these within-subjects effects were modulated by individual differences in demand avoidance versus demand seeking behavior. Using this approach, we observed only qualified support for the cost for control hypothesis, in that though FPN activity was related to effort avoidance, this relationship was not clearly mediated by the cognitive control manipulation, as opposed to other factors. However, we discovered that effort-related modulation of regions corresponding to the “default mode network” were associated with effort-related changes in effort avoidance.

Methods

Participants

Twenty-eight right-handed adults (aged 18-35; 14 female) with normal or corrected to-normal vision were recruited for the behavioral demand avoidance experiment (Experiment 1) using Brown research subject pool. Fifty-six right-handed adults (aged 18-35; 26 female) with

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




Figure 1.1. Schematic of the parametric demand selection task. (a) Example trials of the categorization task with a digit printed on a colored circle. The color cued which task to perform. The response mappings shown below the digits are for the reader's reference and did not actually appear on the screen. (b) An example of six deck cues with their associated effort level pairing,

based on the probability of task switching in 13 trials from that deck. (c) Schematic of block events during the Learning Phase. First, a symbol icon is presented for 1 sec, after which the participant completes 13 categorization trials with the associated cue tiled on the background. (d) Schematic of block events during the Test Phase. In the “Selection” epoch, the participant chooses between two symbol icons that are associated with two different effort levels, with a deadline of 3 sec. In the “Execution” epoch, the participant executes the effort level associated with the selected option, while the selected cue is tiled on the background.

Behavioral Task

In both the Experiment 1 and 2, participants performed a parametric variant of the demand selection task (DST; Fig 1.1.). In an initial Learning phase, the participant associated virtual card “decks”, denoted by particular symbols, with a certain effort level through experience with trials drawn from that deck. Then, in a second Test phase, participants chose which deck they wished to perform (Selection epoch) and then performed the trials drawn from that deck (Execution epoch). A distinguishing characteristic of this version of the DST relative to prior versions is that we varied the effort level parametrically over six levels based on the frequency of task switching required by a deck. This allowed us to characterize the functional form of changes in behavior or fMRI signal due to our cognitive control manipulation. We now provide detail on each of these phases of the behavioral task.

The Task “Decks”

Throughout all phases of the experiment, participants performed trials drawn from virtual decks (Fig 1.1.a). Blocks consisted of sequences of 13 consecutive trials drawn from a given deck. On each trial, the participant categorized a presented digit as either odd/even (parity judgment) or greater/less than 5 (magnitude judgment). The color of a circle (green or blue) surrounding the digit cued whether to perform the parity or magnitude judgment on each trial based on a pre-instructed mapping. The participant indicated their categorization by pressing the left or the right arrow key on the keyboard. The response deadline was 1.5 sec. Trials were

separated by .2 sec. Response mappings and color-to-task mappings were counterbalanced across participants. Digits 1-4 and 6-9 were used with equal frequency across both tasks and were randomized for order of presentation.

In order to manipulate cognitive effort across decks, we varied the frequency of task switching required by a particular deck (Fig 1.1.b). The probability of switching tasks from trial to trial within a sequence drawn from each deck increased across 6 effort levels: 8%, 23%, 38%, 54%, 69%, 85%. Note that the lowest effort level still included one switch. In this way, all effort levels included the same number of parity and magnitude trials. Thus, any differences in effort preference among these could not be attributed to differences in difficulty or effort between the tasks themselves. Likewise, we did not include “pure blocks” of only parity judgments or only magnitude judgments, as doing so would have made these decks qualitatively different from all other decks, in terms of the amount of each task being performed. Having the lowest effort level include one switch ensures that the only parameter that changes between effort levels is the frequency of task-switching. A higher frequency of task switching is associated with a higher experience of cognitive effort (Monsell, 2003), and has been shown to be aversive (Arrington & Logan, 2004). At the beginning of each block, a shape was presented for 1 sec to indicate which deck was being performed, and this shape was also tiled in the background throughout performance of the block. Participants were told that this shape represented the symbol on the back of the virtual card deck from which trials for that sequence were drawn. Each effort level was associated with a particular deck symbol. Participants were not told about this relationship, but could learn through experience that certain decks were more effortful than others to perform. The mappings between deck symbols and effort levels were randomized across participants.

Practice Phase

During an initial “Practice Phase” participants gained familiarity with the trial structure, the categorization decisions, and the color and response mappings. After being given instructions regarding the categorization and response rules (Fig 1.1.a), they practiced two 13-trial blocks. Each task sequence included the presentation of a deck symbol, and the subsequent performance of 13 consecutive task-switching trials. In the first block of the Practice Phase, feedback was presented after the button press as either 'Correct' in pink, 'Incorrect' in yellow, or 'Please respond faster' if the participant failed to answer in 1.5s. In the second block of the Practice phase, feedback was omitted as would be the case in the actual experiment. The deck symbols that participants encountered in the Practice Phase were for practice only and were not presented during the learning or test phases to avoid any association of error feedback with a particular deck symbol.

Learning Phase

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

Test Phase

In the Test Phase (Fig 1.1.d), two decks were presented and the participant chose which they would like to perform (Selection epoch). The participants were told to freely choose between decks prior to a 3 sec deadline. We note that in contrast to other DST variants (Gold et al., 2015), participants in this task were not told about the difficulty manipulation to avoid biasing choices based on participants’ explicit beliefs about effort. Once the participant made their selection, the selected deck turned blue and both options remained on the screen until the end of the 3 sec deadline. In the event of a non-response, the same choice pair was re-presented

at the end of the entire experiment until they made a decision. Each pair was selected from the set of fifteen unique (un-ordered) pair combinations of all six decks, excluding self-pairing (e.g., deck #1 paired with deck #1). Each deck was presented either on the left or the right side of the screen, counterbalanced for location across trials. The Selection epoch was followed by the execution of the selected effort level task sequence (Execution epoch). The sequence of events in this epoch was the same as during the Learning phase. We note that based on the self-selected execution of effort tasks in this phase, there were differences in the number of trials each individual performed at each effort level. Although this variability naturally alters our detection power for corresponding effort trials across individuals in the Test Phase, Learning phase ensured that people established stable behavior and effort associations with each level based on the same number of trials with each effort level prior to the Test phase

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

Post-experimental Debriefing Inventory

Upon the completion of the experiment, participants were asked to respond to 6 questions regarding their understanding of the task. The first 5 questions required text entry. The final question required participants to enter a number. The questions were as follows: 1) What do you think the deck symbols stood for?, 2) Was there any reason for why you selected certain decks in the last phase?, 3) Did you have a preference for any deck?, 4) What do you think the experimental manipulation was?, 5) Imagine somebody else is about to perform the same experiment. What advice would you give them? And, did you follow this advice?, 6) Please rate each deck symbol in terms of its difficulty (out of 6).

Behavioral Data Analysis

Trials with response times below 200 ms were excluded from further analysis. Execution trials on which participants missed the response deadline were also excluded from further analysis (approximately %1 of execution trials in both phases). Response times were calculated using only correct trials.

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

In order to appropriately test for changes in behavioral selection rates across effort levels while also defining group membership (i.e., Demand Avoiders or Demand Seekers) based on the same choice behavior, we conducted a permutation procedure. In order to appropriately test for changes in behavioral selection rates across effort levels while also defining group membership (i.e., Demand Avoiders or Demand Seekers) based on the same choice behavior, we conducted a permutation procedure. First, we estimated likelihood distributions of the slope change in selection rates across effort levels in two different choice-defined groups. Specifically, for each of 10,000 iterations, we conducted the following procedure. (1) We selected a random half of each participant's data and calculated an overall rate of choosing the harder task from those data. (2) We assigned each participant to either the Demand Seeker or Demand Avoider group depending on whether their selection rate was greater or less than 50% respectively. (3) Using the other half of each participant's data, we fit a linear function to selection rates across all effort levels, and extracted the slope (β). (4) We averaged the β values for each group separately. Over 10,000 iterations for each group, this procedure yielded likelihood distributions of selection rate slopes across effort levels for Demand Avoiders and Demand Seekers.

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

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

MRI procedure

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

fMRI Analysis

Functional images were preprocessed in SPM12 (<http://www.fil.ion.ucl.ac.uk/spm>). Before preprocessing, data were inspected for artifacts and variance in global signal (tsdiffana, art_global, art_movie). Functional data were corrected for differences in slice acquisition timing by resampling slices to match the first slice. Next, functional data were realigned (corrected for motion) using B-spline interpolation and referenced to the mean functional image. 1-2 sessions were excluded from 3 other participants prior to behavioral analysis due to movement during data collection in the scanner. Functional and structural images were normalized to Montreal

Neurological Institute (MNI) stereotaxic space using affine regularization followed by a nonlinear transformation based on a cosine basis set, and then resampled into 2x2x2 mm voxels using trilinear interpolation. Lastly, images were spatially smoothed with an 8 mm full-width at half-maximum isotropic Gaussian kernel.

A temporal high-pass filter of 128 (.0078 Hz) was applied to our functional data in order to remove noise. Changes in MR signal were modeled under assumptions of the general linear model (GLM). Two GLMs were devised: a linear effort-level GLM and an independent effort-level GLM. Both GLMs included nuisance regressors for the six motion parameters (x,y,z,pitch,roll,yaw) and four run regressors for the ‘Linear Effort-Level GLM’ and one run regressor for the ‘Independent Effort Level GLM’. The number of run regressors was different across GLMs because ‘Independent Effort GLM’ included the regressor for each effort level separately. In this GLM, it was not possible to have 4 runs for each effort level since some participants did not choose to select from an effort level an entire run. Thus, in order to retain the entirety of the dataset, we collapsed all trials across runs for each effort level.

Linear Effort-Level GLM

The linear effort-level GLM tested which voxels in the brain parametrically increased or decreased linearly with effort level. Two event regressors were used. Execution events were modeled as a boxcar that onset with the presentation of the first trial stimulus of the sequence and ended with the participant’s response to the final item. Thus, the duration of this boxcar varied as a function of response time. Second, the Selection event regressor modeled each selection choice event with a fixed boxcar of three secs. We used parametric modulators on these event regressors to test the linear effect of effort level. The Execution event regressor was modulated by an Effort Level parametric regressor corresponding to the effort level of that task

sequence (1 through 6). The Selection event regressor was modulated by (a) an Effort Level parametric regressor based on the chosen effort level (1 through 6), and (b) a Difference regressor that scaled with the difference between the chosen and the unchosen effort option on that selection trial (1 through 5). Note that as implemented by SPM, each parametric regressor includes only the variance unique to it and shared with those ordered after it. Thus, for example, the Effort Level regressor includes variance explained over and above that shared with the Execution event regressor. The Difference regressor did not yield statistically reliable results and is not discussed further. The Execution and Selection event regressors, along with their parametric modulators, were modeled separately for each scanning run within the GLM.

Independent Effort Level GLM

The independent effort level GLM sought to characterize the signal change related to each effort level independently of each other or of any particular function (e.g., linear). This GLM included twelve event regressors, one for each effort level (1 through 6) by epoch (Execution and Selection). Events in the Execution regressors were modeled as boxcars that onset with the presentation of the first trial stimulus of the sequence and ended with the participant's response to the final item. Events in the Selection regressors were modeled with an 3 sec boxcar at the onset of each choice pair. The Selection onset regressor was further modulated by a parametric Difference regressor that scaled with the difference between the chosen and the unchosen effort option on that selection trial. Boxcars and durations were the same as in the linear effort level model. In this GLM, four run epochs and a linear drift over the whole experiment were included as nuisance regressors.

For both GLMs, SPM-generated regressors were created by convolving onset boxcars and parametric functions with the canonical hemodynamic response (HRF) function and the temporal derivative of the HRF. Beta weights for each regressor were estimated in a first-level, subject-specific fixed-effects model. For group analysis, the subject-specific beta estimates were analyzed with subject treated as a random effect. At each voxel, a one-sample t-test against a contrast value of zero gave us our estimate of statistical reliability. For whole brain analysis, we corrected for multiple comparison using cluster correction, with a cluster forming threshold of $\alpha = .001$ and an extent threshold, k , calculated with SPM to set a family-wise error cluster level corrected threshold of $p < .05$ for each contrast and group (See Table 1.1.). Note that the higher cluster forming threshold helps avoid violation of parametric assumptions such that the rate of false positive is appropriate (Eklund et al., 2016; Flandin & Friston, 2016).

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

Table 1.1. Cluster extent (k) thresholds as computed by SPM for each contrast and demand groups using a cluster forming threshold, $\alpha = .001$).

ROI analysis

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

Independent ROIs

In order to test the cost of control hypothesis, we defined a fronto-parietal control network ROI defined from a functionally neutral group ([Network 12] Yeo et al., 2011) along with a VS ROI (Badre et al., 2014), as VS activity has been consistently observed in cognitive effort literature (Botvinick et al., 2009; Schouppe et al., 2014a). We also included a Default Mode Network ROI from a functionally neutral group ([Network 16] Yeo et al., 2011). A priori DMN regions included ventromedial prefrontal cortex (vmPFC), orbitofrontal cortex, posterior cingulate cortex and parts of precuneus. A priori FPN regions included bilateral lateral prefrontal cortex, bilateral parietal cortex and SMA.

PCA ROIs

Our study is the first to parametrically manipulate implicitly learned effort cost/values in DST in fMRI. Though we *a priori* conceived of these effort levels as being linear, and so showing a linear brain response, we acknowledge that this need not be the case. As such, we wanted to be sure that our *a priori* expectation of a linear response function with increasing effort, as tested in the parametric GLM, did not cause us to systematically miss networks of the brain that code effort according to a different functional form, but that are relevant signals for the behavioral decision. In order to have an unbiased analysis of the neural response function across effort levels, we adopted a data-driven Principal Component Analysis (PCA) approach on the whole brain (126,866 voxels) to explore the shape of the neural functions of different brain areas that respond to effort execution, and their relation to demand avoidance behavior. We ran the PCA over averaged β estimates of each voxel in the whole brain as they were derived from the Independent Effort Level GLM task Execution Level 1-6 onset regressors across all participants. Therefore, PCA was conducted on a total of 126,866 observations and 6 variables.

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

Functional connectivity analysis

For each participant, functional connectivity analysis was implemented in MATLAB using the CONN toolbox (<http://www.nitrc.org/projects/conn>; [Whitfield-Gabrieli & Nieto-Castanon, 2012](#)). CONN uses the CompCor method ([Behzadi et al., 2007](#)) which segments white matter (WM) and cerebrospinal fluid (CSF), as well as the realignment parameters entered as confounds in the first-level GLM analysis ([Behzadi et al., 2007](#)), and the data were band-pass filtered to 0.01 Hz to 0.1 Hz.

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

Next, we focused on two connectivity measures. First, we calculated the average connectivity Between Positive-Negative PC1, within Positive PC1 and within Negative PC1 clusters at each effort execution level, in order to observe the connectivity change in these

measures with increasing effort execution. Next, we collapsed connectivity coefficients for Between Positive-Negative PC1, within Positive PC1 and within Negative PC1 across effort levels and compared average connectivity rates between demand groups. The results of the connectivity analysis did not yield statistically reliable results and they are reported in Supplementary Results 6.

Brain – behavior analysis

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

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

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

In order to test the effects of brain and performance on selection behavior, while accounting for individual variability within demand groups, we adopted the following mixed effects hierarchical regression model:

$$y_i = X_i(\beta + s_{j[i]} + m_{[i]}) + \epsilon$$

$$s_j \sim N(0, \Sigma_S), \quad \epsilon \sim N(0, \sigma)$$

Here, y_i is the i th observed selection rate, X_i is a matrix of predictors (brain activity) and covariates (performance), β_i is a vector of coefficients (as in conventional linear regression), $j[i]$ is the subject of the i th observation, so that $s_{j[i]}$ is a subject-specific perturbation of all of the coefficients.

This fitting procedure was implemented using the R package `blme`, which includes the `lme4` package (Bates et al., 2014) and performs maximum-a-posteriori estimation of linear mixed-effects models.

Results

Demand avoidance behavior

Experiment 1 established the basic behavioral effects of our parametric effort manipulation on performance. Overall, participants were highly accurate on the categorization task across both phases (mean error: 12% in the Learning Phase, 15% in the Test Phase), and

participants missed the deadline on few trials (2.3% of Learning phase trials, $SE=0.003$, 1.4% of Test phase trials, $SE=0.01$).

Both RT and error rates showed a linear effect of effort condition, increasing with a higher proportion of task switches across both switch and repeat trials. There was a significant effect of Effort on error rates, ($F(2.76,74.53)=29.88$, $p < .001$, $\eta_p^2 = .53$) and RT ($F(2.69,72.71)=131.17$, $p < .001$, $\eta_p^2 = .83$). And, the increase in both was linear over effort levels (errors: $F(1,27)=61.17$, $p < .001$, $\eta_p^2 = .69$; RT: $F(1,27)=248.19$, $p < .001$, $\eta_p^2 = .90$). These increases were also significant for both switch and repeat trials tested separately (switch errors: $F(1,27)=29.73$, $p < .001$, $\eta_p^2 = .52$; repeat errors: $F(1,27)=42.59$, $p < .001$, $\eta_p^2 = .61$; switch RT: $F(1,27)=0.37$, $p < .001$, $\eta_p^2 = .61$; repeat RT: $F(1,27)=42.89$, $p < .001$, $\eta_p^2 = .61$).




Figure 1.2. Distribution of participants for the overall probability of selecting the easier task across all decision trials in A) Experiment 1 and B) Experiment 2. In Experiment 1, 25 of the 28 participants (89%) selected the easier task more than 50% of the time overall. In Experiment 2, nearly half of participants (N=24) selected the easier task more than 50% of the time. The mean overall probability of selecting the easier task for each experiment is indicated with a vertical dashed line.




Figure 1.3. Choice behavior for the entire Experiment 1 sample (A), Demand Avoiders in Experiment 2 (B), and Demand Seekers in Experiment 2 (C). The top panels plot the selection rates across effort levels in terms of the probability of choosing that task when it was given as an option. In Experiment 1, the entire sample showed a decreasing tendency to choose higher effort levels. In Experiment 2, Demand Avoiders showed a decreasing tendency to choose a task as more cognitive effort is required, and Demand Seekers exhibited the opposite pattern. The bottom panels plot the decision times across effort levels. In Experiment 1, decision time was mostly unaffected by effort level with only a marginal omnibus effect of effort level on decision time. In Experiment 2, the decision time to select an effort level linearly increased with increasing effort levels for Demand Avoiders and the decision time to choose an effort level was similar across effort levels for Demand Seekers. All positive and negative trends were significant linear effects at $p < .05$. The shaded area plots standard error of the mean.

During the Selection phase, 25 of the 28 participants (89%) selected the easier task more than 50% of the time overall (Fig 1.2.A). Further, there was a significant effect of effort level on choice behavior, ($F(5,135)=7.23, p<.001, \eta_p^2=.21$), such that higher effort levels were associated with higher avoidance rates than lower effort levels (Fig 3A). This pattern across effort levels was fit by a linear effect ($F(1,27)=24.54, p<.001, \eta_p^2=.48$). The probability of selecting the

easier task also significantly changed depending on the difference between effort levels given as options ($F(3,81.09)=6.07$, $p=.001$, $\eta_p^2=.18$), such that larger differences increased the probability of choosing the easier task. This effect was linear across effort levels ($F(1,27)=16.83$, $p<.001$, $\eta_p^2=.38$). Decision time (Fig 1.3.A) was mostly unaffected by effort level with only a marginal omnibus effect of effort level on decision time ($F(2.98,80.56)=2.62$, $p=.06$, $\eta_p^2=.09$) and no effect of difficulty difference.

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

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

The diagnostic feature of demand avoidance in our paradigm is the linear change in selection rates across effort levels, as observed in Experiment 1. Therefore, if the subgroup with overall selection rates above 50% (Demand Avoiders) are engaged in effort avoidance, we should also see a linear change in their selection rates across effort levels as was observed in

Experiment 1. Likewise, those with selection rates below 50% (Demand Seekers) might show the opposite pattern, with a linearly increasing tendency to choose the harder task. However, simply testing the slope across effort levels within each group might be biased by the way the groups are defined, as the same choice data would be used for both computations. Thus, to test this prediction, we performed a permutation procedure that allowed us to generate likelihood distributions for the slope of change in selection rates given an overall tendency (>50%) to avoid versus seek the harder task, as well as a null distribution. This procedure computed the selection rate slope and group membership on independent halves of the data over multiple (10,000) random draws from the data set.




Figure 1.4. The overlap of the likelihood distribution for the Demand Avoider (blue) and Demand Seeker (yellow) groups with the permuted null distribution (purple). The mean demand avoidance slope for Demand Avoiders is negative, and the demand avoidance slope for Demand Seekers is positive, consistent with a respective tendency to avoid versus select a task as function of its effort. These distributions overlapped minimally with the permuted null distribution (Demand avoider: $p = .008$, Demand seeker: $p = .003$).

Fig 1.4. plots the overlap of the likelihood distributions for the Demand Avoider and Demand Seeker groups with the permuted null distribution. First, the mean slope was negative for Demand Avoiders and positive for the Demand Seekers, consistent with a respective tendency to avoid versus select a task as function of its effort. Second, these distributions overlapped minimally with the permuted null distribution (Demand Avoider: $p < .01$, Demand Seeker: $p < .01$). Thus, the likelihood of randomly assigning individuals to two groups and obtaining the demand avoidance slopes we observe using a random half of the data is less than 1%.

Our analysis showed that overall demographics between demand groups was similar. There were no age (Demand Avoiders: $M= 22.12$, $SD= 3.20$; Demand Seekers: $M= 21.45$, $SD= 2.99$) or gender differences (Demand Avoiders: $N_{\text{female}}=11$, $N_{\text{male}}=15$; Demand Seekers: $N_{\text{female}}=14$, $N_{\text{male}}=10$) between demand groups (age: $F(1,49)=0.56$, $p= .46$; gender: $X^2(1, N = 50) = 1.28$, $p=.40$).

Post-experimental debriefing inventory results showed that participants were mostly unaware of the difficulty manipulation across effort tasks. Only 11 out of 50 participants reported that they noticed a difficulty difference between decks. Further, free text responses showed that even the “aware” participants were not completely sure about the exact difficulty ranking across decks. The final inventory asked participants to rate the difficulty of each deck out of 6. Across participants and decks, average rating accuracy was 15%, consistent with chance level or 1 out of 6 rated correctly.

In order to test if overall difficulty ratings predicted effort selections, we ran a hierarchical regression with ratings as the dependent variable and the effort selections as the independent variable. The results showed that post-experimental difficulty ratings significantly

predicted ($M=-0.01$, $SD=0.01$, $t(300)=0.03$) effort selections across all participants. In other words, participants selected those tasks that they rated to be more effortful less often. This was the case even if the difficulty ratings did not match the task-switching probability of the task or individuals' performance decrements during task-execution, both of which are argued to underlie the phenomenology of effort. Therefore, it is possible that the difficulty ratings could have been formed post-hoc in relation to task selections. A follow-up test averaged selection rates based on difficulty ratings across demand avoiding and demand seeking participants separately, in order to observe the functional form of selection rates as ordered by subjective ratings. Selection rates as ordered by subjective difficulty ratings across both groups were at chance level for all difficulty levels, indicating that participants' post-hoc difficulty ratings did not differentiate selection rates at the group level. Future studies should get online difficulty ratings during task execution and directly test how effort selections are related to prior reports of the subjective experience of effort.

RTs but not error rates (Table 1.2.) across effort levels and experimental phases differed across demand groups (errors: $F(1,48)=1.81$, $p=.19$, $\eta_p^2=.04$; RT: $F(1,48)=6.21$, $p=.02$, $\eta_p^2=.12$). Demand Avoiders had faster RTs compared to Demand Seekers in both phases. Both RT and error rates showed a linear effect of effort, indicating that our effort manipulation was successful in increasing task difficulty (errors: $F(1,48)=72.43$, $p<.001$, $\eta_p^2=.60$; RT: $F(1,48)=363.90$, $p<.001$, $\eta_p^2=.88$). RTs but not error rates showed improvements across experimental phases (errors: $F(1,48)=1.01$, $p=.32$, $\eta_p^2=.02$; RT: $F(1,48)=113.49$, $p<.001$, $\eta_p^2=.70$).

		RT				ER			
	Difficulty	Demand Seekers		Demand Avoiders		Demand Seekers		Demand Avoiders	
		Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.
Learning Phase	1	0.54	0.05	0.50	0.09	0.11	0.13	0.07	0.10
	2	0.58	0.06	0.53	0.10	0.13	0.15	0.09	0.10
	3	0.63	0.07	0.57	0.09	0.14	0.15	0.11	0.11
	4	0.65	0.07	0.60	0.10	0.14	0.14	0.12	0.10
	5	0.69	0.08	0.64	0.12	0.17	0.14	0.13	0.11
	6	0.73	0.08	0.65	0.13	0.16	0.13	0.14	0.12
Test Phase	1	0.49	0.06	0.45	0.07	0.13	0.16	0.07	0.08
	2	0.53	0.06	0.47	0.07	0.13	0.14	0.07	0.09
	3	0.55	0.07	0.49	0.08	0.14	0.14	0.09	0.10
	4	0.57	0.07	0.52	0.09	0.15	0.14	0.10	0.10
	5	0.59	0.08	0.53	0.09	0.16	0.14	0.11	0.09
	6	0.60	0.09	0.55	0.11	0.17	0.15	0.10	0.08

Table 1.2. reports descriptive statistics for Error rates (ER) and Response Times (RT(s)) for each demand group, difficulty level and experimental phase.

For Demand Avoiders, there was a significant effect of effort level on choice behavior, $F(5,125)=8.24, p < .001$. Similar to the pilot behavioral study, they were significantly less likely to choose an effort level with increasing effort requirement; there was a significant linear effect, $F(1,25)=31.46, p < .001$. There was a significant effect of effort level on decision time $F(5,125)=3.05, p = .01$. The decision time to select an effort level linearly increased with increasing effort levels, $F(1,25)=6.34, p = .02$. For Demand Seekers, there was a significant effect of effort level on choice behavior, $F(3.22,73.98)=4.67, p < .01$. They were significantly more likely to choose an effort level with increasing effort requirement, there was a significant linear effect, $F(1,23)=34.20, p < .001$. The decision time to choose an effort level was similar ($M= 1.00$ s., $SD= 0.25$) across all effort levels, $F(2.88,66.25)=2.55, p = .07$.

In addition to calculating the probability of selecting an effort level, we also analyzed the probability of choosing the easier task by collapsing all trials across the difficulty difference they yield between options in order to estimate the effort distance effects mostly independent of a particular effort level. The probability of selecting the easier task significantly differed across difficulty difference between effort levels $F(2.31,57.99)=6.94, p < .001$. The probability of selecting the easier linearly increased with increasing difference between effort pairs,

$F(1,25)=15.70$, $p< .001$. The decision time to choose the easier task was similar ($M= 0.89$ s., $SD= 0.20$) across all difference levels $F(2.50,62.49)=1.10$, $p= .35$. The probability of selecting the easier task changed across difficulty difference between effort levels $F(2.88,66.27)=3.58$, $p=.02$). The probability of selecting the easier linearly decreased with increasing difference between effort pairs, $F(1,23)=9.03$, $p< .01$. The decision time to choose the easier task was similar ($M= 0.98$ s., $SD= 0.24$) across all difference levels $F(2.16,49.81)=0.72$, $p= .58$.

In order to test selection rates relative to a common baseline, we recalculated overall probability of selecting the easier task separately for only choices when a higher effort level (effort level 2,3,4,5,6) was paired with the lowest effort level (effort level 1). Consistent with our original grouping procedure, the probability of selecting the easier task across these pairs significantly differed between demand groups, with Demand Avoiders selecting the easiest effort level more than Demand Seekers each time the easiest task was paired with a higher effort level (Demand Avoiders: $M= .72$, $SD= 0.04$; Demand Seekers: $M= .39$, $SD= 0.04$; $F(1,48)=36.76$, $p<.001$, $\eta_p^2 =.43$). Selection rates across demand groups were similar between pairs ($F(3.42,164.36)=0.36$, $p=.81$, $\eta_p^2 =.01$), indicating that participants' preferences for the easiest task were stable regardless of the effort level it was paired with.

Run	Demand Seekers		Demand Avoiders	
	Mean	Std. dev.	Mean	Std. dev.
1	0.44	0.12	0.59	0.11
2	0.42	0.10	0.60	0.12
3	0.41	0.09	0.63	0.13
4	0.44	0.07	0.59	0.17

Table 1.3. reports demand avoidance rates for each demand group across runs during the Test Phase.

In order to examine the effects of time on effort selection behavior, overall probability of selecting the easier task was calculated separately for each of the four blocks of the Test Phase. Probability of selecting the easier task did not change across Test blocks ($F(2.59,124.22)=0.14$, $p=.92$, $\eta_p^2=.003$), indicating that participants' effort selection were consistent across time (Table 1.3.). Furthermore, the consistency in selection behavior did not change between demand groups ($F(2.59,124.22)=1.60$, $p=.20$, $\eta_p^2=.03$).

Additionally, we tested the effects of performance and task-switching probability on effort selection behavior using hierarchical regression analysis. First, Error Rates during task-execution did not predict deck selection for either Group (Demand Avoiders: $M_\beta=-0.28.$, $SD_\beta=0.19$, $t(156) = -1.48$, $p = .14$; Demand Seekers: $M_\beta=0.08.$, $SD_\beta=0.11$, $t(144) = 0.67$, $p = .50$). RT during task execution predicted demand avoidance only in the Demand Avoider group (Demand Avoiders: $M_\beta=-0.44.$, $SD_\beta=0.09$, $t(156) = -2.34$, $p = .02$; Demand Seekers: $M_\beta=0.35.$, $SD_\beta=0.20$, $t(144) = 1.73$, $p = .09$), indicating that Demand Avoiders selected those tasks that yielded the least time-on-task.

Next, we tested the effects of task-switching probability on the selection rates in the presence of RT in order to see if our linear effort manipulation predicted demand avoidance over and beyond RT during task-execution. Demand Avoiders reliably selected those tasks that required the least task-switching, whereas Demand Seekers sought them (Demand Avoiders: $M_\beta=-0.34.$, $SD_\beta=0.03$, $t(156) = -5.82$, $p < .001$; Demand Seekers: $M_\beta=0.24.$, $SD_\beta=0.06$, $t(144) = 4.09$, $p < .001$). The effects of RT were no longer significant when task-switching probability was added to the model (Demand Avoiders: $M_\beta=-0.07.$, $SD_\beta=0.18$, $t(156) = -0.37$, $p = .71$; Demand Seekers: $M_\beta= 0$, $SD_\beta=1.05$, $t(144) = 0$, $p = .99$), indicating that selection behavior was explained better with our linear task manipulation than time-on-task during task execution.

In summary, the behavioral data provide evidence for two groups of participants that differed in their effort choices despite no demographic or performance differences between them except overall faster RTs for Demand Avoiders. Demand Avoiders (Fig 1.3.B) showed a decreasing tendency to choose a task as more cognitive effort is required. Demand Seekers (Fig 1.3.C) exhibited the opposite pattern. This pattern in effort selections was mostly captured by our linear effort manipulation on task-switching probability for both demand groups. We note that based on a behavioral experiment reported in the Supplementary materials (see Supplementary Text), this higher tendency to seek demand in the fMRI group versus the behavioral group may reflect a self-selection bias. Specifically, those who volunteer for fMRI tended to score higher on the Need for Cognition scale, and so might be more likely to be demand seeking. These behaviorally defined groupings were maintained throughout the subsequent fMRI analyses in order to address individual differences.

The functional form of univariate brain activity over effort levels

We investigated the pattern of activation in independent ROIs based on our *a priori* hypotheses, in order to examine their underlying neural function across the execution of increasing effort levels. We defined ROIs in the FPN (Network12; Yeo et al, 2011), DMN (Network16; Yeo et al., 2011), and/or ventral striatum (VS; Badre et al., 2014). Then, using the “Independent Effort Level GLM”, we estimated the activation in the network for each effort level separately. These estimates are plotted in Fig 1.4.

As expected, the FPN ROI showed a linear trend across effort levels (Fig 1.5.B). There was a significant effect of effort on FPN β estimates, $F(5,240)=5.19$, $p<.001$, $\eta_p^2=.10$. The β estimates linearly increased with increasing effort requirements of the task, $F(1,48)=20.85$, $p<$

.001, $\eta_p^2 = .30$. There was no effect of avoidance group on β estimates, $F(1,48)=0.57, p= .45, \eta_p^2 = .01$.

The DMN showed a negative linear trend in activation across effort levels (Fig 1.5.A). There was a significant effect of effort on DMN β estimates, $F(5,240)=4.18, p= .001, \eta_p^2 = .08$. The β estimates linearly decreased with increasing effort requirements of the task, $F(1,48)=9.77, p < .01, \eta_p^2 = .17$. There was no effect of avoidance group on β estimates, $F(1,48)=1.06, p= .31, \eta_p^2 = .02$.

An ROI in VS also showed a negative linear trend (Fig 1.5.C). There was a significant effect of effort on VS β estimates, $F(3.26,156.52)=5.94, p = .001, \eta_p^2 = .11$. The β estimates linearly decreased with increasing effort requirements of the task, $F(1,48)=9.40, p < .01, \eta_p^2 = .16$. There was no effect of avoidance group on β estimates, $F(1,48)=0.32, p= .58, \eta_p^2 = .01$. We note that in all cases a linear function was the curve of best fit relative to alternatives (summarized in Supplementary Table 1.4.), with the exception of VS that, though showing a reliable linear function, was better fit by a quadratic due to an apparent asymptote at level 4.

	Linear	Quadratic	Cubic	Order4	Order5
FPC	$F(1,48)=20.85, p < .001^*$	$F(1,48)=3.51, p= .07$	$F(1,48)=0.09, p= .77$	$F(1,48)=0.28, p= .60$	$F(1,48)=0.09, p= .77$
DMN	$F(1,48)=9.77, p= .003^*$	$F(1,48)=7.26, p= .01^*$	$F(1,48)=0.21, p= .65$	$F(1,48)=0.66, p= .42$	$F(1,48)=1.56, p= .22$
VS	$F(1,48)=9.40, p= .004^*$	$F(1,48)=12.70, p= .001^*$	$F(1,48)=2.02, p= .16$	$F(1,48)=7.06, p= .01^*$	$F(1,48)=1.06, p= .31$

Table 1.4. Curve fits for Independent ROIs: FPC, DMN and VS. Significant fits are presented with an asterisk (*).




Figure 1.5. Estimated beta weights at each effort level for both demand groups from independent ROIs defined in (A) the default mode network (DMN), (B) the frontoparietal network (FPN), and (C) the ventral striatum (VS). Demand Avoiders are indicated in solid lines. Demand Seekers are indicated in dashed lines. There was no effect of avoidance group on β estimates for any of the ROIs. In all cases a linear function was the curve of best fit relative to alternatives (summarized in Supplementary Table 1.4.), with the exception of VS. All linear trends are significant at $p < .05$ for both groups.

In order to ensure that our ROI approach did not exclude regions or networks outside of our *a priori* defined ROIs, we conducted whole brain analyses based on parametric contrasts across effort levels. Figure 1.6. plots corrected whole brain activation from the main univariate contrasts. During task performance (Execution > Baseline), both groups activated a fronto-parietal network commonly seen during cognitive control tasks and showed deactivated (Baseline > Execution) regions of inferior parietal lobule and precuneus that overlap the “default-mode network” (DMN) as well as regions of posterior cingulate cortex and medial prefrontal cortex commonly associated with reward. There was no significant difference between Demand Avoiders and Demand Seekers in these basic task-positive and task-negative activation patterns.

The frontoparietal control network showed a positive correlation with effort level (Parametric+) and conversely DMN and VS negatively correlated with parametric increases (Parametric-) in effortful task execution. There were also no group differences between for the Parametric+ or the Parametric- contrasts (Fig 1.6.C-D; Table 1.5.), indicating that the results of these whole-brain analysis results were consistent with the ROI analyses.

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

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

Table 1.5. Whole brain analysis for both demand groups. (* indicates that higher clustering threshold was used in order to break down big clusters for reporting purposes)




Figure 1.6. Four whole brain contrasts are plotted on a canonical surface with a coronal slice (Y=8) on the far right to show activation in the basal ganglia. Yellow: Demand Avoiders; purple: Demand Seekers; Red: overlap between Demand Avoiders and Demand Seekers. (A) The contrast of all blocks (across effort levels) from the execution phase versus baseline primarily shows the FPN and other typical task active areas. (B) The reverse contrast of baseline > execution phase blocks shows typical task negative activations including the DMN. (C) The positive parametric effort level effect shows regions that demonstrated a linear increase with effort level, including regions of the FPN. (D) The negative parametric effort level effect shows regions that were negatively correlated with effort level including regions of the DMN and ventral striatum. Images are thresholded at $p < .05$ FWE cluster corrected.




Figure 1.7. Summary of the top components identified by the PCA of the activation pattern over effort levels across all voxels in the brain. Left panel: Eigenvalue by each PC component. Right Panel: Variance explained by each PC component. Further analysis focused on the top 3 components as these explained 90% of the variance.

From 126,866 voxels included in the analysis, the PCA located 3 PCs that explained 90% of the variance of the data. The first PC explained 57% of the variance. The percent variance of the data explained by each subsequent PCs reached asymptote by 4-5 PC, and 5 PCs explained 100% of the total variance (Fig 1.7.). We extracted 3000 voxels that loaded the most positively (Positive PCs) and most negatively (Negative PCs) on each component and extracted their β s in order to observe the functional form of each PC.




Figure 1.8. Visualization of PCA components and their neural functions across effort levels. A) Results from an analysis of the first 3000 voxels that load positively on components 1-3. Component1 showed a significant linear trend; Component2 showed a significant quadratic trend; Component3 beta estimates were flat across effort levels. B) Results from an analysis of the first 3000 voxels that load negatively on components 1-3. Component1 showed a significant negative linear trend; Component2 was not fit by any function. Component3 beta estimates showed a quadratic trend across effort levels. On the brain: Red: Component1; Green: Component2; Blue: Component3.




Figure 1.9. overlays the DMN (top) and FPN (bottom) ROIs on canonical brain as they are separately defined a priori and by PCA.

Positive PC1 (Fig 1.8., Table 1.6.) overlapped with the frontoparietal network, and included the left SMA, left superior parietal lobule, left middle frontal gyrus, left inferior frontal gyrus. As we observed in Parametric+ whole-brain analysis, there was a significant effect of effort on FPN cluster β estimates as defined by Positive PC1, $F(3.84, 184.18) = 19.08$, $p < .001$, $\eta_p^2 = .28$. The β estimates from these voxels linearly increased with increasing effort requirements of the task, $F(1, 48) = 69.24$, $p < .001$, $\eta_p^2 = .59$. This converges with the preceding *a priori* ROI analyses and indicates that the most voxel-wise variance in effort execution can be explained by a linear trend in the brain regions that contribute to controlled behavior.

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

Table 1.6. Cluster coordinates for the Positive PCs.

Negative PC1 (Fig 1.8., Table 1.7.) overlapped with the DMN, and included the medial prefrontal cortex, posterior cingulate cortex, left and right angular gyrus, left and right temporal gyrus, and left and right inferior frontal gyrus. The β estimates linearly decreased with increasing effort requirements of the task, $F(1,48)=27.45, p<.001, \eta_p^2=.36$.

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

Table 1.7. Cluster coordinates for the Negative PCs.

Positive PC2 (Fig 1.8., Table 1.6.) included anterior cingulate cortex, right orbitofrontal cortex, caudal segments of the right parietal cortex, left precuneus, and cuneus. There was a significant effect of effort on Positive PC2 β estimates, $F(3.51, 168.37) = 4.05$, $p < .01$, $\eta_p^2 = .08$. The β estimates had a quadratic trend across increasing effort requirements of the task, $F(1, 48) = 19.63$, $p < .001$, $\eta_p^2 = .29$. No cortical or subcortical regions loaded negatively on PC2.

Regions that loaded most positively on PC3 (Fig 1.8., Table 1.6.) included the dorsal portions of the occipital lobe. Consistent with its functional definition, there was not a significant effect of effort on Positive PC3 β estimates, $F(3.78, 181.22) = 2.34$, $p = .06$, $\eta_p^2 = .05$. Despite the visual nature of the task, this component overall showed a task-negative trend across effort levels.

Negative PC3 (Fig 1.8., Table 1.7.) included right inferior frontal gyrus, medial frontal cortex, left and right angular gyrus, brainstem and left insula, regions that are mainly adjacent to the regions observed in Positive PC2. There was a significant effect of effort on Negative PC3 β estimates, $F(3.71,178.25)=5.17$, $p=.001$, $\eta_p^2=.10$. Similar to that of Positive PC2, the β estimates had a quadratic trend across increasing effort requirements of the task, $F(1,48)=29.34$, $p<.001$, $\eta_p^2=.38$.

Overall, PCA results indicate that a linear function explained more than 50% of the voxel-wise variance in the shape of activation function across effort levels with Positive PC1 showing a positive trend and Negative PC1 showing a negative trend. We note that the Positive PC1 has a close correspondence to the *a priori* defined FPN, and the Negative PC1 has a similar relationship to the DMN (see Figure 1.9.). However, these correspondences are not perfect with substantial components of each of the *a priori* networks not showing clear positive or negative linear patterns. The PCA is an unbiased means of defining regions that show a linear pattern across effort levels. Thus, we use the PCA definition for subsequent analyses. However, in order to make its relationship to the *a priori* networks clearer, we refer to these ROI definitions from the PCA as the “Fronto-Parietal Component” and the “Default Mode Component”. We also performed all analyses in the *a priori* defined networks and report these analyses in Supplementary Results 8.

Selection Epoch

Independent ROI analyses were conducted for each effort level of the Selection Phase. The results show that there was no significant effect of effort, a significant effect of group, or an interaction in either of these analyses (Network12: no main effect of effort: $F(4.18,200.92)=0.47$, $p=.77$, $\eta_p^2=.01$, no main effect of group: $F(1,48)=0.17$, $p=.68$, $\eta_p^2=.003$, no effort*group

interaction: $F(4.18, 200.92) = 1.73$, $p = .14$, $\eta_p^2 = .04$; VS: no main effect of effort: $F(4.05, 194.31) = 0.42$, $p = .80$, $\eta_p^2 = .01$, no main effect of group: $F(1, 48) = 0.15$, $p = .70$, $\eta_p^2 = .003$, no effort*group interaction: $F(4.05, 194.31) = 0.65$, $p = .66$, $\eta_p^2 = .01$). Instead, whole-brain analysis showed that effort selections linearly scaled with motor cortex (MNI XYZ = -38, -28, 54; Cluster corrected $p < .05$).

Brain – Behavior Analysis

Having established the functional form of activity change across effort levels, we tested whether brain regions tracking effort execution predicted effort selections. Two correlation techniques were adopted to examine the type of brain-behavior relationship: 1) relationship between brain activity and effort selection rate at each effort level ('Individual task selection'), 2) relationship between change in brain activity and change in effort selection rates across effort levels ('Effort Level').

'Individual task selection' analysis showed that Positive PC1 ROI activity during task-execution was uncorrelated with deck selection in both groups (Demand Avoiders: $M_\beta = -0.6$., $SD_\beta = 0.05$, $t(156) = -1.26$, $p = .21$; Demand Seekers: $M_\beta = 0.05$., $SD_\beta = 0.05$, $t(144) = 0.94$, $p = .35$). These results were consistent between ROI definitions. Thus, we did not observe a reliable relationship between fronto-parietal network and avoidance behavior.

'Individual task selection' analysis showed that Negative PC1 β estimates did not individually predict task selection behavior independent of our effort manipulation in either group (Demand Avoiders: $M_\beta = 0.01$., $SD_\beta = 0.06$, $t(156) = 0.21$, $p = .33$; Demand Seekers: $M_\beta = -0.08$., $SD_\beta = 0.08$, $t(144) = -0.91$, $p = .36$).

'Effort Level' analysis showed that change in Fronto-Parietal Component during task execution across effort levels (i.e., due to the task switching manipulation) did not reliably

predict change in effort selection rates in demand avoiders ($r(24) = .17, p=.41$). For demand seekers, there was a marginal negative correlation between demand avoidance and Fronto-parietal Component change across effort levels ($r(22) = -.40, p=.052$). Thus, to summarize, we did not find evidence that Fronto-Parietal Component activation and task selection was related to the parametric effort manipulation in demand avoiders. By contrast, demand seeking behavior in demand seekers marginally increased with increasing Fronto-Parietal Component activity change (Fig 1.10.). We emphasize that the latter observation was not statistically significant, and so does not provide evidence of a relationship between activation in FPN and effort selection. Nevertheless, even if one were to accept the marginal effect, the non-significant trend is incongruent with the ‘cost of control’ hypothesis, such that greater FPN recruitment during harder tasks yielded greater effort selection in demand seekers. However, the Demand Group x Fronto-Parietal Component slope interaction was statistically unreliable ($b= 0.60, t(46) = 1.75, p = .09$), indicating that the relationship between Fronto-Parietal Component slope and demand avoidance did not change as a function of demand group.

The Default Mode Component, by contrast, did show this relationship (Fig 6). Pearson correlation between Default Mode Component β estimate slopes and demand avoidance behavior across Demand Avoider participants was significant ($r(24) = .49, p = .01$). In other words, those individuals who had shallower Default Mode Component slopes (i.e., more positive, less deactivated) across effort levels during task execution showed greater demand avoidance. This relationship was absent for Demand Seekers ($r(22) = -.28, p = .19$). The Demand Group* Default Mode Component slope interaction was statistically reliable ($b= 0.99, t(46) = 2.44, p = .02$). In follow up tests, however, we did not find that these differences between the Fronto-Parietal Component and Default Mode Component ROIs were statistically reliable. Thus, we find

evidence for a relationship between effort selection and the change in activation across effort levels in Default Mode Component in Demand Avoiders. However, though we do not find evidence of this relationship for Fronto-Parietal Component, we also do not have evidence that this difference distinguishes these brain networks.




Figure 1.10. Principal Component cluster correlations with demand avoidance behavior are plotted for (A) Positive (“Fronto-Parietal Component”) and (B) Negative PC1 (“Default Mode Component”). Results for Demand Avoiders are plotted in red. Results for Demand Seekers are plotted in black. The ‘Effort level’ analysis tested whether the change in effort selection rates was correlated to the change of brain activity in each PCA ROI during task execution across effort levels. Default Mode Component showed a positive relationship between the slope of change in this network across effort levels and effort selections. This was only the case in the Demand Avoider group. Error bars plot standard error of the mean. * $p < .05$.

Independent ROIs' relationship to behavior

Network12: Neither for Demand Avoiders nor Demand Seekers Network12 Yeo ROI activity during task-execution predicted deck selection (Demand Avoiders: $M_{\beta}=-0.10$, $SD_{\beta}=0.08$, $t(156) = -1.31$, $p = .19$; Demand Seekers: $M_{\beta}=0.04$., $SD_{\beta}=.09$, $t(144) = 0.46$, $p = .64$).

Consistent with our whole brain ROI analysis, Pearson correlation between Network12 β estimate slopes and demand avoidance behavior across Demand Avoiders, $r(24) = .30$, $p=.14$ or Demand Seeker participants was not significant, $r(22) = -.21$, $p=.32$, indicating that individual variability in demand avoidance behavior cannot be explained by activity change across effort levels in the Network12 ROI.

Network16: Neither for Demand Avoiders nor Demand Seekers Network16 Yeo ROI activity during task-execution predicted deck selection (Demand Avoiders: $M_{\beta}=-0.01$, $SD_{\beta}=0.38$, $t(156) = -0.17$, $p = .87$; Demand Seekers: $M_{\beta}=-0.08$., $SD_{\beta}=.47$, $t(144) = -0.78$, $p = .43$).

In congruence with our whole brain ROI analysis, Pearson correlation between Network16 β estimate slopes and demand avoidance behavior across Demand Avoider participants was significant, $r(24) = .40$, $p=.04$, indicating that those individuals who had shallower Network16 β s across effort levels during task execution showed greater demand avoidance. As expected, Pearson correlation between Network16 β estimate slopes and demand avoidance behavior across Demand Seeking participants was not significant, $r(22) = -.18$, $p=.38$. Unlike Negative PC1, the Demand Group*Network16 slope interaction was not statistically reliable ($b= 0.95$, $t(46) = 1.72$, $p = .09$).

Unlike Negative PC1, Network 16 slope did not predict the selection slope after the effects of ER and RT were removed (Network 16 slope: $b = .54$, $t(22) = 1.80$, $p = .09$, $R^2=.07$),

indicating that change in Network 16 activity during effort execution did not predict change in effort selections over and beyond what could be attributed to the change in performance.

VS: Neither for Demand Avoiders nor Demand Seekers *VS* ROI activity during task-execution predicted deck selection (Demand Avoiders: $M_{\beta}=0.06$., $SD_{\beta}=0.07$, $t(156) = 0.87$, $p = .39$; Demand Seekers: $M_{\beta}=-0.09$., $SD_{\beta}=0.11$, $t(144) = -0.77$, $p = .44$).

Pearson correlation between *VS* β estimate slopes and demand avoidance behavior across Demand Avoider, $r(24) = .22$, $p=.28$ or Demand Seeker participants was not significant, $r(22) = .01$, $p=.97$, indicating that individual variability in demand avoidance behavior cannot be explained by activity change across effort levels in the *VS* ROI and that the observed modulation of Parametric- ROI on demand avoidance behavior in Demand Avoider group cannot be explained by the effect of *VS* on avoidance behavior.

In order to complement our findings at the subject level, we also ran a group level analysis that regressed each subject's demand avoidance rates onto the linear relationship between BOLD activity and effort levels across the whole-brain. Therefore, each participants' overall demand avoidance rates were entered as covariate into the group level analysis that estimated betas at each effort level for the contrasts that tested for a parametric linear relationship between task execution and brain activity. This analysis showed that there were no brain regions that significantly correlated with demand avoidance rates across participants, indicating that linear changes in brain activity as a function of our effort manipulation were independent of effort choices.

Additionally, we tested for the relationship between connectivity within and between Fronto-Parietal Component and Default Mode Component and demand avoidance, in order to see if individual differences in connectivity predict effort selections. As FPN recruitment linearly

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

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

	p(Easier Task Selection)	Demand Avoidance Slope
Within-FPC Connectivity	$r(48)=.14, p=.33$	$r(48)=.13, p=.39$
Within-DMN Connectivity	$r(48)=.12, p=.41$	$r(48)=.11, p=.46$
FPC-DMN Connectivity	$r(48)=-.08, p=.59$	$r(48)=-.08, p=.56$

Table 1.8. Brain connectivity – behavior correlations for Within-FPC, Within-DMN and FPC-DMN connectivity, as clusters defined by PCA.

Thus, neither connectivity within nor between these networks correlated with behavioral avoidance rates (Table 1.8.). Thus, together with our analysis on connectivity change across

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

Effortful task-execution naturally comes with performance costs such as increased error rates and response times. Thus, the higher likelihood of time on task or errors might lead to demand avoidance as well. In order to differentiate the role of performance measures from that of Default Mode Component in demand avoidance, we next regressed out the role of performance from brain-behavior relations. A stepwise multiple regression was conducted to evaluate whether Default Mode Component slope, RT slope and ER slope were necessary to predict selection slope in Demand Avoiders. At step 1, the Default Mode Component slope entered into the regression equation and was significantly related to selection slope (Default Mode Component slope: $b = .59$, $t(24) = 2.78$, $p = .01$). The multiple correlation coefficient was .49, indicating approximately 24% of the variance of the selection slope could be accounted for by Default Mode Component slope. At step 2, ER and RT slopes were regressed out of the selection slope, and Default Mode Component slope entered into the regression on the residuals of ER and RT slopes. Default Mode Component slope still reliably predicted the selection slope even after the effects of ER and RT were removed (Default Mode Component slope: $b = .58$, $t(22) = 2.41$, $p = .03$, $R^2 = .16$), indicating that change in Default Mode Component activity during effort execution predicted change in effort selections over and beyond what could be attributed to the change in performance. However, we do note that the *a priori* ROI definition of DMN no longer predicted demand avoidance when the effects of performance were controlled ($p = .09$). As discussed below, this difference likely arose due to the substantially broader ROI for the *a priori* definition.

Next, we tested the effects of performance and task-switching on Fronto-Parietal Component activity. For Demand Avoiders, task-switching probability predicted Fronto-Parietal Component recruitment over and beyond performance measures ($M_{\beta}=0.20$, $SD_{\beta}=0.06$, $t(153.86) = 3.10$, $p < .01$). For Demand Seekers, both ER and task-switching probability explained Fronto-Parietal Component recruitment (ER: $M_{\beta}=-0.90$, $SD_{\beta}=0.32$, $t(15.5) = -2.76$, $p = .01$; TS: $M_{\beta}=0.24$, $SD_{\beta}=0.07$, $t(60.43) = 3.51$, $p < .001$). Comparison of the models with ER and without ER showed that, the inclusion of ERs in the model predicting Fronto-Parietal Component significantly improved goodness of fit compared to the model that only included task-switching probability ($X^2(2) = 12.09$, $p < .01$), indicating that for Demand Seekers, Fronto-Parietal Component recruitment was best explained by a combination of task-switching probability and error signals. Thus, for demand seekers, effort tasks that yielded greater task-switching probability and smaller error rates predicted greater Fronto-Parietal Component recruitment during task-execution. In other words, FPN recruitment increased during blocks participants showed greater accuracy on difficult trials.

Overall, we have observed that change in Default Mode Component predicts change in effort selection as a function of effort levels in Demand Avoiders over and beyond what can be explained by performance measures ER and RT. Fronto-Parietal Component activation did not predict effort selections in Demand Avoiders, and marginally predicted demand seeking behavior in Demand Seekers.

Discussion

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

Consistent with the idea that effort might register as disutility, we observed that reward network, including VS, vmPFC and PCC, linearly reduced activity as a function of increasing effort execution, providing support for the ‘cost of effort’ hypothesis and consistent with prior reports (Schouppe et al., 2014a; Botvinick et al., 2009). We also observed a saturating decrease in VS activation as effort levels increased, consistent with observations by Botvinick et al. (2009). However, VS activity during task execution did not predict effort selections (Supplementary Results 8), indicating that effort costs, as tracked by VS during task-execution, do not directly influence demand avoidance behavior in DST paradigms in the absence of monetary reward. Nevertheless, we only scanned during the “Test phase” when the association of effort with the different task conditions had already been established. It is reasonable to hypothesize that the costs computed by VS may be differentially important during the “Learning phase” when effort associations are being acquired. This will be an important target of future research.

To test the hypothesis that effort costs arise from engagement of the cognitive control system, we manipulated effort by varying cognitive control demands in a task-switching paradigm. Prior task-switching literature observed greater activity in the dorsal striatum and frontoparietal cortex during the performance of task-switching compared to the performance of task-repeat trials (Vallesi et al., 2015; Yin et al., 2015; Piguet et al., 2013; Gu et al., 2008; McCool et al., 2008). Consistent with the literature that implicates FPN activity in tasks demanding cognitive control, we observed that most of the variance in our data could be explained by a linearly increasing activation function in this network with increasing task-switching probability. However, in our study, we showed that blocks that involved the performance of greater task-switching trials yielded greater FPN recruitment, but not dorsal striatum. This discrepancy might be due to differences in fMRI designs between our study and the literature. Studies that implicated dorsal striatum adopted an event-related fMRI design which separated task-switching trials from task-repeat trials. However, our study adopted a block design during the effort execution phase, which resulted in weaker ability to detect the time course of the hemodynamic response during task-switching trials. Thus, the neural response during blocks of greater task-switching trials still included trials that tasks repeat.

Additionally, the neural function underlying effort execution took a linear trend for DMN and FPN. However, VS activity was best explained by a quadratic trend. It is possible that this quadratic trend observed in VS activity tracks the task-switching probability based entropy across task levels. For example, the averseness associated with task uncertainty might override the averseness associated with the overall task-switching probability of a task. Thus, the cost of effort, as tracked by VS, might be task uncertainty. In order to test this hypothesis, we re-ordered selection rates according to the level of task uncertainty of each effort level and tested for a

linear trend across subjects. We observed that there was no significant trend across selection rates over task levels as re-ordered by uncertainty level of the task, suggesting that task uncertainty does not underlie effort costs.

It was shown that expected value of different options is represented at specialized brain regions, which then get transformed into common expected subjective values (ESV) for comparison in the dopaminergic neurons in the vmPFC, and partially in VS and PCC (Levy & Glimcher, 2011). All three brain regions have been implicated in our study, where we observed a declining trend of activity across the execution of manipulated effort levels in both demand groups. If these regions simply tracked the subjective values of each effort level and participants were solely making their effort decisions based on these values, then we would expect only demand avoiders to show a reduction in the recruitment of this expected value network. However, we did not find any significant group differences in the activity of any of these regions. It is possible that trial-trial effort choices are not represented by overall subjective values. Future studies should adopt temporally precise techniques to record brain activity on a trial-by-trial basis. On the other hand, participants might be making their choices based on more than one decision parameter alone. While ESVs might track the cost of executing each effort option, the learnability of the effort level might constitute the motivational value of that option. For example, it was suggested that while phasic dopaminergic signals in SN and VTA project to NAcc and VMPFC in order to track changes in unexpected external rewards, salience-coding neurons project to NAcc and dLPFC in order to drive curiosity and direct attention to the challenging option (DeYoung, 2013). Alternatively, curiosity-driven neural activity might be tracked by tonic, and not phasic dopaminergic activity, given that the reward is not externally generated, and hence it is not unexpected (Domenico & Ryan, 2017). Therefore, in the absence

of external reward, both the averseness of task performance and the subjective salience of the effort option might drive the activity of the dopaminergic reward system, albeit in different modes of activity. Future studies should manipulate or record subjective reports of curiosity in addition to averseness of an option and test the involvement of the dopaminergic reward system in representing these parameters.

Under the cost hypothesis, we would expect that effort seekers would be different from Demand Avoiders in four ways: less FPN activity, smaller error rates, and shorter response times or greater VS activity. These neural and behavioral markers would indicate that Demand Seekers either exerted less effort than Demand Avoiders or that they valued effort exertion more than Demand Avoiders. However, the results show that there was no difference in error rates or neural recruitment between demand groups. Moreover, contrary to the expectations of the cost hypothesis, Demand Avoiding participants had overall faster response times compared to Demand Seeking participants, and if anything, Demand Seekers who showed greater demand seeking behavior showed an increasing trend in FPN recruitment across the execution of effort levels. These results suggest that behavioral and neural proxies for cognitive control recruitment may not obligatorily register as something costly. Instead, difficult tasks that yield greater change in FPN recruitment might be preferred by demand seeking individuals.

Further, the linear increase in FPN activity due to the cognitive control manipulation was not predictive of an increase in demand avoidance behavior. And, we found that for Demand Seekers, performance-related factors like error likelihood predicted FPN activity even after the effects of our experimental cognitive control manipulation were controlled. Thus, we observe only partial support for the cost of control hypothesis, at least as it might stem from engagement of the FPN network. Performance indices like error-likelihood and RT likely correlate with

cognitive control demands in a task, but they could relate to other factors, as well. Thus, to avoid circularity in relating cognitive control to FPN activation and demand avoidance, the task switching manipulation provided an independent definition of cognitive control. And using this definition of cognitive control demands, we did not find evidence that changes in activation in FPN attributable to cognitive control were related to demand avoidance.

We note that our observations are not necessarily inconsistent with prior reports. A previous study that also utilized a task-switching paradigm to manipulate cognitive effort (McGuire and Botvinick, 2010) showed that a change in FPN activity predicted demand avoidance. However, this study used a binary comparison, where a change in brain activity and task selections between low and high effort conditions was not corrected by the change in intermediate task levels. By contrast, our parametric manipulation provided greater sensitivity to separate changes in FPN activation due to the cognitive control manipulation versus other factors related to performance.

Additionally, our study is the first to test the neural mechanisms underlying effort selection and effort execution in the same experiment. Whole-brain analysis of the Selection Epoch showed that execution and the selection of effort did not yield activity in the same brain regions. While the execution of effort yielded parametric effects in FPN, DMN and VS, the selection of effort yielded parametric effects only in the primary motor cortex, suggesting that the selection and the execution of cognitive effort recruit different neural systems. These differences might help explain the observed discrepancies in the literature regarding differential the engagement of brain networks depending on whether one is performing or anticipating an effortful task. Future studies should attend to temporal dynamics of effort performance.

There are no previous studies that parametrically manipulated implicit cognitive effort

demands in the absence of monetary reward as we did here. However, some studies have manipulated explicit cognitive effort demands using reward discounting procedures (Chong et al., 2015; Apps et al., 2015; Massar et al., 2015; Libedinsky et al., 2013; Westbrook et al., 2013). For example, Cognitive Effort Discounting (COG-ED) paradigm (Westbrook et al., 2013) shows that participants parametrically discount task value by increasing demands on cognitive effort. This is consistent with our finding that probability of selecting a task parametrically decreases for demand avoiding participants. However, unlike the discounting procedures, we also show that the same task demands were found increasingly valuable by demand seeking participants. This individual difference suggests that effort is not obligatorily aversive or coded as disutility, but rather effort signals can be interpreted positively or negatively depending on the participants.

The discrepancies in these observations might be due to three notable differences between this procedure and the one employed in discounting paradigms, beyond our inclusion of fMRI participants. First, in COG-ED, participants were explicitly aware of the difficulty difference between levels, as the instructions regarding each effort level significantly differed due to the nature of the effort task. In our version of DST, the task instructions were the same across effort levels, and as post-experimental debriefing inventories showed most participants were not explicitly aware of a difficulty difference between effort levels. Second, in COG-ED, participants made their choices without immediately executing the chosen option, while our procedure required the participants to execute the effort level immediately following their selections. Third, in COG-ED, participants made a choice between an option that yields low effort and low reward, and an option that yields large effort and large reward. Our task was performed in the absence of explicit feedback or reward. Directly testing these dimensions of decision making may be fruitful directions for future research.

Further, a notable advantage of the effort discounting procedures such as COD-ED is that they can estimate the subjective point of equality for each effort level by requiring choices between options that yields low effort and low reward, and an option that yield large effort and large reward. However, in doing so, COG-ED procedure must assume that the inclusion of extrinsic reward does not influence the intrinsic evaluation of cognitive effort. The advantage of the current task is that it aims to test the intrinsic cost/value associated with effort in the absence of feedback or incentives, and thus is directly testing the influence of intrinsic motivation that would drive voluntary effort selections in the absence of monetary reward. Nevertheless, without a discounting approach, we are unable to assess the unit of exchange between effort and monetary reward across effort levels.

A related open question concerns whether demands on cognitive control are themselves effortful, or rather, cognitive control tends to consume more time and it is the time on task that the brain taxes as costly. Kool et al. (2010) showed that participants avoided switching between tasks even when staying on task meant longer task durations. Consistent with those findings, we observed that the effect of task-switching probability predicted effort selections, however in opposite directions for each demand group. While Demand Avoiders avoided those tasks that yielded greater task-switching, Demand Seekers chose them. Future studies should aim at equating error-rates and time-on-task across effort levels in order to empirically differentiate the effects of task performance from cognitive control recruitment on effort decisions.

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

showed the highest demand avoidance rate. The relationship between DMN and effort selections persisted even after performance measures such as RT and ER were controlled. We note that though FPN did not show these effects, as discussed above, we also did not find a significant difference between these networks. So, we cannot conclude that this is a qualitative difference such that DMN is more determinant of effort selection than FPN. These potential differences may be an important avenue for future research.

It is not evident from our study what drives this relationship between DMN activity, though prior studies of this network may offer some clues. While FPN recruitment has been shown to underlie cognitive control, DMN has been associated with a wide range of internal mental functions, including self-referential thoughts, mind-wandering, and episodic future planning (Buckner et al., 2008; Weismann et al., 2006). Inhibition of DMN has been argued to support task-related attention (Spreng, 2012), while the inability to inhibit DMN has been related to dysfunctional cognitive processes such as rumination (Nejad et al., 2013) and pathologies such as depression (Anticevic et al., 2012; Lemogne et al., 2012). Accordingly, in our study, those participants who showed higher effort avoidance rates could have had increased difficulty deactivating DMN activity or relied more on the processes it entails, which in turn might have registered as a cost. Alternatively, recent studies (Guggenmos et al., 2016, eLife, Gherman & Philiastides, BioRxiv) showed that in the absence of external feedback, confidence might register as reward. This confidence signal has been shown to recruit the same brain regions that are associated with reward processing. It is possible that DMN recruitment in our study also tracks task-related confidence signals. However, under this assumption, the observation that demand avoiding participants with greater DMN recruitment at greater effort levels would incur greater effort costs would imply that effort avoidance increases as task-related confidence increases for

difficult tasks. Future studies should seek to replicate this discovery and to determine what factor drives this change in DMN activity across effort levels.

It is possible that this DMN activity is related to reward and/or value processing that predicts effort selections. A common observation in the literature is that vmPFC positively tracks subjective value and predicts subjects' preferences between choices (Rushworth et al., 2011). Given this functional association, the 'cost of effort' hypothesis would predict that individuals will avoid tasks that yield the least activity in the reward network. However, in our study, we observed that individuals who showed the least reduction in this network showed the greatest demand avoidance. In addition to linear reductions in vmPFC and PCC activity with increasing effort, independent ROI definitions of VS showed that VS reduces its activity across increasing effort, however, DMN but not VS predicted demand avoidance behavior (see Supplementary Results 8). These results together suggest that while reward network tracks effort in congruence with the 'cost of effort' hypothesis, this cost does not predict effort selections.

From a learning account, DMN de-activation could be associated with the speed of automatization of task strategies. Demand avoiding participants who have not fully automatized task-specific policies, as tracked by DMN, might experience greater effort costs. On the other hand, individual differences in DMN de-activation during task-execution could be associated with differences in the employment of task-irrelevant processes. Demand avoiding participants who engage in task-irrelevant thoughts more often might experience greater effort costs. Regardless of the underlying mechanism that mediates the relationship between DMN and demand avoidance, it is important to note that DMN - demand avoidance relationship was unique to demand avoiders, although DMN showed a similar trend across the execution of effort levels in both groups. If the same neural processes underlie effort-based decisions in both demand

avoiding and demand seeking participants, we would have expected DMN to explain effort decisions in both groups. However, it is possible that demand avoiding and demand seeking participants differ from each other in the way they utilize the same neural process during decision-making, e.g. DMN de-activation.

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

Finally, individual differences are an important variable considered here that has not received much attention in the literature. In our fMRI task, we observed high individual variability in effort avoidance, such that roughly half the participants were better characterized as Demand Seekers than Demand Avoiders. This rate of variability has not been reported previously in the literature and was not the case in our own behavior-only task (Experiment 1).

The conflicting results might be due to at least three reasons: 1) self-selection bias in fMRI experiments, 2) context effects in fMRI settings, 3) time-of-year effects for subject pool participants who showed a tendency to participate in our experiments around their finals. The behavioral participants in our behavior-only study mostly consisted of undergraduate participants who volunteered for course credit. However, our fMRI participants consisted of a more variable sample who volunteered for fMRI scanning at different months of the year in response to ads placed around campus. A recent study showed that PET scan volunteers significantly scored

higher in sensation-seeking trait compared to behavior-only volunteers (Oswald et al., 2013). Sensation-seeking is also a trait that positively correlated with Need for Cognition (NfC), a self-report inventory that tests effort avoidance and that negatively correlates with Effort Discounting scores (Schuller, 1999).

We confirmed in a follow up study that an independently recruited sample of volunteers responding to advertisements for fMRI scored significantly higher on NfC compared to subject pool participants (see Supplementary Text). Of course, this study does not provide direct evidence that the fMRI participants reported in this paper had higher NfC, as we did not administer this to them at the time they were scanned. Thus, we can only provide the general observation that NfC is higher among participants responding to ads for fMRI. Future studies should directly test the relationship between the demand avoidance scores in this parametric version of DST and NfC.

The demonstration of individual differences is central to the current study for three reasons. First, we think that individual differences in effort expenditure behavior and brain activity has not been fully explored in the effort literature, and ours is the first study to show that the same performance and brain activity utilized by Demand Avoiders to avoid effort, can be used by Demand Seekers to choose effort. Second, for the half of our sample that sought effort, increased frequency of task-switching and FPN recruitment registered as something valuable rather than costly. This observation was counter to expectations and strains against strong versions of the ‘cost of effort’ hypothesis. Third, the results show that the DMN-behavior relationship we observe in Demand Avoiders is not present in Demand Seekers, which indicates that DMN-based activity can register as an exclusive cost signal that is tracked by demand avoiding individuals.

Regardless of the source of the individual differences, we were able to test how the factors affecting effort-based decisions differed across these groups. Although there were no differences in brain activity and error rates between demand groups during task execution, Demand Avoiders avoided those tasks that yielded greater task-switching, while Demand Seekers chose them. This discrepancy in the way separate demand groups utilized the same information indicates that FPN, and task performance could influence effort-based value-based decisions generally, even if not necessarily in terms of a cost signal. On the other hand, reduced DMN de-activation across effort levels influenced effort avoidance only in demand avoiding participants, indicating that change in DMN activity across effort levels entered into effort-based decisions exclusively as a cost signal. However, note that the findings of our PCA-based ROI definition of DMN and a priori definition of DMN are not completely congruent. While both definitions showed that those participants who showed less DMN inhibition across effort levels showed greater demand avoidance rates, this relationship persisted only for the PCA-defined DMN when the effects of performance were controlled. The effect was only marginal using the a priori DMN definition. This discrepancy is likely due to differences in the brain regions these ROIs encompass. The a priori DMN definition covers a broader region than that encompassed by the PCA regions.

A potential caveat given that our effort task parametrically manipulated implicit effort costs, is that a participant with consistent but arbitrary rankings might show a linear slope across effort selection while they also showed a bias towards or against selecting the easier task. One source of such stable preferences unrelated to effort could stem from a preference for the symbol shapes that cued effort levels. However, symbol orders were randomized across participants, making it unlikely that a group of participants would systematically prefer symbol shapes that

also correlated with our effort manipulation. Further, these preferences would also evidently correlate with changes in brain during execution that we report here. Thus, effort avoidance provides a parsimonious account of the full set of results we present here. Nevertheless, we do note that it is possible that effort selections could have been driven by a fixed, but arbitrary, preference ranking.

In conclusion, we adopted a parametric version of DST in order to test the ‘cost of control’ hypothesis and explore the neural mechanisms that underlie effort execution. We have observed that the reward network reduces activity in response to executing more effortful tasks, in congruence with the ‘cost of effort’ hypothesis. DMN but not FPN predicted effort avoidance behavior in demand avoiding participants, indicating that control recruitment as indexed by FPN recruitment does not underlie demand avoidance. As behavioral task performance can represent the required cognitive control to perform well at a task, increased time-on-task and error-likelihood could constitute costs associated with the opportunity-cost of time and error-avoidance. However, neither error rates or response times predicted effort selections in either Demand Avoiders or Demand Seekers. Additionally, high individual variability in effort avoidance behavior in our task showed that the direction greater task-switching probability influenced effort-based decisions depended on the demand group, indicating that proxies of cognitive control do not exclusively influence effort-based decisions as a cost. On the other hand, it was shown that reduced DMN de-activation was an exclusive cost signal that differentiated Demand Avoiders from Demand Seekers, promising to be a new avenue for future effort-based research.

Supplementary text

Trait differences in effort avoidance between subject pool and fMRI volunteers

In order to see if there are trait differences in effort avoidance behavior between behavioral subject pool participants and fMRI volunteers, we administered the Need for Cognition (NfC) inventory on an independent sample of Psychology subject pool participants and fMRI volunteers. None of the participants in this study went through the demand avoidance procedure explained in the main text. The sample consisted of new study participants who were invited to the lab to complete a brief survey. The only difference between participants was how they contacted the lab for research, either by responding to posted advertisements for fMRI experiments or by seeking course credit through the Psychology subject pool.

Participants: Eleven participants (aged 18-35; 7 female) with normal or corrected to-normal vision were recruited for the behavioral demand avoidance experiment (Experiment 1). All participants were free of neurological or psychiatric conditions, were not taking drugs affecting the central nervous system, and were screened for contraindications for MRI. Participants provided informed consent in accordance with the Research Protections Office at Brown University. Participants were recruited either from CLPS research participation system (SONA) fulfilling course credits or called back upon contacting our lab to volunteer for a paid fMRI study. Subject pool participants were compensated by 1 course credit/hr and fMRI volunteers were compensated for \$10/hr for their participation.

Methods: Participants completed NfC inventory following the performance of fMRI or behavioral tasks. NfC Scale (Cacioppo et al., 1984) is a questionnaire that measures an individual's tendency for exerting mental effort. In this questionnaire, participants rate their agreement with 18 statement on a 9-point scale. The scale values range between 'very strong disagreement' to 'very strong agreement'. The statements involve the performance of problem

solving or thinking, such as 'Thinking is not my idea of fun'. Out of the 18 statement, 9 of them are reverse scored. A high score on this scale indicates an individual's tendency to pursue mentally effortful activities.

Results: Independent samples t-test results showed that fMRI volunteers ($M = 21.20$, $SD = 10.61$) scored significantly higher on NfC compared to subject pool participants ($M = 2.17$, $SD = 10.40$), $t(9)=3.00$, $p=.02$. The difference of 19.03 scale units indicated a large effect, (Cohen's $d=1.79$).

Chapter 3: Testing the neural systems underlying cognitive effort avoidance during effort learning

Introduction

People avoid cognitively effortful tasks. This tendency is evidenced in demand selection paradigms, where participants are asked to choose between two tasks that differ in the demand level of cognitive effort. In these paradigms, the probability of selecting the easier task increases with increasing experience in both effort tasks (Kool et al., 2010). Accordingly, it was argued that what underlies demand avoidance behavior is the learned cost associated with exerting cognitive control (Shenhav et al., 2013; Botvinick, 2007). This hypothesis was supported by neural evidence illustrating that reward network brain activity in VS following effort execution was smaller compared to that of the easier task (Botvinick et al., 2009), and demand avoidance in DST positively correlated with frontoparietal control network activity during effort execution across participants (McGuire & Botvinick, 2010). Moreover, Nagase et al. (2018) exhibited that the cost of both cognitive effort options is learned incrementally via trial-by-trial experience and activates anterior cingulate cortex and frontoparietal regions that are usually associated with cognitive control.

However, not all effort decisions are binary. More often in real life, we are confronted with multiple contexts that vary in the cognitive demands they require. Our previous study (Sayali & Badre, 2018) modified the demand selection paradigm by parametrically manipulating task-switching probability across six effort levels, and separating effort learning from effort selection. Our results indicated that reward network negatively and frontoparietal control network positively scaled with increasing cognitive effort demands, although neither network activity predicted effort selections. Instead, change in default-mode network activity across effort

levels during effort executions underlie demand avoidance behavior, indicating that presence of multiple versus binary effort options might yield differences in the neural mechanisms explaining demand avoidance. However, this study tested the relationship between reward network and effort selections after effort costs had already been learned. Thus, it is not yet known what neural mechanisms underlie the learning of multiple effort options and whether these mechanisms underlie demand avoidance behavior.

Moreover, it was illustrated that (Schouppe et al., 2014a) in the absence of reward, voluntary selection between effort options activate the reward network differently than that of imposed effort anticipation, indicating that the way in which participants engage with effort might influence the neural mechanisms that underlie effort learning. However, Nagase et al. (2018) observed brain activity during the learning of cognitive effort tasks which the participant voluntarily selected on a trial-by-trial basis, and therefore it is not yet known what neural mechanism underlie the learning of effort cost when effort tasks are imposed.

In order to address these gaps in the literature, the current study observed brain activity during the learning of six parametrically varying cognitive effort tasks, using the parametric version of DST. In the current study, the learning of effort options has been separated by effort selections and thus, during the learning of effort costs, execution of each effort task was imposed rather than voluntary. Consistent with previous literature, we predict to observe changes in control and reward networks, with increasing imposed effort experience, that explain demand avoidance.

Methods

Participants

Previous study (Sayali & Badre, 2018) results demonstrated that fMRI volunteers score higher than subject pool participants on a self-report inventory (NfC) that measures daily engagement in mentally challenging activities. Based on this study, we determined a cutoff score of 16/72 on NfC to recruit only demand-avoiding participants. Three hundred seventeen (317) adults were pre-screened by Need for Cognition (NfC) Inventory, in person. All participants were compensated for their time either monetarily or with course credit depending on their enrollment method. Additionally, during their visit, we also gathered health information regarding their eligibility for fMRI. One hundred nineteen (119) individuals scored lower than 16/72 on NfC, and thus passed our demand avoidance criteria. Among the 119, 78 of them both met the fMRI eligibility criteria, agreed to be re-contacted, and had not participated in a previous study that tested effort-based decision-making. Among the 78 eligible individuals, only 29 of them agreed to participate in our fMRI study. Among the 29 participants, 5 of them cancelled their appointment on the day of their scan. Therefore, a total of 24 right-handed adults (aged 18-35; 12 female) who scored lower than 16/72 on NfC, with normal or corrected to-normal vision were recruited for the fMRI experiment. Two participants withdrew from the fMRI study and did not complete all experimental phases. Two participants reported during debriefing that they failed to follow instructions. One participant was excluded due to head movement greater than our voxel size across all sessions. Thus, in total, data from 5 participants were excluded prior to analysis. Therefore, 19 participants were included in the behavioral and fMRI analyses. All participants were free of any neurological or psychiatric conditions, were not taking drugs affecting the central nervous system, and were screened for contraindications for MRI.

Participants provided informed consent in accordance with the Research Protections Office at Brown University.

Behavioral Task

Participants performed a parametric variant of the demand selection task (DST; Fig 1). In an initial Learning phase, the participant associated virtual card “decks”, denoted by particular symbols, with a certain effort level, through experience with trials drawn from that deck. Then, in a second Test phase, participants chose which deck they wished to perform (Selection epoch) and then performed the trials drawn from that deck (Execution epoch). A distinguishing characteristic of this version of the DST, relative to prior versions, is that we varied the effort level parametrically over six levels based on the frequency of task-switching required by a deck. This allowed us to characterize the functional form of changes in behavior or fMRI signal due to our cognitive control manipulation. We now provide detail on each of these phases of the behavioral task.

The Task “Decks”

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

1-4 and 6-9 were used with equal frequency across both tasks and were randomized for order of presentation.

In order to manipulate cognitive effort across decks, we varied the frequency of task switching required by a particular deck (Fig 1b). The probability of switching tasks from trial to trial, within a sequence drawn from each deck, increased across 6 effort levels: 8%, 23%, 38%, 54%, 69%, 85%, respectively. Note that the lowest effort level still included one task-switch. In this manner, all effort levels included the same number of parity and magnitude trials. Thus, any differences in effort preference among these could not be attributed to differences in difficulty or effort between the tasks themselves. Likewise, we did not include “pure blocks” of only parity judgments or only magnitude judgments, as doing so would have made these decks qualitatively different from all other decks, in terms of the amount of each task being performed. Having the lowest effort level include one task-switch ensures that the only parameter that changes between effort levels is the frequency of task-switching. A higher frequency of task-switching is associated with a higher experience of cognitive effort (Monsell, 2003), and has been shown to be aversive (Arrington & Logan, 2004). At the beginning of each block, a shape was presented for 1 second to indicate which deck was being performed, and this shape was also tiled in the background throughout the performance of the block. Participants were told that this shape represented the symbol on the back of the virtual card deck from which trials for that sequence were drawn. Each effort level was associated with a particular deck symbol. Participants were not told about this relationship, but could learn through experience that certain decks were more effortful than others to perform. The mappings between deck symbols and effort levels were randomized across participants.

Practice Phase

During an initial “Practice Phase” participants gained familiarity with the trial structure, the categorization decisions, and the color and response mappings. After being given instructions regarding the categorization and response rules (Fig. 1a), they practiced two 13-trial blocks. Each task sequence included the 3 sec long presentation of a deck symbol, and the subsequent performance of 13 consecutive task-switching trials. In the first block of the Practice Phase, feedback was presented, after the button press, as either 'Correct' in pink, 'Incorrect' in yellow, or 'Please respond faster' if the participant failed to answer in 1.5 sec. In the second block of the Practice phase, feedback was omitted, as would be the case in the actual experiment. The deck symbols that participants encountered in the Practice Phase were for practice only and were not presented during the learning or test phases to avoid any association of error feedback with a particular deck symbol.

Learning Phase

In the Learning Phase (Fig. 1c), participants learned the association between 6 deck symbols and an effort level. Each deck was performed 15 times, in random order. This phase was performed in the fMRI and participants used a MRI-compatible button box to indicate their decisions. The Execution trials were optimized as blocks. The Anticipation and Execution events were separated in time by a jittered time interval (mean 4 sec) so that signal related to each could be analyzed independently. The Learning phase was separated into two, approximately 20 minute-long scanning runs.

Test Phase

In the Test Phase (Fig. 1d), two decks were presented and the participant chose the one they would like to perform (Selection epoch). The Test Phase took place outside the magnet, in the behavioral testing room. The participants were told to freely choose between decks prior to a 3 sec deadline. We note that in contrast to other DST variants (Gold et al., 2015), participants in this task were not told about the difficulty manipulation, in order to avoid biasing choices based on participants' explicit beliefs about effort. Once the participant made their selection, the selected deck turned blue and both options remained on the screen until the end of the 3 sec deadline. In the event of a non-response, the same choice pair was re-presented at the end of the entire experiment, until they made a decision. Each pair was selected from the set of fifteen unique (un-ordered) pair combinations of all six decks, excluding self-pairing (e.g., deck #1 paired with deck #1). Each deck was presented either on the left or the right side of the screen, counterbalanced for location across trials. The Selection epoch was followed by the execution of the selected effort-level task sequence (Execution epoch). The sequence of events in this epoch was the same as during the Learning phase. The Test phase was separated into four, approximately 15 minute-long blocks. In each run, each pair was presented 3 times in a pseudo-randomly intermixed order, making a total of 180 decision trials across 4 blocks.

Behavioral Data Analysis

Trials with response times below 200 ms were excluded from further analysis. Execution trials on which participants missed the response deadline were also excluded from further analysis (approximately %1 of execution trials in both phases). Response times were calculated using only correct trials. Task performance was analyzed with effort level and experimental phases (Learning or Test Phase) as within-subject variables. One participant did not select from

the fifth effort level during the Test Phase, and thus was not included in the effort*phase performance analysis.

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

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

MRI procedure

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

with padding, while visual stimuli were rear projected and viewed with a mirror attached to the head coil.

fMRI Analysis

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

A temporal high-pass filter of 128 (.0078 Hz) was applied to our functional data in order to remove noise. Changes in MR signal were modeled under assumptions of the general linear model (GLM). Two GLMs were devised: a linear effort-level GLM and an independent effort-level GLM. Both GLMs included nuisance regressors for the six motion parameters (x,y,z,pitch,roll,yaw) and four run regressors for the ‘Linear Effort-Level GLM’ and one run regressor for the ‘Independent Effort Level GLM’. The number of run regressors was different across GLMs, because ‘Independent Effort GLM’ included the regressor for each effort level separately.

Linear Effort-Level GLM

The linear effort-level GLM tested which voxels in the brain parametrically increased or decreased linearly with effort level. Two event regressors were used. Execution events were

modeled as a boxcar that onset with the presentation of the first trial stimulus of the sequence and ended with the participant's response to the final item. Thus, the duration of this boxcar varied as a function of response time. Second, the Selection event regressor modeled each anticipation event with a fixed boxcar of three sec. We used parametric modulators on these event regressors to test the linear effect of effort level. The Execution event regressor was modulated by: (a) an Effort Level parametric regressor corresponding to the effort level of that task sequence (1 through 6); and (b) a Prediction Error parametric regressor that scaled with the difference between the Expected Value and Actual Value of the exerted effort on that trial. The Anticipation event regressor was modulated by: (a) an Effort Level parametric regressor based on the presented effort level deck symbol (1 through 6); and (b) an Expected Value regressor that scaled with the estimated expected value of the upcoming effort execution. Note that as implemented by SPM, each parametric regressor includes only the variance unique to it and shared with those ordered after it. Thus, for example, the Effort Level regressor includes variance explained over and above that shared with the Execution event regressor. The Execution and Anticipation event regressors, along with their parametric modulators, were modeled separately for each scanning run, within the GLM.

Independent Effort Level GLM

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

a 3 sec boxcar at the onset of deck symbol presentation. Boxcars and durations were the same as in the linear effort level model. In this GLM, two run epochs and a linear drift over the whole experiment were included, as nuisance regressors.

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

ROI analysis

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

Independent ROIs

We defined a fronto-parietal control network ROI defined from a functionally neutral group ([Network 12] Yeo et al., 2011) along with a VS ROI (Badre et al., 2014), as VS activity has been consistently observed in the cognitive effort literature (Botvinick et al., 2009; Schouppe

et al., 2014a). We also included a Default Mode Network ROI from a functionally neutral group ([Network 16] Yeo et al., 2011). A priori DMN regions included ventromedial prefrontal cortex, orbitofrontal cortex, posterior cingulate cortex and parts of precuneus. A priori FPN regions included bilateral lateral prefrontal cortex, bilateral parietal cortex and SMA.

We additionally defined dACC ROIs (Shenhav et al., 2014, 2016) in order to test the involvement of dACC in learning effort costs, as this region has been shown to be involved in tracking the expected value of effort.

Finally, we defined a primary motor cortex (PMC) ROI (Moore et al., 2000) in order to test the involvement of this region in anticipating the upcoming effort level execution during the Learning phase. In our previous study, this ROI showed an increasing trend across the selection of the effortful tasks during the Test Phase.

PCA ROIs

Our previous study (Sayali & Badre, 2018) was the first to parametrically manipulate implicitly learned effort cost/values in DST in fMRI. In this study, in order to have an unbiased analysis of the neural response function across effort levels, we adopted a data-driven Principal Component Analysis (PCA) approach on the whole brain (126,866 voxels), in order to explore the shape of the neural functions of different brain areas that respond to effort execution during the Test Phase and their relation to demand avoidance behavior. In order to observe the relationship between the activity of these brain regions during the learning of effort costs and subsequent demand avoidance behavior, we defined fronto-parietal control and DMN ROIs based on the previous study's Test Phase PCA.

Brain – behavior analysis

In order to understand which brain regions underlie effort execution, during Learning, that predict effortful task selection, all ROIs during the Learning Phase were related to the selection behavior during the Test Phase. We analyzed this relationship in two ways: 1) correlating the change in effort selection rates to the change of brain activity during task execution across effort levels ('Effort level'); and 2) correlating individual effort selection rate to the brain activity during task execution at that effort level ('Individual task selection'), unranked with respect to task switching frequency. The former analysis focuses on activity change in a brain region due to the task-switching manipulation. The latter considers how overall activity levels, independent of a linear task-switching manipulation, correlate with effort avoidance.

For the 'Effort level' analysis, we computed the slope of the linear function for the selection rates, β estimates of a ROI across effort levels. We multiplied the slope we obtained from the selection rates by -1, in order to indicate the decreasing likelihood to select more effortful tasks, which is conceptualized as 'demand avoidance' in our task. We then correlated the demand avoidance slopes and the β estimate slopes for all demand avoiders, in order to address the role of this region in explaining individual variability in demand avoidance behavior.

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

In order to test the effects of brain and performance on selection behavior, while accounting for individual variability within demand groups, we adopted the following mixed effects hierarchical regression model:

$$y_i = X_i(\beta + s_{j[i]} + m_{[i]}) + \epsilon$$

$$s_j \sim N(0, \Sigma_S), \quad \epsilon \sim N(0, \sigma)$$

Here, y_i is the i th observed selection rate, X_i is a matrix of predictors (brain activity) and covariates (performance), β_i is a vector of coefficients (as in conventional linear regression), $j[i]$ is the subject of the i th observation, so that $s_{j[i]}$ is a subject-specific perturbation of all of the coefficients.

This fitting procedure was implemented using the R package `blme`, which includes the `lme4` package (Bates et al., 2014) and performs maximum-a-posteriori estimation of linear mixed-effects models.

Results

Demand avoidance behavior

Overall, participants were highly accurate during task execution across both phases (mean error: 13% in the Learning Phase, 13% in the Test Phase), and participants missed the deadline on only a few trials (2.6% of Learning phase trials, $SE=0.01$, 1.2% of Test phase trials, $SE=0.03$).

RTs, but not error rates improved across phases (RT: $F(1,17)=43.43$, $p<.001$, $\eta_p^2=.72$; errors: $F(1,17)=0.22$, $p=.65$, $\eta_p^2=.01$). Both RT and error rates showed a linear effect as a result of effort condition, increasing with a higher proportion of task switches across both switch and repeat trials. There was a significant effect of effort on error rates ($F(5,85)=20.07$, $p<.001$, $\eta_p^2=.54$) and RT ($F(2,14,36.37)=143.19$, $p<.001$, $\eta_p^2=.89$). The increase in both was linear over effort levels (errors: $F(1,17)=40.82$, $p<.001$, $\eta_p^2=.71$; RT: $F(1,17)=215.79$, $p<.001$, $\eta_p^2=.93$). These increases were also significant for both switch and repeat trials, tested separately, with the exception of Switch RT (switch errors: $F(1,17)=36.10$, $p<.001$, $\eta_p^2=.68$; repeat errors: $F(1,13)=63.46$, $p<.001$, $\eta_p^2=.83$; switch RT: $F(1,17)=2.05$, $p=.17$, $\eta_p^2=.11$; repeat RT: $F(1,17)=104.90$, $p<.001$, $\eta_p^2=.86$).

As only those participants who self-reported to be effort avoidant by the NfC inventory participated in this experiment, participants, on average, avoided the harder task 67% ($SE=0.02$) of the time (Figure 2.1.A). Demand avoidance, observed in the current study, was significantly different than chance across participants ($t(18)=6.87$, $p<0.001$). The probability of selecting the easier task did not reliably change across Test blocks, although there was a marginal effect of block on demand avoidance ($F(3,54)=2.73$, $p=.53$, $\eta_p^2=.13$), indicating that participants' effort selection were relatively consistent across time. This demand avoidance trend was greater than

the demand avoidance observed in our previous studies (Sayali & Badre, 2018), for which we did not use a pre-screening method, indicating that overall NfC scores are indicative of demand avoidance in our paradigm. However, NfC scores across participants did not correlate with the overall probability of selecting the easier task ($r(17)=-.26, p=.29$), indicating that self-reported NfC scores do not directly reflect between-subject variability in demand avoidance, as tested in our paradigm. The lack of a direct match between NfC reports and demand avoidance scores in our task might be due to the implicit nature of effort decisions in our design, which we discuss further in the Discussion.




Figure 12.1. Demand avoidance across participants. 1A) Overall probability of selecting the easier task. All participants selected the easier task more than 50% of the time. 1B) Probability of selecting an effort level across participants. The probability of selecting a task significantly decreased with increasing effort level.

Consistent with the diagnostic feature of this paradigm, there was a significant effect of effort level on choice behavior ($F(5,90)=8.86$, $p<.001$, $\eta_p^2=.33$); as such, higher effort levels were associated with higher avoidance rates than were lower effort levels (Figure 2.1.B). This pattern across effort levels was fit by a linear effect ($F(1,18)=47.88$, $p<.001$, $\eta_p^2=.73$). The probability of selecting the easier task also significantly changed depending on the difference between effort levels given as options ($F(2.30,41.35)=5.14$, $p=.001$, $\eta_p^2=.22$), thereby, larger differences increased the probability of choosing the easier task. This effect was linear across effort levels ($F(1,18)=13.48$, $p=.002$, $\eta_p^2=.43$). Decision time was similar across effort levels ($F(5,85)=1.74$, $p=.06$, $\eta_p^2=.09$) and there was no effect of difficulty difference on decision times.

Additionally, we tested the effects of performance and task-switching probability on effort selection behavior using hierarchical regression analysis. In order to include the entire sample in this analysis, one participant's missing 5th effort level RT and ER values, during the Test Phase, were estimated by fitting the first 4 effort levels' performance measures by a linear function. Across participants, RTs, but not error rates, during task-execution in the Learning Phase predicted effort selections in the Test Phase (RT: $M_\beta=-0.81$., $SD_\beta=0.21$, $t(114) = -3.83$, $p < .001$; errors: $M_\beta=0.08$., $SD_\beta=0.19$, $t(114) = 0.43$, $p = .67$). The same pattern was true for the Test Phase RTs and error rates (RT: $M_\beta=-0.56$., $SD_\beta=0.20$, $t(114) = -2.83$, $p < .001$; errors: $M_\beta=-0.17$., $SD_\beta=0.19$, $t(114) = -0.90$, $p = .37$), indicating that participants selected those tasks that yielded the least time-on-task.

Next, we tested the effects of task-switching probability on the selection rates in the presence of task performance, in order to see if our linear effort manipulation predicted demand avoidance over and beyond RT during task-execution. The effects of RT in either phase were no longer significant when task-switching probability was added to the model. The only significant

predictor was task-switching probability (Learning Phase: $M_{\beta}=-0.51$, $SD_{\beta}=0.09$, $t(114) = -5.21$, $p < .001$; Test Phase: $M_{\beta}= -0.51$, $SD_{\beta}=0.08$, $t(114) = -5.97$, $p < .001$), indicating that effort selection behavior was explained better with our linear task manipulation than task performance, during task execution in either experimental phase.

In summary, consistent with our previous study (Sayali & Badre 2018), we showed that NfC pre-screening scores are predictive of participants' overall tendency to avoid effortful tasks in our design. However, individual NfC scores are not directly predictive of the within-subject variability in demand avoidance, indicating that our design measures additional effort-related processes that are not explained by self-reports. Consistent with our previous reports of demand-avoiding participants, participants avoided effortful tasks that yielded slower response times. However, this effect of time-on-task was abolished in the presence of task-switching probability, indicating that participants' choices were mainly driven by our control manipulation across effort levels.

The functional form of univariate brain activity over effort levels

We investigated the pattern of activation in *a priori* ROIs retrieved from our previous study, in order to examine their underlying neural function across the execution of increasing effort levels across blocks in the Learning Phase and compare it to the neural function we have observed in our previous study that observed brain activity in the Test Phase (Sayali & Badre, 2018). We defined 'Frontoparietal component' and 'Default-Mode component' ROIs based on the principal component analysis results of the Test Phase in our previous study. In the previous study, we adopted a PCA approach to the whole brain analysis to identify voxels in the whole brain that share a similar underlying activation functional form across the execution of effort levels. From 126,866 voxels included in the analysis, the first PC explained 57% of the variance,

from which we extracted 3000 voxels that loaded the most positively (Frontoparietal component) and most negatively (Default-Mode component) on each component and extracted their β s, in order to observe the functional form of each PC. We also defined ROIs in the FPN (Network12; Yeo et al, 2011), DMN (Network16; Yeo et al., 2011), Primary Motor Cortex (Moore et al., 2000), and/or ventral striatum (VS; Badre et al., 2014). Then, using the “Independent Effort Level GLM”, we estimated the activation in the network for each effort level separately.

Consistent with prior reports, the Frontoparietal component ROI during task execution showed a linear trend across effort levels (Figure 2.2.A). There was a significant effect of effort on Frontoparietal component β estimates, $F(5,90)=11.12$, $p< .001$, $\eta_p^2 =.38$. The β estimates linearly increased with increasing effort requirements of the task, $F(1,18)=26.02$, $p< .001$, $\eta_p^2 =.59$. Additionally, Frontoparietal component β estimates decreased across blocks, $F(1,18)=11.29$, $p= .003$, $\eta_p^2 =.39$, indicating that prolonged effortful task execution reduced frontoparietal network recruitment. These results were consistent with independent ROI definition of FPN (linear effort effect: $F(1,18)=6.09$, $p= .02$, $\eta_p^2 =.25$; session effect: $F(1,18)=10.89$, $p= .004$, $\eta_p^2 =.38$), albeit with relatively smaller effect sizes (Figure 2.2.C). There was no effect of effort on either Frontoparietal component or FPN during the Anticipation epoch (Figure 3A, frontoparietal component: $F(3.51,63.21)=1.22$, $p= .31$, $\eta_p^2 =.06$; Figure 3C, FPN: $F(3.51,63.23)=0.95$, $p= .46$, $\eta_p^2 =.05$).




Figure 2.2. Univariate ROI results during task execution. 2A) Frontoparietal component, as defined by Sayali & Badre (2018), displayed a significant linear trend across the execution of effort levels. Frontoparietal component recruitment declined across task blocks. 2B) Default-Mode component, as defined by Sayali & Badre (2018), demonstrated a significant linear trend across the execution of effort levels. 2C) FPN, as defined Yeo et al. (2011), exhibited a significant linear trend across the execution of effort levels. 2D) DMN, as defined Yeo et al. (2011), did not show a reliable effect of effort. 2E) VS, as defined by Badre et al. (2014), displayed a declining trend across the execution of effort levels. 2F) PMC as defined by Moore et al. (2000) exhibited a quadratic trend across the execution of effort levels. PMC recruitment decreased across experimental blocks. 2G) Overlays the DMN (top) and FPN (bottom) ROIs on the canonical brain as they are separately defined by Yeo et al. (2011) and Sayali & Badre (2018).

Consistent with prior reports, the Default-Mode component showed a negative linear trend in activation across effort levels (Figure 2.2.B). There was a significant effect of effort on Default-Mode component β estimates, $F(5,90)=3.31$, $p=.01$, $\eta_p^2=.16$. The β estimates linearly decreased with increasing effort requirements of the task, $F(1,18)=11.59$, $p=.003$, $\eta_p^2=.39$. These results were not consistent with the independent definition of DMN (Figure 2.2.D) effect of effort level: $F(5,90)=1.21$, $p=.31$, $\eta_p^2=.06$), indicating that task execution in our design recruits a specific component of DMN that connectivity-based definitions of DMN do not necessarily explain (Figure 2.2.G). There was no effect of block on β estimates, $F(1,18)=1.19$, $p=.29$, $\eta_p^2=.06$, indicating that Default-Mode component recruitment did not change as a function of effort experience. There was a marginal effect of effort on both Default-Mode component and DMN during the Anticipation epoch (Figure 2.3.B, Default-Mode component: $F(5,90)=2.25$, $p=.06$, $\eta_p^2=.11$; Figure 3D, DMN: $F(5,90)=2.34$, $p=.05$, $\eta_p^2=.12$).




Figure 2.3. Univariate ROI results during task anticipation. 2A) Frontoparietal component as defined by Sayali & Badre (2018) was similar across the anticipation of effort levels. 2B) Default-Mode component as defined by Sayali & Badre (2018) displayed a marginal linear trend across the anticipation of effort levels. 2C) FPN as defined Yeo et al. (2011) was similar across the anticipation of effort levels. 2D) DMN as defined Yeo et al. (2011) demonstrated a marginal effect of effort. 2E) VS as defined by Badre et al. (2014) was similar across the anticipation of effort levels. 2F) PMC as defined by Moore et al. (2000) exhibited a significant declining trend across the anticipation of effort levels.

An independent ROI in VS was tested. This ROI showed a declining trend across effort execution in the Test Phase in our previous study. Also in the Learning phase, VS did not show a reliable negative linear trend (Figure 2.2.E). There was a marginal effect of effort on VS β estimates, $F(5,90)=2.14$, $p = .07$, $\eta_p^2 = .11$, although the effect size is similar across studies

(previous study $\eta_p^2 = .11$), indicating that with greater sample size this effect could have been reliable. There was no effect of experimental block on β estimates, $F(1,18)=1.36$, $p = .26$, $\eta_p^2 = .07$. There was no effect of effort on VS during the Anticipation epoch (Figure 3E), $F(5,90)=1.16$, $p = .33$, $\eta_p^2 = .06$.

An independent ROI in Primary Motor Cortex (PMC) was also tested. This ROI showed an increasing trend across effort selection in the selection epoch of the Test Phase in our previous study. However surprisingly, in the Learning phase, PMC showed a declining trend during the anticipation of increasingly effortful tasks (Figure 2.3.F). There was a marginal effect of effort on PMC β estimates, $F(5,90)=2.92$, $p = .02$, $\eta_p^2 = .14$. PMC β estimates linearly decreased across the anticipation of increasingly more effortful tasks, $F(1,18)=8.52$, $p = .01$, $\eta_p^2 = .32$. There was no effect of experimental block on β estimates, $F(1,18)=1.45$, $p = .24$, $\eta_p^2 = .08$. Inconsistent with findings of our prior study, there was also an effect of effort on PMC during task-execution, $F(1,18)=6.93$, $p = .02$, $\eta_p^2 = .28$. However, the effect of effort on PMC during task-execution (Figure 2.2.F) was best fit by a quadratic trend, $F(1,18)=9.76$, $p = .01$, $\eta_p^2 = .35$, in which the 5th and 6th effort levels yielded smaller β estimates compared to mid-effort levels, but not low-effort levels. There was also an effect of block on β estimates, $F(1,18)=2.16$, $p = .02$, $\eta_p^2 = .28$, in which PMC recruitment declined during task execution across task blocks.

Overall, univariate brain activity results indicate that *a priori* definitions of FPN and DMN, based on our previous study, display similar β estimates of brain activity during the Learning of effort tasks, while connectivity-based independent definitions of the same networks yield either smaller or unreliable effects, indicating that the task-specific portions of these bigger networks are consistently recruited. As such, consistent with prior findings, FPN recruitment increases, and DMN recruitment decreases with effort execution during the Learning Phase. FPN

recruitment also decreases with increased effortful task-execution experience in the Learning Phase. Consistent with prior reports, we found a declining trend of effort execution on VS activity during the learning of effort tasks; although this effect was not reliable, the effect sizes were comparable across studies, indicating that greater sample size could have yielded a reliable effect. Neither FPN nor VS exhibited a reliable effect of effort during the anticipation of effortful tasks, suggesting that the anticipation and the execution of effort rely on different mechanisms. One surprising finding concerns PMC recruitment during effortful task anticipation. This brain region previously showed an increasing trend during the selection of the more effortful task, in our previous study. However, here we show that PMC recruitment decreases during the anticipation of increasingly more effortful tasks, indicating that voluntary selection and the anticipation of imposed effortful tasks yield different neural responses.

Brain – Behavior Analysis

Having established the functional form of activity change across effort levels, we tested whether brain regions tracking effort execution predicted effort selections. Two correlation techniques were adopted to examine the type of brain-behavior relationship: 1) relationship between brain activity and effort selection rate at each effort level ('Individual task selection'), 2) relationship between change in brain activity and change in effort selection rates across effort levels ('Effort Level'). All analyses were conducted on all definitions of FPN, DMN, and VS.

Our previous study showed that Fronto-Parietal Component during task execution in the Test Phase across effort levels did not reliably predict change in effort selection rates in either group while DMN did for demand avoiding participants. However, during the Learning Phase, none of FPN, DMN, or VS ROI definitions in neither experimental phase (including the change

in β estimates across experimental blocks) predicted effort selections or demand avoidance. Thus, we report the results of these analyses in Supplementary Results.

Discussion

The current study aimed to test the brain mechanisms involved in learning the value of imposed effortful tasks and how learned values of effort guide effort decisions. In order to infer the mechanisms associated with demand avoidance, we first pre-screened participants with NfC, an inventory that measures subjective reports of daily engagement in effortful tasks and then observed brain activity in the Learning phase of our parametric version of DST. The ‘cost of control’ hypothesis of effort proposes that effort inherits an implicit cost, which is derived from recruited cognitive control during effortful task performance. Thus, we predicted that the value of effort would be learned by reward and control-related brain regions, and inform demand avoidance rates.

First, the current study aimed to test the involvement of previously observed brain regions during the learning of effort. Our parametric version of DST separates effort-based decision-making into two phases. In the first ‘Learning’ phase, participants learned the association between effort tasks and symbols without making selections. In the ‘Test’ Phase, participants chose between two symbols that were associated with implicitly learned effort levels on each trial and perform the associated effort task. Our previous study (Sayali & Badre, 2018) observed brain activity during the Test phase, after effort tasks had already been learned, and demonstrated that FPN activity increases and DMN activity decreases with the performance of increasingly effortful tasks. However, DMN, but not FPN, predicted demand avoidance. The current study provides consistent evidence that FPN increases and DMN decreases with the performance of increasingly effortful tasks during the learning of effort. The results indicate that similar brain mechanisms are involved during the execution of effort task in both Learning and Test phases of the parametric version of DST. However, the recruitment of neither network during the Learning

phase predicted demand avoidance in the Test phase. It is possible that DMN is recruited as a cost signal when demand selections are being made, or when effort learning has been finalized. Future studies should test the relationship between DMN recruitment and effort costs across imposed and free-choice trials, within the same experimental phase.

Our previous study also illustrated that reward-related brain activity decreased with the performance of increasingly effortful tasks, in the parametric version of DST, after effort tasks had already been learned. However, reward-related brain activity, as indexed by VS, during task performance did not predict demand avoidance. The current study aimed to test the involvement of the same brain region in effortful task performance during the learning of effort, and its involvement in predicting subsequent demand avoidance behavior. Consistent with the idea that effort registers as an implicit cost, VS, a region implicated in reward processing, showed a negative trend across effortful task performance, during the learning of effort tasks, with a comparable effect size to that of the previous study that observed VS after effort tasks had been learned. If effort costs were learned with experience, we would predict VS to reduce its activity with increasing effort experience. However, VS recruitment during task execution did not change across experimental blocks. Moreover, there was no modulation of VS activity by effortful task anticipation either, indicating that the expected value of the subsequent effort task is not represented in VS. These results are consistent with Prevost et al. (2010) who observed brain activity during a physical effort discounting paradigm, and demonstrated that VS did not scale with the subjective value of effort, indicating that VS signal might not be discounted by the expected cost of effort.

Additional ROI results of reinforcement learning model of effort costs showed that VS was not engaged in the learning of efforts costs on a trial-by-trial basis either. These results contradict

Nagase et al. (2018), which observed that reward network activity tracks the expected cost of the chosen option. The lack of VS engagement in the learning of effort costs in our study might be due to a lack of effort selection during the time brain activity was observed. For example, Schouppe et al. (2014a) showed that the direction of VS engagement during the anticipation of effortful tasks changes as a function of whether the upcoming effort task is voluntarily selected or imposed. Future studies should directly test the involvement of VS in effort learning as a function of voluntary selection.

‘Cost of control’ hypothesis (Shenhav et al., 2013; Botvinick et al., 2007) suggests that people avoid effortful tasks to the extent those tasks recruit brain regions that are associated with cognitive control. Consistent with this hypothesis, a previous report (McGuire & Botvinick, 2010) illustrated that subsequent self-reports of demand avoidance can be predicted by average FPN (but not dACC) during the performance of effort tasks. Contrary to these findings, a recent study (Nagase et al., 2018) observed that dACC, as well as FPN, is recruited in tracking the expected cost of effort. It was argued that dACC is recruited during the learning of effort costs only when effort selections are made. Our study tests this assumption by observing brain activity during the learning of effort tasks, when effort tasks are imposed rather than selected, and by relating brain activity to demand avoidance behavior in a separate Test phase.

Building upon this theory, the Expected Value of Control (EVC; Shenhav et al., 2013) theory suggests that the same conflict signal informs reward-processing regions of the brain as a cost signal, and reduces the value of the task to the extent of required control engagement. Thus, EVC would predict the involvement of FPN and dACC during the anticipation epoch of our design. In accordance with this hypothesis, Nagase et al. (2018) observed frontoparietal-insular cortex to scale with the expected cost of the chosen cognitive effort option. However, our study did not

observe any neural activity during the anticipation of effort, even when specific FPN and dACC ROIs were tested. These results suggest that dACC and FPN engage in conflict monitoring aspects of effort learning, but not in the representation of expected cost when effort trials are imposed rather than voluntarily selected.

The discrepancy between our findings and the previous reports might be due to differences in experimental design and implemented brain-behavior analysis. McGuire and Botvinick (2010) first reported the relationship between average FPN activity and demand avoidance behavior in a task with only two effort levels. As both our previous and current study observed, average FPN activity does not predict effort selection in a parametric version of DST whereby neural function underlying effort exertion is corrected by intermediate effort levels. Instead, the involvement of FPN in effort-based decision-making might be more nuanced, such that FPN relates to effort decisions only during the learning of effort values on a trial-by-trial basis.

Second, Nagase et al. (2018) also used a binary effort paradigm, and observed effort-learning-related brain activity when participants were making decisions regarding which effort task to choose. In our study, participants learn each effort level without making selections, thus the learning of the value of effort is not affected by demand avoidance. Our results suggest dACC is not exclusively involved in the demand decision process, but instead might be involved in a broader effort learning mechanism.

One surprising finding was regarding the involvement of primary motor cortex in the anticipation of effortful tasks. Our previous study found that primary motor cortex increased activity with the anticipation of increasingly effortful tasks, when effort tasks are voluntarily selected. Our current study observed that primary motor cortex decreased activity with the anticipation of increasingly effortful tasks when those tasks are imposed. Furthermore, in both

studies, where the effort task was voluntarily selected or imposed, primary motor cortex was the only predictor of the upcoming task difficulty. Additional ROI analyses concluded that neither region associated with decision-making such as FPN, dACC or VS were engaged during the anticipation of effort.

There are at least two hypotheses regarding the engagement of sensorimotor system in decision-making. First, it was argued that sensorimotor system represents competing options, and is involved in response deliberation (Cisek, 2007; Shadlen et al., 2008). Second, it is possible that sensorimotor system is not instrumental in decision-making itself, but instead represents the information retrieved from upstream regions. Supporting the former hypothesis, recent animal evidence (Thura & Cisek, 2014) demonstrated that primary motor cortex represents the voluntary commitment to a choice by displaying a growing urgency signal for the selected option while suppressing others. Research on human participants (Gluth et al., 2012) also indicated that motor cortex represented the value of the selected option when participants were deciding to buy or reject a stock offer. Additionally, a physical effort discounting paradigm (Prevost et al., 2010) observed that left primary motor cortex activity was positively correlated with the subjective value of the effortful reward, indicating that motor cortex is involved in motor preparation for task engagement as a function of the value of effort.

Thus, it is possible that in our study, primary motor cortex activity represents the value of the upcoming effort task by a growing commitment signal when the decision is voluntary and suppresses it when it is imposed. Furthermore, we have not observed the involvement of FPN in the anticipation of effort, suggesting that primary motor cortex does not simply represent the decision process spill-over from upstream regions. Future studies should directly test the causal

role of motor cortex involvement in cognitive effort-based decision-making by using brain stimulation techniques.

Finally, our previous study (Sayali & Badre, 2018) showed that participants who volunteered to participate in fMRI had greater NfC scores compared to subject pool participants. Accordingly, in our previous study, we had found that half of our fMRI sample consisted of demand seekers. In order to test the learning of effort costs, the current study aimed to recruit only demand-avoiding participants. Thus, we used NfC inventory to pre-screen participants before fMRI recruitment. As the entire sample exhibited demand avoidant behavior, we demonstrate that NfC scores are predictive of a group level demand avoidance behavior in the parametric version of DST. However, individual NfC scores did not predict overall demand avoidance scores across individuals, indicating that between-subject variability in demand avoidance, as tested in our design, cannot be explained by self-reports of demand avoidance, as measured by NfC.

This observation contradicts the relation between NfC scores and an effort-discounting task (Westbrook et al., 2013) scores that adopted working memory retrieval to parametrically manipulate effort demands. In this study, participants were explicitly aware of the difficulty difference between options when making selections. However, our parametric version of DST tests for demand avoidance without the explicit awareness of an effort manipulation. This is evidenced by a lack of post-experimental subjective reports of awareness. However, NfC is an explicit measure of daily engagement in mentally challenging tasks. It is possible that participants' self-reports of effort avoidance do not directly explain their implicit tendency to avoid cognitive effort, as operationalized here as task-switching probability.

Additionally, during the subject recruitment of this study, we have observed that demand-avoiding participants overall were either not willing to participate in an fMRI study, did not want to be re-contacted for future studies, or did not keep their appointment (21%) at a rate that was not previously observed by researchers for the recruitment of other studies, in our lab. Future studies should attend to this observation when testing demand-avoiding participants in the fMRI, as the majority of the participants recruited for fMRI research might consist of demand-seeking individuals due to self-selection bias.

One of the limitations of the current study is the representativeness of the sample. Prescreening survey indicated that around half of all participants who participated in NfC were categorized as demand avoiding. However, only a small portion of demand avoiding individuals agreed to participate in our fMRI study. Our previous study also showed that participants who volunteer for fMRI testing also score higher on NfC compared to participants who only volunteer for behavioral testing. Together, these observations suggest that demand avoiding participants who volunteered for fMRI testing in our study might not be representative of demand avoiding group characteristics.

Another limitation of the current study is regarding the manipulation of cognitive control demands in our study. We adopted a task-switching paradigm to vary task difficulty. Task-switching manipulation has been shown to have subtle effects on difficulty awareness (Sayali & Badre, 2018), by which it stands in contrast to other manipulations such as working memory retrieval (Westbrook et al., 2013) or arithmetic summation (Nagase et al., 2018). Additionally, parametric manipulation of difficulty in our study minimized the subjective experience of a difference between task levels, potentially making effort-learning more difficult in terms of conscious processing. Thus, an important question concerns whether the number of trials during

the Learning Phase was enough for a reliable acquisition of association between deck symbols and difficulty levels. Given this caveat, current study was not able to reliably computationally model learned effort costs/values, as opposed to studies (Nagase et al., 2018) that adopted an arithmetic summation task to manipulate task levels. While this study was able to reliably model effort learning, it tested for effort learning in a binary (easy vs. hard) design and only modeled performance measures such as RT and ER and objective measures such as task difficulty. Across models, they assumed that participants aimed at maximizing their response time, accuracy performance, or their progress across trials. However, it was observed that calculating the deviation of each task-switching block's latency and accuracy, in relation to the participants' overall accuracy and overall latency, would improve estimates of executive control, and explain individual variability in other cognitive control tasks, such as the antisaccade analogue task performance. Thus, one way to increase the validity of their model would have been via implementing an Δ_{accuracy} variable that updates the current trial's accuracy, in relation to the participant's maximum overall accuracy. For example, participants might also engage in a strategy that prioritizes one over the other, such as the speed-accuracy trade-off (Samavatyan & Leth-Steenson, 2009). Accordingly, it was argued that (Hughes et al., 2014) using only RTs or ERs to model recruited cognitive control might lead to biased, or unreliable estimates. Instead, it is recommended to use a compound score that includes both the RT and the ER in order to measure individual differences in executive control. Future studies should parametrically manipulate arithmetic tasks to manipulate task levels and directly manipulate trial-by-trial accuracy, in order to incorporate deviations in accuracy into models of task performance for estimating cognitive control.

In conclusion, consistent with our prior reports, we observed that average FPN activity during task execution does not predict demand avoidance. This indicates that the involvement of FPN in demand avoidance might be more nuanced than representing the overall cost of effortful tasks in the absence of immediate decision-making regarding effort. Consistent with this interpretation, we showed that FPN is not involved in representing the expected cost of the upcoming trial either. Instead, primary motor cortex anticipated the difficulty of the upcoming trial by decreasing its activity, in the current study, when effort tasks are imposed. This observation was in contrast with our previous study in which we observed that PMC increased recruitment with voluntarily selection of the more effortful task. Future studies should test the causal role of PMC in cognitive effort avoidance.

Chapter 4: Testing the role of dopamine status during effort learning and selections in Parkinson's Disease patients

Introduction

Effort avoidance is operationalized as the preference for the easier task when given an option of which task to perform. This phenomenon can be measured in the laboratory using a paradigm called 'Demand Selection Task' (DST; Kool et al., 2010). In this task, participants choose between two options that require different degrees of cognitive effort, manipulated in different ways, such as the frequency of task-switches, the number of items to maintain in memory, or the presence of conflicting stimuli.

Recent work (Botvinick et al., 2001; Kool et al., 2010; McGuire & Botvinick, 2010) has established that when given a choice most participants prefer the less effortful task, while subjective reports following effort exertion indicate that increased cognitive effort is accompanied with feelings of frustration and negative affect. The 'cost of control' account of cognitive effort (Expected Value of Control (EVC) theory, Shenhav et al., 2013) proposes that the recruitment of control is determined as a function of the probability of an outcome minus the cost of the control required to obtain the reward. The currency of this cost can be a function of higher error rates, increased response conflict, or a greater opportunity cost of time during effort performance. For example, in DST, individuals avoided tasks more to the degree that the task yielded greater response times during task performance (Kool et al., 2010).

Based on these results, the 'cost of control' hypothesis of effort argues that using cognitive control to complete a task is penalized by an intrinsic cost, which reduces the reward value of that task's outcome via dopaminergic activity. Furthermore, it has been argued that conflict monitoring brain regions inhibit dopamine release during effort, via their connections

with the striatum, the input nuclei of BG, situated in the dopaminergic system (Botvinick et al., 2001). The transient pause in dopamine bursts experienced during control-demanding events, leads to avoidance learning via dopaminergic frontostriatal loops, and consequently to the habit of demand avoidance. Consistent with this model, the ventral striatum was seen to decrease its activation following the performance of an effortful task during the reward receipt (Botvinick et al., 2009).

A recent study that adopted a probabilistic learning task while manipulating the degree of effort required to reach the reward state, demonstrated that effort diminishes the reward value of the task or enhances aversion to punishment, depending on the basal dopaminergic state of the individual, as manipulated by a dopaminergic antagonist (Cavanagh et al., 2014). A causal test of the cost hypothesis was performed on PD patients using the same paradigm, indicating that the cost of effort was enhanced for patients who were off their dopaminergic medication, confirming the causal role of dopamine in learning effort costs (Cavanagh et al., 2017). Therefore, all studies that tested the effect of dopaminergic manipulation on cognitive effort avoidance did so in the presence of external reward. However, it was observed that in the absence of reward, voluntary selection of the more difficult task yields greater ventral striatum engagement (Schouppe et al., 2014a).

PD results in the death of DA neurons in the substantia nigra, a structure situated in the nigrostriatal pathway, so the mesolimbic ‘reward’ system remains relatively intact (Dauer & Przedborski, 2003). Thus, DA medication in PD patients overstimulates the mesolimbic DA system, and leads to established alterations in reward-based decision-making (Santangelo et al., 2013). For example, the modulation of DA intake in PD patients influences the sensitivity to negative consequences when making decisions in reinforcement learning paradigms (Frank et al.,

2004). However, it is unknown whether the dopaminergic involvement during effort-based decision-making primarily supports the learning or the expression of effort costs, or both. All studies that tested the effect of dopaminergic manipulation on cognitive effort avoidance did so without separating the learning from the test phase. To provide causal evidence for the role of DA on learning effort costs, we will test PD patients on and off their DA medication while they learn (Learning) versus while they make decisions (Test) based on effort (Figure 3.1.a).

Under the ‘cost of control’ hypothesis, we make three main predictions (Figure 3.1.b). First, being off-dopaminergic medication will increase demand avoidance rates, compared to age-matched control participants and participants on-medication. Thus, demand avoidance rates will scale as $PD_{OffOff} > Control > PD_{OnOn}$. Second, given previous research (Cavanagh et al., 2017), we predict that diminished dopamine will increase the cost of effort during learning. Thus, the cost will be higher for those patients off their medication during the Learning Phase, resulting in a main effect of dopamine on effort learning. Third, diminished dopamine might also increase sensitivity to effort costs during decision-making, resulting in a main effect of medication on both phases. Accordingly, those participants who are off their medication on both days will show the greatest demand avoidance.

Finally, the interactive effects of differential dopamine states on effort learning and decision-making are not yet known. Therefore, we will additionally explore interaction effects between phases, comparing PD Group 2 demand avoidance rates to PD Group 4, and comparing PD Group 1 demand avoidance rates to PD Group 3.




Figure 3.1. A) Depiction of dopaminergic drug manipulation in PD patients across experimental phases, B) demand avoidance predictions for PD patients under the ‘cost of control’ hypothesis.

Methods

Participants

For Experiment 1, eleven right-handed adults (aged 18-35; 5 female) with normal or corrected to-normal vision were recruited for the behavioral demand avoidance experiment using the Brown research subject pool. All participants were compensated by course credit.

For Experiment 2, forty-eight Parkinson’s disease (PD) patients and fifteen age-matched healthy controls were recruited from the Movement Disorders Clinic at Rhode Island Hospital and Butler Hospital. All patients were diagnosed with PD by a movement disorder specialist with greater than 90% diagnostic certainty using UK Brain Bank criteria. All patients were on a dopaminergic therapy with a minimum levodopa equivalent daily dose (LEDD) of 220 mg. Out of all PD patients, 8 of them were using a dopamine agonist (either pramipexole (.25-.75 mg) or ropinirole (.5 mg)). The rest of the patients were using levodopa (Carbidopa-Levodopa 25-100). All patients had a mini-mental status score (MMSE) of ≥ 24 . Patients with atypical parkinsonism, poor response to levodopa, concomitant diagnosis of another degenerative brain disorder, MMSE score below 24, aphasia or acalculia, or an active psychiatric disorder were

excluded. All patients were in good overall health with good or corrected vision enabling them to see the screen and perform the study for 30 minutes in the off-medication state. All age-matched healthy adults were free of neurological or psychiatric conditions and were not taking drugs affecting the central nervous system. Six patients withdrew from the study on the second day and did not complete all experimental phases. Four patients failed to follow instructions. Three patients reported comorbid anxiety or depression, and were excluded from the study. Three participants' data were corrupt due to technical failure. In total, data from 16 participants were excluded. Thus, 34 PD patients (aged 52-87; 9 female) and 13 age-matched control participants (11 female) were included in the analyses. All participants were paid a total of \$25 for their participation. Participants provided informed consent in accordance with the Research Protections Office at Brown University.




Figure 3.2 Schematic description of the modified demand selection task. A) An example of four deck cues with their associated effort level pairing, based on the probability of task-switching in 13 trials from that deck. B) Example trials of the categorization task with a digit printed on a virtual card. The color cued which task to perform. The response mappings shown below the digits are for the reader's reference and did not actually appear on the screen. C) Schematic of block events during the Learning Phase. First, a symbol icon is presented for 1 sec, after which the participant completes 13 categorization trials with the associated card tiled on the

background. D) Schematic of block events during the Test Phase. First, the participant chooses between two virtual decks that are associated with two different effort levels, with a deadline of 3 sec. Next, the participant executes the effort level associated with the selected option, while the selected card is tiled on the background.

Behavioral Task

Participants performed a modified version of the demand selection task (DST; Fig 3.2) across two consecutive days before noon on each day. With a full factorial design, levodopa intake was alternated across experimental phases for PD patients. Depending on their group, patients were asked to withhold their medication in the morning. These patients were tested at least 12 hours after their last medication intake, during a practically defined ‘Off’ period. Upon their arrival on the first day, all participants were assessed by MMSE scale, a 30-point questionnaire that measures cognitive impairment.

On the first day, participants performed the Learning phase. In this phase, the participant associated virtual card “decks” – denoted by particular symbols – with a certain effort level through experience with trials drawn from that deck. Then, on the second day, in the Test phase, participants chose which deck they wished to perform (Selection epoch) and then performed the trials drawn from that deck (Execution epoch). A distinguishing characteristic of this version of the DST, relative to prior versions, is that we had two of four decks yielding the same easy deck, and the other two yielding the same difficulty task. This allowed us to compare selection rates between decks pairs of the same difficulty, as well as of different difficulties. Any bias toward a deck, in pairs of same difficulty, would indicate that the participant used heuristics that are independent of our effort manipulation. We now provide detail on each of these phases of the behavioral task.

The Task “Decks”

Throughout all phases of the experiment, participants performed trials drawn from virtual decks (Fig 3.2.A). Blocks consisted of sequences of 13 consecutive trials drawn from a given deck. On each trial, the participant categorized a presented digit as either odd/even (parity judgment) or greater/less than 5 (magnitude judgment). The color of a circle (green or blue), surrounding the digit, cued whether to perform the parity or magnitude judgment on each trial based on a pre-instructed mapping. The participant indicated their categorization by pressing the left or the right arrow key on the keyboard. The response deadline was 3 sec. Trials were separated by .2 sec. Response mappings and color-to-task mappings were counterbalanced across participants. Digits 1-4 and 6-9 were used with equal frequency across both tasks and were randomized for order of presentation.

In order to manipulate cognitive effort across decks, we varied the frequency of task switching required by a particular deck (Fig 3.2.A). The probability of switching tasks from trial to trial, within a sequence drawn from each deck, increased across the 2 effort levels: 23%, and 69%, respectively. Note that the lowest effort level still included two switches. A higher frequency of task switching is associated with a higher experience of cognitive effort (Monsell, 2003), and has been shown to be aversive (Arrington & Logan, 2004). At the beginning of each block, a shape was presented for 1 sec to indicate which deck was being performed, and this shape was also tiled in the background throughout the performance of the block. Participants were told that this shape represented the back of the virtual card deck from which trials for that sequence were drawn. Each effort level was associated with a particular deck. Participants were not told about this relationship, but could learn through experience that certain decks were more effortful than others to perform. The mappings between decks and effort levels were randomized across participants.

Practice Phase

During an initial “Practice Phase” participants gained familiarity with the trial structure, the categorization decisions, and the color and response mappings. After being given instructions regarding the categorization and response rules (Fig 3.2.B), they practiced two 13-trial blocks. Each task sequence included the 1 sec long presentation of a deck symbol, and the subsequent performance of 13 consecutive task-switching trials. In the first block of the Practice Phase, feedback was presented after the button press as either 'Correct' in pink, 'Incorrect' in yellow, or 'Please respond faster' if the participant failed to answer in 3 sec. In the second block of the Practice Phase, feedback was omitted, as would be the case in the actual experiment. The deck symbols that participants encountered in the Practice Phase were for practice only and were not presented during the Learning or Test Phases to avoid any association of error feedback with a particular deck symbol.

Learning Phase

In the Learning Phase (Fig. 3.2.C), participants learned the association between 4 decks and an effort level. Each deck was performed 15 times in random order. The Learning Phase was separated into six, approximately 5 minute-long runs.

Test Phase

In the Test Phase (Fig 3.2.D), two decks were presented and the participant chose which they would like to perform (Selection epoch). The participants were told to freely choose between decks prior to a 3 sec deadline. We note that in contrast to other DST variants (Gold et al., 2015), participants in this task were not told about the difficulty manipulation to avoid biasing choices based on participants’ explicit beliefs about effort. Once the participant made their selection, the unselected deck disappeared from the screen and the selected deck remained

on the screen for the duration of the 3 sec interval. In the event of a non-response, the same choice pair was presented again at the end of the entire experiment until they made a decision. Each pair was selected from the set of six unique (un-ordered) pair combinations of all four decks, excluding self-pairing (e.g., deck #1 paired with deck #1). Each deck was presented either on the left or the right side of the screen, counterbalanced for location across trials. The Selection epoch was followed by the execution of the selected effort level task sequence (Execution epoch). The sequence of events in this epoch was the same as during the Learning Phase. The Test Phase was separated into two, approximately 7 minute-long blocks. In total, each pair was presented 6 times in a randomly intermixed order, making a total of 36 decision trials across two blocks.

Inventory

Upon the completion of the Test Phase, all participants completed three surveys: (1) Need for Cognition (NfC), a subjective inventory that measures daily engagement in cognitively challenging activities; (2) Visual Analogue Fatigue Scale (VAFS), a subjective inventory that measures daily Parkinson's Disease fatigue symptoms; and (3) fatigue severity scale (FSS), a subjective inventory that measures overall fatigue in daily activities.

Behavioral Data Analysis

Trials with response times below 200 ms were excluded from further analysis (less than 1% of execution trials in both phases for all participant groups in Experiment 1 and Experiment 2). Execution trials on which participants missed the response deadline were also excluded from further analysis (less than 1% of execution trials in both phases for Experiment 1 participants, approximately 2% of execution trials in both phases for PD patients, and less than 1% for age-matched controls in Experiment 2). Response times were calculated using only correct trials.

Choice behavior was assessed by calculating the probability of selecting a deck across all selection trials on which that deck was presented as an option during the Test Phase.

‘Demand Avoidance’ was calculated by averaging the probability of selecting the easier task in pairs with difficulty difference.

‘Bias’ was calculated by subtracting the probability of selecting the first deck from the probability of the second deck in pairs without a difficulty difference. Note that, the probability of selection across easy and difficult decks adds up to 1, and thus, calculating the difference between any easy and difficult pair yields the same value.

In Experiment 1, data were analyzed using a mixed-design analysis of variance (ANOVA) (within subject factors: Deck type and Difficulty level). If the sphericity assumption was violated, Greenhouse-Geisser correction was used. Significant interactions were followed by simple effects analysis, the results of which are presented with False Detection Rate (FDR) correction. Alpha level of .05 was used for all analyses. Error bars in all figures stand for within-subject error.

‘Demand Avoidance General Linear Model (GLM)’

In order to test the hypothesis that dopaminergic medication intake during Learning and/or Test Phases has an effect on demand avoidance behavior, the ‘Demand Avoidance General Linear Model (GLM)’ was devised. In this GLM, Demand Avoidance for each participant was entered as the dependent variable. Each PD patient group and age-matched controls were assigned a unique group value. This group value of the participant was entered as the predictor. Three contrasts tested the predictions under the ‘cost of control’ hypothesis and two contrasts explored the interactions between Learning and Test Phases.

Relationship between task-performance and effort-selection

In order to test the effects of performance and task-selection on selection behavior, while accounting for individual variability within experimental groups, we adopted the following mixed effects hierarchical regression model:

$$y_i = X_i(\beta + s_{j[i]} + m_{[i]}) + \epsilon$$

$$s_j \sim N(0, \Sigma_S), \quad \epsilon \sim N(0, \sigma)$$

Here, y_i is the i th observed selection rate, X_i is a matrix of predictors (brain activity) and covariates (performance), β_i is a vector of coefficients (as in conventional linear regression), $j[i]$ is the subject of the i th observation, so that $s_{j[i]}$ is a subject-specific perturbation of all of the coefficients.

This fitting procedure was implemented using the R package `blme`, which includes the `lme4` package (Bates et al., 2014) and performs maximum-a-posteriori estimation of linear mixed-effects models.

Results

Experiment 1

Experiment 1 established the basic behavioral effects of our modified version of DST. Overall, undergraduate participants were highly accurate on the categorization task across both phases (mean error: 5% in the Learning Phase, 5% in the Test Phase), and participants missed the deadline on few trials (0.6 % of Learning phase trials, $SE=0.001$, 0.1% of Test phase trials, $SE=0.0003$).

Both RT and error rates increased with greater effort demands, increasing with a higher proportion of task-switches across both switch and repeat trials. There was a significant effect of Effort on error rates, ($F(1,10)=8.45$, $p=.02$, $\eta_p^2=.46$) and RT ($F(1,10)=43.74$, $p<.001$, $\eta_p^2=.81$). Meanwhile, deck type did not have an effect on task performance (errors: $F(1,10)=2.27$, $p=.17$, $\eta_p^2=.18$; RT: $F(1,10)=2.15$, $p=.17$, $\eta_p^2=.18$). RT and ER increases for the harder task were also significant for switch error rates and repeat RTs, tested separately (switch errors: $F(1,10)=5.69$, $p=.04$, $\eta_p^2=.36$; repeat errors: $F(1,27)=0.4$, $p=.84$, $\eta_p^2=.004$; switch RT: $F(1,27)=3.21$, $p=.10$, $\eta_p^2=.24$; repeat RT: $F(1,27)=43.01$, $p<.001$, $\eta_p^2=.81$). RTs, but not error rates improved across phases. There was a significant effect of phase on RT ($F(1,10)=91.10$, $p<.001$, $\eta_p^2=.90$), but not on error rates ($F(1,10)=3.06$, $p=.11$, $\eta_p^2=.23$).

During the Test phase, 10 out of 11 participants (90%) selected the easier task more than 50% of the time overall (Figure 3.5). The probability of selecting the easier task across participants was significantly different than chance, $M_\beta=0.67$., $SD_\beta=.16$, $t(10) = 3.6$, $p = .005$. Furthermore, there was a significant effect of effort level on choice behavior, ($F(1,10)=13.06$, $p=.005$, $\eta_p^2=.57$), but not of deck type, ($F(1,10)=0.24$, $p=.64$, $\eta_p^2=.02$), such that participants chose the easier task regardless of which deck color yielded the easier task. Participants did not

show a bias toward the easy decks in pairs without a difficulty difference ($M_{\beta}=0.09$., $SD_{\beta}=.44$, $t(10) = 0.67$, $p = .52$), indicating that participants only chose the easier task when there was a difficulty difference between options. Decision time to choose a task was the same across decks and effort levels (deck: $F(1,10)=0.60$, $p=.81$, $\eta_p^2=.01$; effort: $F(1,10)=2.99$, $p=.11$, $\eta_p^2=.23$).

We tested the effects of Learning Phase and Test Phase task performance and task-switching probability in predicting effort selections. Consistent with previous reports (Demand Avoiders in Sayali & Badre, 2018), undergraduate participants avoided tasks that yielded greater task-switching probability ($M_{\beta}=-0.49$., $SD_{\beta}=0.14$, $t(44) = 4.92$, $p = .001$). There was no effect of Learning or Test Phase RTs or error rates.

Overall, Experiment 1 showed that our modified version of DST was successful in manipulating cognitive control recruitment, as indexed by RTs and error rates, and establishing demand avoidance behavior in an undergraduate sample as a function of task-switching probability. Furthermore, Experiment 1 allowed us to detect the observed effect size of effort level on choice behavior ($\eta_p^2=.57$). A power analysis revealed that on the basis of the mean, and the within-subject comparison effect size observed in Experiment 1, an N of approximately 17 in each experimental group (68 total patients and 17 age-matched controls) would be required to obtain statistical power at the recommended .80 level (Cohen, 1988).

Experiment 2

Experiment 2 data collection is still ongoing, however, the data for a total of 34 PD participants and 13 age-matched control participants have been analyzed and reported below.

Experiment 2 tested the same experimental paradigm on PD patients and age-matched controls. PD patients were separated into 4 experimental groups based on their medication intake. ‘OnOn’ group participants were on their medication during both phases of the

experiment. ‘OnOff’ group participants were on their medication during the Learning Phase, and off their medication during the Test Phase. ‘OffOn’ group participants were off their medication during the Learning Phase, and on their medication during the Test Phase. ‘OffOff’ group participants were off their medications during both phases of the experiment.

We tested demographic, personality trait and fatigue severity differences between groups and control participants. PD groups and control participants did not differ in terms of their age ($M_{\beta}=69.67$, $SD_{\beta}=7.23$, $F(4,44)=0.55$, $p=.69$) or level of education they received ($M_{\beta}=15.28$, $SD_{\beta}=5.82$, $F(4,41)=2.39$, $p=.07$). However, PD patients consisted mostly of males, as opposed to control participants who mostly consisted of females (OnOn: $N_{\text{female}}=1$, $N_{\text{male}}=6$; OnOff: $N_{\text{female}}=2$, $N_{\text{male}}=6$; OffOn: $N_{\text{female}}=3$, $N_{\text{male}}=4$; OffOff: $N_{\text{female}}=3$, $N_{\text{male}}=7$; Control: $N_{\text{female}}=11$, $N_{\text{male}}=2$; gender: $X^2(4, N = 45) = 13.15$, $p=.01$). As PD affects more males than females, our control participants mainly consisted of PD patients’ partners.

There were no differences in NfC scores among groups ($F(4,39)=1.63$, $p=.19$), indicating that PD patients, across experimental groups, as well as control participants, reported similar levels of daily engagement in effortful activities. As expected, Visual Analogue Fatigue Scale (VAFS) scores differed across PD patients and control participants ($F(4,38)=2.77$, $p=.04$), with control participants scoring higher on this scale compared to all PD groups except OffOn group. As a low score on this measure means worse fatigue severity, PD patients displayed higher global fatigue levels compared to control participants. Consistently, Fatigue Severity Scale (FSS) also differed across participants ($F(4,39)=1.63$, $p=.001$), with control participants reporting lower fatigue levels in the past week compared to all PD groups except OffOn group.

Overall, we showed that PD patient groups and control participants did not differ in terms of age or education. As PD affects more males than females, PD participants consisted of mostly

males, and age-matched controls consisted mostly of females. Participants were matched in terms of their self-reported tendency to avoid mental effort. As fatigue is a common symptom of PD, PD patients reported significantly higher levels of fatigue compared to age-matched control participants.

Task performance for all patients and participants improved across phases and worsened with effort requirements of the task. There was an effect of experimental phase on RTs (Figure 3.3) but not on error rates, (RTs: $F(1,42)=30.84$, $p<.001$, $\eta_p^2=.42$; error rates: $F(1,42)=3.26$, $p=.08$, $\eta_p^2=.07$). There was an effect of task difficulty on both RTs and error rates, (RTs: $F(1,42)=206.24$, $p<.001$, $\eta_p^2=.83$; error rates: $F(1,42)=42.53$, $p<.001$, $\eta_p^2=.50$). RTs, but not error rates, changed across participant groups (RTs: $F(4,42)=3.56$, $p=.01$, $\eta_p^2=.25$; error rates: $F(4,42)=1.59$, $p=.19$, $\eta_p^2=.13$). OffOff group participants overall had significantly slower response times compared to OnOn group and Control participants, indicating that being off- PD medication in either phase resulted in overall slower response times.




Figure 3.3 Response times during task execution in Learning and Test Phases for each experimental group.

We have tested three hypotheses by devising five contrasts in our Demand Avoidance GLM. The first contrast compared the effect of being off-medication during both phases to age-matched controls and being on-medication on both days, with the prediction that demand avoidance would scale as: $PD_{OffOff} > Control > PD_{OnOn}$. This contrast was statistically insignificant, $F(1,42)=0.62$, $p=.43$, $\eta_p^2=.02$. The second contrast compared the main effect of medication status during the Learning Phase. This contrast was also statistically insignificant, $F(1,42)=1.09$, $p=.30$, $\eta_p^2=.03$. The third contrast compared the main effect of medication status during both the

Learning and the Test Phases. This contrast was also statistically insignificant, $F(1,42)=0.79$, $p=.38$, $\eta_p^2=.02$.

Additionally, we aimed to explore interaction effects of dopaminergic state between experimental phases. Thus, the fourth contrast tested for an interaction between Learning and Test Phase by comparing OnOff group to OffOff. This contrast was marginal, $F(1,42)=3.87$, $p=.05$, $\eta_p^2=.08$. The fifth contrast also tested for an interaction between Learning and Test Phase, this time by comparing OnOn group to OnOff. This contrast was not significant, $F(1,42)=0.08$, $p=.78$, $\eta_p^2=.002$.

The results show that participants off their medication did not show greater demand avoidance compared to age-matched controls and participants on their medication. Additionally, there were no main effects of the Learning Phase, contrary to the predictions of ‘dopamine cost hypothesis’ of effort. However, there was an interaction between the medication status intake across Learning and Test Phases. Specifically, participants who were off their medication during Test, but on their medication during Learning selected the easier task more often than participants who were off their medication during both phases of the experiment (Figure 3.4).




Figure 3.4 Plots demand avoidance rates in PD patients as a function of dopaminergic intake during Learning and Test Phase. The interaction between Learning and Test Phases was marginal.

Overall, during the Test phase, 38% of PD patients and 54% of age-matched controls selected the easier task more than 50% of the time (Figure 3.5.). None of the participant groups, except for OffOff group participants, had easier task selection rates that reliably differed from chance ($M_{\text{OnOn}}=0.45.$, $SD_{\text{OnOn}}=.13$, $t(6) = -1.0$, $p = .36$; $M_{\text{OnOff}}=0.55.$, $SD_{\text{OnOff}}=.21$, $t(9) = 0.76$, $p = .47$; $M_{\text{OffOn}}=0.48.$, $SD_{\text{OffOn}}=.08$, $t(6) = -0.76$, $p = .48$; $M_{\text{Control}}=0.51.$, $SD_{\text{Control}}=.18$, $t(12) = .11$, $p = .92$). OffOff group participants significantly selected the harder task more than 50% of the time ($M_{\text{OffOff}}=0.41.$, $SD_{\text{OffOff}}=.12$, $t(10) = -2.37$, $p = .04$). There was no significant effect of effort level ($F(1,42)=0.77$, $p=.38$, $\eta_p^2 = .02$), or of deck type, on choice behavior, ($F(1,42)=0.16$, $p=.70$, $\eta_p^2 = .004$). There was no effect of group, difficulty or deck type on decision times. No participants (all $ps>.05$), except the OnOff group participants, showed a bias toward the easy deck in pairs

without a difficulty difference ($M_{\beta}=-0.18.$, $SD_{\beta}=.24$, $t(9) = -2.40$, $p = .04$), indicating that most participants did not prefer a task in pairs with or without a difficulty difference.

RT analysis showed that participants in the Off-Off group showed the slowest overall RTs in the Learning Phase as well as showing demand seeking behavior in the Selection Phase. It is possible that participants with slower RTs have greater room to benefit from RT improvements, which might boost the value of the difficult task. In order to test this assumption, we median-split (Median=1.06 sec) the dataset for participants with slow vs. fast RTs across all patient groups and calculated the slope of RT change across the Learning Phase trials of both difficulty levels for all participant groups. Repeated measures ANOVA results for difficulty*group showed that there was no main effect of difficulty ($F(1,45)=2.24$, $p=.14$) but there was a main effect of participant group ($F(1,45)=4.59$, $p=.04$). The average RT slope for participants with slower RTs was -0.01 ($SE=0.001$), and the average RT slope for participants with faster RTs was -0.013 ($SE=0.001$). The results indicate that participants with faster RTs showed greater RT improvements in the Learning Phase, arguing against the possibility that participants with slower RTs had greater room for task improvement.




Figure 3.5 Graphs Demand Avoidance rates in modified DST for all PD patients, age-matched controls (Experiment 2), and undergraduate students (Experiment 1).

Finally, we tested the effects of Learning Phase and Test Phase task performance and task-switching probability in predicting effort selections for each participant group separately. No groups, except the OffOff group participants, selected tasks as a function of task-switching probability. OffOff group participants selected tasks that yielded greater task-switching probability ($M_{\beta}=0.25$, $SD_{\beta}=0.08$, $t(40) = 3.19$, $p = .003$). There was no effect of Learning or Test Phase RTs or error rates on task selections.

Discussion

The present research aimed to test the causal role of dopaminergic medication intake during effort learning and effort selection in PD patients. The ‘cost of control’ hypothesis predicts that cognitive control registers as disutility, which is tracked by the mesolimbic dopamine pathway activity in the brain. Dopaminergic medication in PD patients overstimulates the mesolimbic dopamine system, and leads to established alterations in reward-based decision-making. Thus, we predicted that dopamine intake in on-medication PD patients would reduce the cost of effort, and lead to reduced demand avoidance behavior compared to participants who are dopamine-depleted. Preliminary results did not find evidence for this hypothesis; rather, we found that dopamine-depleted patients displayed demand-seeking behavior. Furthermore, it was shown that effort is registered as a cost during reinforcement learning, and this cost can be amplified by dopaminergic imbalance in healthy adults (Cavanagh et al., 2014) and PD patients (Cavanagh et al., 2017). However, all studies to date that tested for the effect of dopamine on cognitive demand avoidance did so in the presence of external reward and without separating the learning and test phases. We aimed to test the role of dopamine in learning vs. the decisions regarding effort, in PD patients, by administering two epochs of effort-based decision-making, effort learning and effort decisions, on two consecutive days, while alternating dopaminergic medication intake with a full-factorial design between PD subjects, in the absence of external reward. Preliminary results did not show an effect of dopaminergic intake during learning, rather, we observed that the influence of dopaminergic state during effort learning on demand avoidance

depended on dopaminergic state during effort-decisions. We consider each of these findings in turn.

To test the hypothesis that effort costs are modulated by mesolimbic dopamine and ask whether dopamine intake affects the learning versus selection regarding cognitive effort, we devised a modified version of the DST paradigm. In this modified DST, we separated the Learning and Test Phases into two consecutive days. During the Learning Phase, participants learned the association between virtual deck symbols and subsequent cognitive control demands as manipulated by task-switching probability. During the Test Phase, participants chose between pairs of virtual decks and performed the effort level associated with the selected deck. Experiment 1 showed that in a sample of healthy, undergraduate students, this modified version of DST yielded demand avoidance behavior at a rate consistent with the previous literature (Kool et al., 2010), indicating that the modified version of DST could be reliably used for testing cognitive effort avoidance.

Next, using this modified version of DST, we alternated dopaminergic medication intake across Learning and Test Phases between PD patients. ‘Cost of control’ hypothesis suggests that mesolimbic dopamine plays a motivating role in alleviating effort costs. Consistent with this hypothesis, previous research (Cavanagh et al., 2014) showed that response conflict reduces the reward value or boosts the cost of cognitive tasks depending on the striatal dopamine level of the participant. In this study, during a training phase, participants learned the degree of response conflict and reinforcement each task yielded, and selected the most rewarding task in task pairs during a test phase. Healthy, young participants were administered a D2 agonist (cabergoline) or placebo during two sessions separated by one week. Cabergoline administration increased aversion to punishment, as manipulated by the degree which response conflict discounted the

reward value of the task. This effect was enhanced in participants who showed the greatest dopaminergic diminishment due to manipulation. Moreover, Cavanagh et al., (2017) confirmed the cost of diminished dopamine in PD patients by showing that unmedicated PD patients displayed greater cost of conflict, an effect which was remedied by medication intake. Since these studies administered medication on the same day that both effort learning and effort selection took place, an equivalent experimental condition in our study would be the patient group who were off-their dopamine medication on both Learning and Test Phases. Thus, we predicted that PD patients who are off their dopamine medication on both days would display the greatest demand avoidance behavior. However, in our study, we observed that participants, whose dopamine levels were diminished on both the Learning and Test Phases, exhibited the greatest demand seeking behavior. Thus, we did not provide support for the ‘cost of effort’ hypothesis and did not replicate the ‘cost of conflict’ results in a dopamine-depleted patient sample.

The discrepancy between our findings and the previous research might be due to differences in the operationalization of cognitive effort, presence of extrinsic reinforcements and/or the timing of measures. First, in our task, we manipulated cognitive effort demands by manipulating task-switching probability, while previous research manipulated response conflict in a modified Simon Task. It is not yet known whether the cost of Simon Task and task-switching probability are registered similarly.

Second, the current research tested for effort selection in the absence of explicit reinforcers, such as monetary reward or feedback. Thus, the only potential cost/value of the task in our study was intrinsic. It is possible that the motivation for effort selections are differentially affected by the presence and absence of reward. Previous research (Schouppe et al., 2014a,

Botvinick et al., 2013) showed that the reward network of the brain is recruited differentially depending on the presence of external reward.

Third, prior research manipulated dopamine levels during the same session that effort learning and selection took place, while the current research alternated dopamine intake across two consecutive days. It is possible that the effects of dopaminergic state during effort learning interact with the effects of dopaminergic state during effort selection. Consistent with this interpretation, preliminary results indicated that dopaminergic state, neither during effort learning nor during effort selection, had a main effect on effort selections. Instead, PD patients who had diminished dopamine levels during effort selection showed greater effort seeking behavior, if they also learned effort tasks in a dopamine-diminished state. On the other hand, patients who had diminished dopamine levels during effort selection showed greater demand avoidance behavior if they learned effort tasks in a dopamine-enhanced state.

Alternatively, ‘opportunity cost of dopamine’ hypothesis (Cools, 2015) of effort argues that attending to the current task comes at the cost of missing out on rewarding alternative tasks. In this framework, the reward value of the alternative task is a function of the average reward rate of the task context, which has been observed to be tracked by tonic dopaminergic levels and increased response vigor (Niv et al., 2007). Thus, the ‘opportunity cost of dopamine’ hypothesis predicts that individuals with relatively greater dopamine levels will experience a greater opportunity cost of time, due to greater expected reward rate of alternative tasks, and exhibit greater demand avoidance. A direct test of this account would predict that dopamine-enhanced patients will exhibit the greatest demand avoidance behavior and the fastest response times. For example, a recent report (Frobose et al., 2017) illustrated that prolonging catecholamine

transmission by cabergoline administration increased demand avoidance, although did not affect task performance.

In the current study, consistent with the response vigor predictions, we observed patients off their medication on both days to have the slowest response times, and the highest demand-seeking behavior. Additionally, patients on their medication on both days showed the fastest response times; however, inconsistent with the opportunity cost account, they did not show greater demand avoidance behavior. Furthermore, the change in response times did not account for changes in effort selection rates in any patient group. Instead, patients off their medication reliably selected those tasks that yielded greater task-switching probability.

Previous research indicated that PD patients who are off their medication learn to avoid negative outcomes better than positive outcomes, while PD patients who are on their medication learn better from positive outcomes than negative outcomes (Frank et al., 2004). Thus, it is possible that the ‘the cost of dopamine’, during the learning of effort costs, is evaluated differentially according to the dopaminergic state of the participant, during effort selections. This extended account of the ‘opportunity cost of dopamine’ hypothesis would predict that PD patients who experience greater opportunity cost of time, due to enhanced dopamine levels during Learning, would be more sensitive to this cost during Test if they were in a dopamine-diminished state during Test. On the other hand, participants who were in a dopamine-diminished state during Learning might be exempt from this cost, and thus display less effort avoidant behavior. However, this extended account of ‘opportunity cost of dopamine’ fails to explain the effort seeking behavior in PD patients who were dopamine-depleted in both Learning and Test phases, since according to this account, dopamine-depleted patients during Test would

be more sensitive to learned costs, rather than rewards during decision-making. These unexpected results might be due to two reasons.

First, despite the fact that the modified version of DST yielded demand avoidance behavior in a young, undergraduate sample, preliminary results showed that elderly controls in Experiment 2 did not show demand avoidance, rather they were insensitive to the effort manipulation. Thus, it is possible that, regardless of dopamine status, age-matched PD patients were originally insensitive to effort costs. Given this baseline, reduced sensitivity to the cost of effort might appear as greater effort seeking behavior.

Second, it is possible that in the absence of external reward, voluntary effortful task performance might register as intrinsically rewarding (Schouppe et al., 2014a) or effortless/boring tasks could be intrinsically costly. Cognitive and computational approaches (Oudeyer & Kaplan, 2007) to intrinsic motivation argue that challenging tasks are intrinsically rewarding due to yielding greater interestingness or reduced boredom. Additionally, a recent report showed that tasks that induce boredom increase reward sensitivity (Milyavskaya et al., 2018), indicating that boredom itself might be aversive or yield an intrinsic cost. PD patients who were exempt from the ‘opportunity cost of dopamine’ during effort learning, due to diminished dopamine levels, could have relied on this ‘boredom’ signal associated with the easier task when making selections. In other words, engaging in boring tasks might be intrinsically aversive for PD patients who are in a dopamine-depleted state.

Alternatively, engaging in cognitively effortful tasks might be intrinsically beneficial for PD patients. While some cognitively challenging situations have been shown to worsen PD symptoms via stress (Giladi & Hausdorff, 2006), video games that target cognitive and motor skill training via trial-by-trial performance gains have been shown to elicit greater remediation of

motor symptoms than physical exercise therapy alone (Pompeu et al., 2012). The beneficial effects of dance therapy on PD symptoms are argued to be via the cognitive engagement required during coordinated motor movements (Foster et al., 2013). Thus, the ‘dopaminergic-gain hypothesis of cognitive effort’ argues that the inclusion of a cognitive component in motor training might improve brain health via greater neurogenesis and increased synaptic strength within the cortico-striatal loop by reversing dendritic spine loss in striatal dopamine D2 neurons (Petzinger et al., 2013). Consistent with the dopaminergic-gain hypothesis of effort, studies on ‘flow’ states observed that striatal BOLD activity during the optimally challenging tasks was greater than that of the easy and difficult tasks (Ulrich et al., 2014; Ulrich et al., 2015), leading to an inverted U-shaped relationship between striatal BOLD and effort. Moreover, people who were more prone to ‘flow’ induction had greater dopamine D2 receptor availability (de Manzano et al., 2013), the pathway of which, incidentally, is more active in dopamine-depleted individuals. Future research should observe the relation between dopaminergic activity and challenging task performance in PD patients.

The current research has three major limitations. First, based on the effect size observed in Experiment 1, we aimed to recruit 68 PD patients and 17 age-matched control participants for Experiment 2. Due to challenges in patient recruitment, drop-outs, and removals due to confounding mood disorders and cognitive impairment in PD patients, we were only able to analyze data from approximately half of the target sample size. Thus, we did not have the necessary statistical power for reliable hypothesis testing.

Second, preliminary Experiment 2 results showed that almost half of the elderly control sample selected the harder task, indicating that our modified version of DST was not able to replicate greater effort avoidance rates in elderly, as was observed in the literature (Westbrook et

al., 2013). The discrepancy in our findings might be due to differences in effort tasks between studies. Participants in Westbrook et al. (2013) were aware of the effort manipulation, as task instructions differed across effort levels. In the current task, task instructions were the same across effort levels, and participants were not aware of the effort manipulation. There are at least two studies (Gold et al., 2015; Dunn & Gaspar, 2017) reporting that awareness of the difficulty difference between task options is necessary for displaying effort sensitivity in DST. However, it is not clear why this awareness should specifically affect the elderly, and not undergraduate students, nor PD patients who were off their dopaminergic medication on both days. Future studies should directly test the role of effort-awareness in DST in elderly population.

Third, previous research demonstrated that the recency of PD diagnosis mediates the relation between dopaminergic medication intake and experienced cost of conflict (Cavanagh et al., 2017). Our study did not control for diagnosis recency. Future work will retrospectively obtain these measures, in order to utilize it as a covariate of demand avoidance behavior.

Furthermore, we observed a similar demand seeker/demand avoider ratio in our current study with older adults as in our previous study on fMRI volunteers (Sayali & Badre, 2018). Our previous study observed that fMRI volunteers scored significantly higher on NfC compared to subject pool participants. Similarly, most PD patients and age-matched control participants' NfC scores were in this range. Thus, it is possible that, similar to fMRI volunteers, older adults who volunteer for study participation might have a greater need for cognition. However, it is important to note that there were no group differences in NfC scores across participants. Thus, the difference in demand avoidance rates between groups cannot be attributed to NfC alone. Future studies should attend to overall effort seeking tendencies in elderly when testing for effort avoidance.

Overall, the preliminary results did not provide support for the ‘cost of control’ hypothesis. On the contrary, we observed that dopamine-depleted participants exhibited the greatest demand seeking behavior. Although preliminary, the results provide partial support for the ‘opportunity cost’ hypothesis of dopamine, which predicts enhanced dopamine to increase the cost of effort. Accordingly, learning effort costs in a dopamine-enhanced state could have increased the cost of effort, while making selections based on this cost in a dopamine-depleted state could have increased participants’ sensitivity to this cost. Additionally, greater demand seeking behavior in dopamine-depleted PD patients might provide support for the ‘dopaminergic gain hypothesis’ of effort. Future work will aim at completing the dataset.

Chapter 5: General Discussion

Across all studies, we aimed to test the neural mechanisms underlying cognitive effort avoidance. Although the aversion towards cognitive operations, such as thinking and problem solving, has been colloquially acknowledged for decades, the first empirical demonstration of cognitive effort avoidance is fairly recent (DST; Kool et al., 2010). Since then, multiple attempts have been made to explain mechanism underlying cognitive effort avoidance. Accordingly, it has been suggested that error avoidance (Dunn et al., 2017), time-on-task, opportunity cost (Kurzban et al., 2013), response conflict (Botvinick et al., 2007), metabolic limitations of the brain (Gailliot et al., 2007), or limited multitasking capabilities of the neural system (Musslick et al., 2015) could underlie the aversion associated with cognitive effort.

Regardless of the source of the effort cost, it has been widely accepted that cognitive effort registers as a cost. The theory that has received the most empirical testing is the ‘cost of control’ hypothesis of cognitive effort (Shenhav et al., 2013; Botvinick et al., 2007), arguing that there is an intrinsic cost associated with recruiting cognitive control. The support for this theory indicated that effort avoidance positively correlates with frontoparietal control recruitment (McGuire & Botvinick, 2010), and negatively correlates with ventral striatal activity (Botvinick et al., 2009). Furthermore, it was shown that dopamine administration might reduce the cost associated with effort (Cavanagh et al., 2014).

However, there has also been incongruent evidence demonstrating that ventral striatum increased its activity (Schouppe et al., 2014a), as opposed to the expected decrease, when selecting the harder task, and that greater control recruitment did not predict effort avoidance (Schouppe et al., 2014b). Likewise, it was exhibited that a dopamine reuptake blocker increased, rather than decreased, effort avoidance (Frobose et al., 2017).

An evaluation of the cognitive effort avoidance literature hitherto determined that the discrepancy between the results of these experiments might be due to: 1) differences in task characteristics, 2) differences in task instructions, and 3) differences in the timing of measures. In order to address these issues in the literature, we developed two new experimental paradigms that adopted and modified the DST, and tested the involvement of neural mechanisms in distinct experimental epochs of effort-based decision-making.

Task-characteristics

Presence of external rewards

First, across all studies, we avoided the inclusion of performance feedback, as well as explicit rewards, following effort execution. Prior research had inconsistently tested for effort avoidance in the presence or the absence of reward (Botvinick et al., 2009; Schouppe et al., 2014a; Cavanagh et al., 2014; Frobose et al., 2017). However, it was argued that voluntary selection of a harder task might yield an intrinsic value in the absence of external reward (Schouppe et al., 2014a). Incidentally, using our version of DST in the absence of external rewards, we reported the highest rate of effort seeking behavior in the literature (Chapter 2). In that study, we also demonstrated that effort-seeking participants selected the tasks that yielded the greatest task-switching frequency, as well as an increase frontoparietal recruitment, both of which are indices of cognitive control. Therefore, we argue that in the absence of performance feedback and external reward, the same control indices that register as a cost, might be valued. Further, under the cost hypothesis tested in Chapter 4, Parkinson's Disease patients off of their medication were expected to display the greatest demand avoidance behavior, due to the posited role of dopamine in relieving effort costs under the 'cost of control' hypothesis. However, preliminary results of our study (Chapter 4) indicated that Parkinson's Disease patients who were

off of their medication exhibited the greatest demand seeking behavior, in the absence of reward. The results indicate that the presence of external reward(s) might interact with the intrinsic evaluation of cognitive effort. Future studies should test cognitive effort avoidance in the presence and absence of external rewards.

Estimation of neural activity

Secondly, our first study design aimed at developing a parametric version of DST. The fMRI literature on cognitive effort avoidance thus far has adopted a subtraction approach (Schouppe et al., 2014; Botvinick et al., 2009). These studies inconsistently demonstrated that reward network both increases and decreases its activity, in relation to effortful task execution. We aimed at characterizing the neural function underlying effort selection, by implementing intermediate effort options and correcting for the neural activity change between extreme low- and high-effort conditions. By using a parametric approach, we were still able to find dose-dependent changes in frontoparietal cortex, as well as ventral striatum, in relation to cognitive effort execution, yet unable to replicate the involvement of these regions in predicting effort avoidance, as reported in prior research that had adopted a subtraction logic. The results indicate that a subtraction logic might be insufficient in estimating the relationship between brain and behavior in cognitive effort tasks.

Type of cognitive task

Thus, in our parametric version of DST, we manipulated cognitive effort across six levels by increasing the frequency of task-switching during task-execution. We adopted task-switching as our measure of cognitive control requirement in all experiments, in order to control for task-characteristics across all studies, as well as to match the first cognitive effort avoidance study reported in the literature (DST; Kool et al., 2010).

Across all studies, we aimed to test hypothesis that cost of control is what underlies effort avoidance. In order to test this hypothesis, we parametrically manipulated the degree to which a task required cognitive control. One way to operationalize cognitive control demands of a task is via task-switching. It was shown that task-switching yields greater response times, error rates and greater BOLD activity in brain areas that are associated with cognitive control, compared to task-repeats. Therefore, we manipulated the probability of task-switching across task levels. We predicted that if cognitive control is what underlies effort avoidance, those tasks that yield greater probability of task-switching, thus greater cognitive control demands, would be avoided more compared to those tasks that yield greater task-repeats.

However, the trade-off between stability and flexibility influences the allocation of control to the current task. For example, in a task-switching task, the cost of switching to the other task increases with greater allocation of control to the initial task. Thus, in tasks with greater task-switching requirements, control allocation can be optimized by favoring flexibility and allocating control to the irrelevant task (Musslick et al., 2018). This type of control allocation is evident in reduced switch-costs and greater error-rates in tasks with greater task-switching frequencies, suggesting that participants adopt a flexible control strategy with greater demands on task-switching. In other words, participants reduce the amount of control allocated to the current task when the task require greater switching. This observation is incongruent with the assumption that tasks with greater task-switching frequency should automatically lead to greater demands on cognitive control. If we assume that switch-costs and not manipulated task-switching probabilities represent the amount of control allocated to the relevant task, we would expect that those participants who incurred the greatest switch costs to show the greatest demand avoidance rates. In order to test this prediction, we calculated switch costs across subjects and

compared average switch costs across task levels between demand avoiding and demand seeking participants. We found that there were no significant group differences ($F(1,48)=0.16, p=.69$) in switch costs across levels, indicating that demand avoiding participants did not exhibit greater control allocation compared to demand seeking participants, as defined here by switch costs.

These results are congruent with the observation that demand avoiding and demand seeking participants showed the opposite effort selection pattern across objectively manipulated task levels. The results overall indicate that neither objective control demands of the task as manipulated by task-switching probability, nor individual control allocation as defined by switch-costs underlie effort avoidance in our task.

Task-instructions

As task instructions were kept the same across all effort levels, participants were not aware of a difficulty manipulation. We reported these results in Chapter 2. This observation stands in contrast with the task-instructions adopted in various effort-discounting paradigms (Westbrook et al., 2013), as well as certain versions of the DST paradigm (Gold et al., 2015). Recent reports indicated that awareness and general intelligence might mediate effort avoidance in DST (Gold et al., 2015); moreover they argued against an implicit learning mechanism underlying effort avoidance. Further, self-report measures of cognitive effort engagement, as measured by NfC, have been demonstrated to correlate with effort-discounting scores (Westbrook et al., 2013), while we exhibited that NfC can distinguish demand avoiders from demand seekers, but fails to explain between-subject variability in demand avoidance rates. Thus, our parametric version of DST relied on implicit learning mechanisms instead of meta-cognitive beliefs and mechanisms that might influence effort decisions.

Our post-experimental questionnaires showed across three studies that participants did not report awareness of an effort-related difference between task levels, although choice behavior was consistent with avoiding (or seeking) decks with more task-switching. This discrepancy between task selections and effort-awareness suggests that the mental experience of effort might not be required to drive effort avoidance. This observation stands in contrast with theories (i.e. opportunity cost of effort) that suggest that subjective experience of effort corresponds to the calculation of the costs and benefits of effort. According to these theories, this subjective experience is necessary to reallocate mental resources away from effortful tasks that are not deemed valuable. The averseness of a task is argued to rise to consciousness due to decrements in task-performance during effortful tasks. However, it is important to note that the phenomenology of effort and task-performance can be dissociated. Indeed, this is a long-standing question in psychology. Some of the earliest studies of mental fatigue (Arai, 1912) reported only a weak relationship between the experience of mental fatigue and actual performance declines. Furthermore, it was shown that an ACC lesion patient was able to exert the required cognitive control necessary to complete a Stroop task, and displayed poorer task-performance when the task got more difficult, but the patient did not report any changes in the subjective feelings of effort (Naccache, 2005). This example illustrates that deployment of cognitive control and reduced performance is not always coupled with effort awareness. This observation is consistent with the view that most cognitive operations are unconscious. On the other hand, demand avoidance might be affected by subjective evaluations of one's self-efficacy instead of decisions regarding effort allocation. For example, it was shown that decrements in math performance were related to math anxiety and not reports of effort allocation (Bandalos et al., 1995). It is possible that the subjective feelings of effort, as defined by performance decrements, are tied to

perceived competence at task and rises to consciousness as an epiphenomenon. Additionally, individuals' task goals might determine their attribution styles (Dweck & Leggett, 1988). It was shown that people with learning goals tend to attribute their failures to internal factors such as effort allocation, while people with performance goals tend to attribute their failures to external factors such as luck.

Thus, individual differences in task goals, self-perceived efficacy might also determine on which basis individuals avoided demanding tasks. For example, COG-ED task adopted a working memory paradigm that explicitly associated each task with its corresponding difficulty level. This paradigm stands in contrast to DST paradigms that adopt task-switching designs, where the association between task difficulty and context is learned implicitly. Explicit effort avoidance designs run the risk of confounding averseness associated with control intensity with metacognitive beliefs about performing difficult tasks, i.e. self-perceived efficacy. It was shown that subjective feelings of efficacy can be decoupled from effort exertion in predicting performance decrements, indicating that beliefs about one's competencies can influence task performance. Thus, in this paradigm, participants might rely on their self-perceived efficacy in performing difficult tasks in addition to learned estimates of effort requirements. For example, it was shown that self-efficacy for mental multiplication was correlated to problem solving efficiency (Hoffman, 2010), indicating that beliefs about competency might influence task performance. Similarly, beliefs about one's capacity to perform under difficult contexts might influence their likelihood to engage with such tasks, independent of one's actual performance on the task. On the other hand, in effort avoidance designs that adopt implicit learning mechanisms, we expect that effort avoidance will be driven by individual differences in the speed of learning, rather than the overall control intensity the task demands.

Timing of measures

In our design, we separated the effort-learning and effort-choice components of decision-making into separate experimental phases. A recent study, which used the original DST that includes effort learning and effort decisions in the same epoch, showed that response stickiness might constitute a problem when sampling from different effort options (Frobose et al., 2017). Specifically, diseased populations, such as Parkinson's Disease, might not adequately select from a task, and thus might fail to learn the cost/value associated with that option. In order to address this problem, while developing a design that allows the observation of separate components of effort-related processes, we separated effort learning from effort decisions. Accordingly, participants first learned the cost/value associated with each effort task by experiencing each task without making a choice. During this Learning phase, a symbol cued the execution of each task. Following the presentation of the cue participants executed the associated effort task. During the Test phase, participants chose between two symbols that are associated with two effort tasks, and executed the chosen effort task. The inclusion of a separate Learning phase allowed us to ensure that all participants sampled from each effort level with equal frequency, and that they acquired a preference for an effort task prior to making effort selection. This presupposition was confirmed by our observation that participants' effort avoidance rates were stable across time during the Test phase. We replicated this observation across two studies, with the findings reported in Chapter 2 and Chapter 3. Importantly, this observation stands in contrast with the findings of the original DST paradigm (Kool et al., 2010), which demonstrated that demand avoidance rates increased across the course of the experiment, as the original DST did not separate effort learning from effort decisions. Separating the Learning and Test phases also enabled us to alternate dopaminergic medication intake during effort learning versus effort decisions in

Parkinson's Disease patients. The preliminary results (Chapter 4) indicated that the effect of dopamine intake during effort decisions might change as a function of dopaminergic status during effort learning. This suggests that the same neural mechanisms might impact effort learning and effort decisions, differentially.

For our fMRI designs, we further separated the Learning and Test phases into its sub-components. We separately optimized for effort anticipation (Anticipation epoch) and effort execution (Execution epoch) during the Learning phase, and we separately optimized for effort selection (Selection epoch) and effort execution (Execution epoch) during the Test phase. Prior research did not distinguish between these epochs when testing for neural mechanisms underlying effort avoidance. While some studies observed brain activity following effort execution (Botvinick et al., 2009), others observed during effort anticipation (Schouppe et al., 2014a). Incidentally, the results of these studies were incongruent with regards to the involvement of aforementioned brain areas. In both fMRI studies (Chapter 2 & Chapter 3), we demonstrate that effort execution and effort anticipation, as well as effort execution and effort selection do not rely on the same neural mechanisms. Furthermore, we exhibited that the involvement of the same brain region varies across epochs of effort anticipation and effort selection, indicating that the timing of brain measures, during different epochs of effort-based decision-making, might differentially impact the recruitment and the involvement of brain regions. Separating the Learning phase from the Test phase further enabled us to observe effort learning in the absence of decision-making. A previous study observed brain activity in a task where effort learning and effort decisions simultaneously took place (Nagase et al., 2018). We demonstrated that in the absence of decision-making, the cost of effort is not tracked during the Anticipation epoch. Instead, control- and conflict-monitoring-related brain regions are recruited

only when task performance is worse than expected, implying that the brain might represent cost/benefit calculations only when a voluntary decision is due, and not when effort is imposed.

Thus, across all studies, we showed that task characteristics, task instructions, and timing of measures affect effort-based decision-making and the employment of neural mechanisms during effort tasks. As such, discrepancies in the existing literature can be reconciled by controlling for these parameters across studies. Next, we interpret the findings of our studies in relation to the existing accounts of cognitive effort avoidance: 1) Processing cost, and 2) Resource cost.

Processing cost

The most prominent account of cognitive effort avoidance argues that the cost of effort is the recruited control mechanisms during effortful tasks (Shenhav et al., 2013). The cost associated with controlled task performance reduces the expected value of the task, and determines the likelihood of exerting effort. Under this hypothesis, we expected to see greater frontoparietal network and diminished reward network recruitment during effort execution, and the degree of their change to predict greater demand avoidance behavior. It's important to note that the cost-of-control theory only assumes that effort costs scale with control intensity without making any predictions about the neural mechanisms underlying where and how these costs are registered. Given the association between cognitive control and FPN, in this work, we extrapolate the cost-of-control theory and predict that effort costs will scale with FPN recruitment in order to test the hypothesis that the cost of effort is cognitive control. Our prediction is consistent with Kurzban (2010)'s proposal that lateral prefrontal cortex is engaged by cognitive effort and changes in the recruitment of this region during effortful tasks inform cost/benefit signals tracked by Anterior Cingulate Cortex.

In Chapter 2 and Chapter 3, we showed partial support for these predictions. During both the Learning and Test phases, we showed that frontoparietal network increases, and ventral striatum decreases activity across the performance of increasing effort. However, neither change in frontoparietal network nor in ventral striatum predicted change in effort avoidance. Instead, a trend-level relationship indicated that demand-seeking participants who exhibited the greatest demand-seeking behavior displayed the greatest increase in frontoparietal network recruitment for the harder tasks. Furthermore, demand-seeking participants selected those tasks that yielded greater task-switching frequency. Taken together, these results argue against the intrinsic cost associated with cognitive control.

The cost/value associated with effort, it is argued, is learned via reinforcement learning mechanisms, while employing brain regions that are associated with cognitive control and reward processing on a trial-by-trial manner (Botvinick et al., 2007). Under this presupposition, we expected to see reduced reward network activity with prolonged experience with effortful task execution in demand avoiding participants. As expected, on average, ventral striatum reduced its activity with increasing effort requirements during the Learning phase; however, there was no change in its activity across time.

Parkinson's Disease patients who are off their dopaminergic medication have been found to learn better from punishment (Frank et al., 2004), via greater dopaminergic dips during aversive events. A recent study illustrated that Parkinson's Disease patients who were off their dopaminergic medication showed greater cost of response conflict in a probabilistic reward learning task (Cavanagh et al., 2017). Under the same avoidance learning account and the cost of control hypothesis, we expected to see diminished dopaminergic tone to predict greater demand avoidance behavior in the absence of reward. However, preliminary results have indicated that

patients with diminished dopaminergic tone displayed greater demand seeking behavior. These results argue against the mechanistic role of dopamine in relieving the intrinsic cost of cognitive control.

Overall, the results argue against the processing cost accounts by demonstrating that indices of cognitive control do not necessarily register as a cost signal, either on average, on a trial-by-trial basis, or via modulations in dopaminergic processing.

Resource Cost

Unlike the processing cost accounts, the resource cost accounts of cognitive effort do not attribute the cost of effort to a specific process. As such, *biological, motivational and performance cost* accounts have distinct predictions regarding the role of performance measures, such as error likelihood and time-on-task on effort avoidance.

Biological cost accounts argue that during the experience of effort, biological resources get depleted over time. Under this account, we predicted that demand avoidance would increase with prolonged task durations. However, across two studies, we demonstrated that demand avoidance is stable across four blocks of the Test phase, providing counter-evidence for the ego-depletion observation in the literature.

Performance cost accounts argue that the cost of cognitive effort is the greater error likelihood and/or time-on-task during effortful task execution. Under this account, we predicted that participants would choose tasks that yield the least error rates and/or response times. However, across two studies (Chapter 2 & Chapter 3), the results indicated that demand-avoiding participants' selections were best explained by task-switching probability, and not task performance. Also at the group level, performance cost accounts would predict participants to display the greatest demand-avoiding behavior, which have the slowest response times or the

highest error rates. However, we illustrated that demand-avoiding participants overall had faster response times compared to demand-seeking participants. Furthermore, there were no differences in error rates across demand groups.

The opportunity cost accounts argue that the cost of effort will scale with the number of alternative tasks in the environment, including internal processes, such as mind-wandering and self-referential thinking (Kurzban et al., 2013). Under this account, we expected the recruitment of brain networks that are associated with mind-wandering to predict demand avoidance behavior, during effortful task execution. Consistent with this prediction, the Chapter 2 results illustrated that demand-avoiding participants were distinguished from demand-seeking participants by reduced DMN inhibition during effort execution, indicating that demand-avoiding participants, who least suppressed DMN activity during effort execution, exhibited the greatest demand avoidance.

Moreover, opportunity cost accounts argue that the cost of cognitive control is equal to the subjective value of other alternative rewarding tasks. As the number of alternative tasks, and their subjective value might be impossible to determine, an extension of this account suggests that average reward rate of the environment can be tracked by tonic dopamine (Niv et al., 2007). Accordingly, if the average reward rate of the environment – as tracked by tonic dopamine – is high, the cost of investing in an effortful task will be equally high. Thus, an agent with high opportunity cost of dopamine will be expected to show greater response vigor, as well as greater demand avoidance (Cools, 2015). Under this account, we would expect participants who display the greatest response vigor (in terms of response times) to show the greatest demand avoidance behavior. We provide partial support for this account. While demand-avoiding participants in Chapter 2 showed faster response times compared to demand-seeking participants, response

times did not predict effort selections over and beyond task-switching probability in any of the studies.

Finally, it has been suggested that dopaminergic tone at the striatum might operate at the expense of dopaminergic tone in the prefrontal cortex. According to this account, while greater prefrontal dopamine might promote goal-stabilization and greater task performance, greater striatal dopaminergic tone would increase the opportunity cost of time, tipping the system away from goal-stabilization and leading to effort avoidant behavior. Thus, under the opportunity cost of dopamine hypothesis, we would expect Parkinson's Disease patients who were on their dopaminergic medication to exhibit the greatest demand avoidance behavior, and the fastest response times. We provide partial support for this prediction. While dopamine-depleted patients showed the slowest response times, and the greatest demand-seeking behavior, effort selections were best accounted by task-switching probability, not response times. Additionally, although patients who were on their medication during both the Learning and Test phases showed the fastest response times, participants who were on their medication during Learning and off their medication during the Test phase demonstrated a trend toward greater demand-avoiding behavior, lending evidence against the opportunity cost hypothesis.

Overall, across all studies, we provided impartial support for the *processing* and *resource cost* accounts of cognitive effort, indicating that the existing hypotheses of cognitive effort cannot fully account for all effort decisions. Instead, the current research program suggests that indices of cognitive effort might register as utility instead of a cost, depending on the individual and/or dopaminergic tone, in the absence of external reward. While effort exertion might hold an intrinsic cost that can be tracked by reduced dopamine, the benefits of cognitive effort might yield gains in dopaminergic functioning in some contexts (Di Domenico & Ryan, 2017). The

conditions under which the benefits of cognitive effort outweigh its costs are currently unknown. Yet, understanding these conditions is crucial for developing clinical interventions for populations that are characterized by motivational and/or dopaminergic anomalies. Thus, the question remains as to what mechanism underlies this individual variability in the intrinsic valuation of cognitive effort.

Chapter 2 demonstrated that almost half of the sample consisted of demand-seeking participants, despite displaying similar brain recruitment during task performance. The only differentiating factor between demand-seeking and demand-avoidant individuals was slower response times observed in demand seekers. A closer look at task-performance shows that demand seekers were also relatively less accurate than demand avoiders (Table 1.2.), indicating that demand seekers displayed overall poorer task performance compared to demand avoiders. These group differences in task performance might be due to either deficient task engagement in demand seekers, or a reflection of group differences in structural capacity limits. The former possibility is incongruent with the observation that these participants significantly selected the harder task. On the other hand, Butler et al. (2010) showed that task performance during task-switching negatively correlated with working memory capacity. Thus, it is possible that demand-seeking individuals overall had greater limitations on a global capacity. Under the EVC account of effort allocation, this restriction on control capacity must reflect the likelihood of engaging in control. On the contrary, we observed that participants, who might have reduced control capacity, to seek greater control recruitment. Individual differences in working memory capacity have been demonstrated to predict striatal dopamine synthesis capacity in the striatum (Cools et al., 2008), indicating the possibility that demand seekers differed from demand avoiders in having lower dopaminergic tone. A similar observation was made for PD patients in Chapter 4.

PD patients with diminished DA showed the greatest demand-seeking behavior, as well as the worst task performance, across experimental groups. Importantly, in both experiments, time-on-task did not explain choice behavior, indicating that time costs did not underlie demand-seeking behavior within individuals; yet, it explained group characteristics in predicting effort choices. These results open the possibility that effortful task exertion might have subjective utility for individuals who have greater room for performance improvement.

Empirical Delineation of Cost accounts

In summary, the opportunity cost account of cognitive effort avoidance predicts that effort avoidance will increase as the value of or the orientation to the alternative task increases. The average reward rate of the context registers as the opportunity cost of attending to the current task, increasing the likelihood to favor flexible control strategies over stability. Thus, the opportunity cost hypothesis can be tested by manipulating average reward rate of the environment, the value of competing task options, or comparing the effects of tasks that require flexible and stable control strategies on demand avoidance.

The opportunity cost account predictions already run counter to the predictions of control costs, and thus can be distinguished from them easily. While control cost account predicts that demand avoidance will scale with the control intensity of the current task, opportunity cost incorporates the value of the alternative task into the decision-making process. However, note that it is possible that the value of the alternative task will influence the control strategy by favoring flexibility over stability. For example, in a task-switching task where the probability of task-switching is high, switch costs decrease as a sign of reduced control allocation to the current task. Thus, control cost account would predict that reduced control intensity during high task-

switch tasks should reduce demand avoidance. Opportunity cost, on the other hand, would predict greater demand avoidance with greater value of the alternative task.

Another account that can explain demand avoidance is the speed of learning a controlled task. According to this account, it is not the overall control demands of the task that determines its cost, but the speed of which it gets automatized. This account would predict that trial-trial improvements in task performance would underlie the cost/value of the task. While the opportunity cost account assumes that greater response vigor implicates greater average reward rate of the environment, and thus greater cost of attending to the current task, the learning account would predict that the task that yields faster response times will be selected only if the slope of response of times across trials is greater than that of the task with slower response times. Our studies showed that demand avoiders had faster response times across all effort levels compared to demand seekers, supporting the opportunity cost account. In addition, demand seekers showed greater learning slopes at all effort levels compared to demand avoiders, supporting the learning account. A similar picture emerged with PD patients, out of which participants off-their medication on both days who sought the harder task, showed the greatest learning slopes for the harder task, as well as, the slowest response times compared to other patient groups. It is important to note that the slope of learning might be a function of overall response times. For example, participants who initially had slower response times might have greater room for improvement. Our designs did not permit testing for these effects separately. Future studies should manipulate or model trial-by-trial improvements in easy versus difficult tasks, where the response times are initially equated.

Thus, we argue that in the absence of external reward, cognitive effort might acquire an intrinsic value depending on the learning opportunity it provides for the individual. This learning

opportunity might register as a dopaminergic gain during effort exertion, the subjective value of which would depend on the individual's striatal dopaminergic tone.

Future Directions: Testing the 'Dopaminergic gain hypothesis of optimal effort'

While the majority of cognitive effort literature has focused on the cost of effort, it is also the case that individuals attribute value rather than cost to effort exertion. For example, a recent study showed that ants not only seek high effort options more than low effort ones, but also dispose pheromones that mark the high effort options as memorable, in order to communicate the value of such options to their team members (Czackes et al., 2018). This phenomenon, although paradoxical to the cost of effort, is congruent with the role of dopamine in effort-based decision-making, in contexts where control-demanding activity yields a learning opportunity. For example, in non-human primates, the intrinsic value of information has been illustrated to be encoded by the same dopaminergic neurons that track primary rewards and underlie the preference for seeking immediate information about the reward value of the upcoming stimulus, even when this information-seeking behavior did not influence the size of the upcoming reward (Bromberg-Martin & Hikosaka, 2009). This indicates that the cost of engaging in information seeking behavior – although this behavior did not directly influence the reward outcome – was outweighed by the value of information. Therefore, the control costs in effortful tasks might be counteracted with the value of learning in tasks where an increase in effort is coupled with a learning opportunity.

In humans, states of 'flow' constitute a context in which people enjoy mentally effortful tasks through a learning opportunity. During tasks that induce flow, people enjoy activities that are neither too easy nor too difficult, and thus they prefer ones that are optimally challenging, given their current skills and performance ('optimal effort'), that maximize learning (Nakamura

& Csikszentmihalyi, 2009). During tasks that induce flow, people reported greater positive affect, and increased gains in learning. These observations stand in contrast with the cost of effort hypothesis of cognitive effort, which posits that recruited cognitive control during cognitive effort is associated with negative affect and reduced task value. Thus, in order to reconcile these two seemingly disparate observations, it has been initially hypothesized (Bromberg-Martin and Hikosaka, 2009) that optimally challenging tasks must yield reduced cognitive control activity (e.g. FPN) compared to difficult tasks and hence they are exempt from the intrinsic cost associated with control. However, it was recently shown that optimally challenging tasks recruit similar levels of the FPN to that of difficult effort tasks (Ulrich et al., 2015) indicating that the level of cognitive control employment does not distinguish optimal from difficult effort. What distinguished optimal effort experience from that of easy and difficult effort conditions, in the absence of external reward, was reduced self-referential processing, which correlated with reduced default mode network recruitment (DMN) and greater striatal activity during task performance (Ulrich et al., 2014; Ulrich et al., 2015).

Consistent with the dopaminergic-gain hypothesis of effort, striatal blood oxygenation level dependent (BOLD) activity during the optimal effort condition was greater than that of the easy and difficult tasks, leading to an inverted U-shaped relationship between striatal BOLD and effort. Moreover, people who were more prone to ‘flow’ induction had greater dopamine D2 receptor availability (de Manzano et al., 2013). Such optimal effort conditions have been induced empirically by dynamically adjusting the difficulty of the task to match participants’ current performance levels. For example, if the participant correctly solved the last two trials, the difficulty of the current trial was increased by one level, resulting in an average of 50% accuracy across trials. Thus, in contrast to the experience of too easy and too difficult trials, where

participants were certain that their performance would be either always correct or incorrect, during optimal effort conditions two task parameters were different: 1) participants were not certain of the outcome of their performance ('outcome uncertainty'); and 2) participants' performance on the current trial had a causal effect on the difficulty of the next trial ('controllability'). During these conditions, participants reported greater positive affect, and increased gains in learning, suggesting that optimal effort conditions lead to increased value and greater learning opportunity due to: 1) greater outcome uncertainty; and 2) greater controllability of the task environment. In short, these optimal effort environments are characterized by maximal controllable (reducible) uncertainty about the outcome.

Consistent with the role of outcome uncertainty in reward processing, it was shown that phasic dopaminergic activity keeps track of reward probability, while slow dopaminergic changes tracked reward uncertainty (Fiorillo et al., 2003), indicating that dopamine neurons are instrumental in conveying both parameters of reward. This dual property of dopaminergic neurons helps explain risk-taking and addictive behaviors such as gambling, where the reward is variable rather than fixed. Nevertheless, the same signal would be beneficial in natural learning contexts. For example, studies on artificial intelligence demonstrated that learning could be initiated in robots as a function of the task improvement the learner seeks to maximize by manipulating outcome uncertainty. These learning-progress-driven mechanisms have been shown to initiate spontaneous learning of body movements in robots, as well as the discovery of how to communicate amongst each other (Oudeyer et al., 2016). Similarly, human research on curiosity has indicated that people become increasingly more willing to wait with increasing uncertainty about a lottery outcome, illustrating that the intrinsic value of new information increases as a function of outcome uncertainty and outweighs the cost of waiting (van Lieshout

et al., 2018). These results suggest that task uncertainty might increase the subjective value of effort and the opportunity for learning.

Thus, tasks of ‘flow’ induction utilize this value gain associated with outcome uncertainty. Consequently, optimal effort conditions are designed by dynamically adjusting task difficulty, so that participants’ accuracy on optimal trials is around 50%. This stands in contrast to the experience of easy and difficult tasks, in which the accuracy is above 90% or less than 10%, respectively. Thus, the ‘flow’ experience is always collinear with outcome uncertainty. However, while outcome uncertainty might be necessary to increase the subjective value of effort, it might not be sufficient. Alternatively, it is possible that the observed gains in reward activity are a result of increased arousal during optimal effort conditions that yield greater outcome uncertainty. For example, it was observed (Geana et al., 2016) that tasks that yield too little or too much information are also perceived as more boring than contexts that yielded a useful, but uncertain level of information content, leading to a U-shaped function of boredom. Thus, it is necessary to dissociate the effects of optimal arousal from those of optimal effort, in which the hypothesized gain in value is due to a learning opportunity and not uncertainty-based arousal. To this end, controllability – the extent to which one’s actions determine their outcomes – has also been previously highlighted as a key metaparameter influencing the benefit of control (Boureau et al., 2015).

Thus, future work will aim to clarify the role of dopamine in optimal effort by investigating the effects of dopamine manipulation on effort valuation and the associated reward network activity in tasks that vary in outcome uncertainty and controllability. In order to reconcile individual variability in effort seeking behavior with dopaminergic tone, future work will assess the effects of sulpiride, which is selective for the D2 receptor, because the benefits of

optimal effort, as identified by the ‘flow proneness’ scale, have been shown to depend on D2 receptor availability (de Manzano et al., 2013) and D2 receptor gene polymorphism (Gyurkovics et al., 2016). The beneficial effects of dance therapy on PD symptoms are also argued to be via the cognitive engagement required during coordinated motor movements (Foster et al., 2013). Thus, the inclusion of a cognitive component in motor training might improve brain health via greater neurogenesis and increased synaptic strength within the cortico-striatal loop by reversing dendritic spine loss in striatal (BG nuclei) D2 neurons (Petzinger et al., 2013), subfamily of dopamine neurons that are typically associated with motor inhibition. This dopaminergic gain hypothesis of cognitive exercise is consistent with the observed effects of working memory training on increasing striatal dopamine release in healthy adults (Backman et al., 2011). Specifically, blocking the D2 receptor with sulpiride should reduce the ‘precision’ of the value function (decreasing gains during the optimal effort) in low-dopamine participants, while paradoxically enhancing it in high-dopamine participants. This is based on the observations that D2 receptor agents are more likely to act pre- rather than post-synaptically in high versus low-dopamine subjects (Cools et al., 2009; Frank et al., 2006). Thus, we predict that the modulation of the ‘precision’ of the value function across participants will predict changes in the selection (i.e. valuation) of the optimal effort condition.

References

- Anticevic, A., Cole, M. W., Murray, J. D., Corlett, P. R., Wang, X. J., & Krystal, J. H. (2012). The role of default network deactivation in cognition and disease. *Trends in cognitive sciences*, 16(12), 584-592.
- Apps, M. A., Grima, L. L., Manohar, S., & Husain, M. (2015). The role of cognitive effort in subjective reward devaluation and risky decision-making. *Scientific reports*, 5, 16880.
- Arai, T. *Mental Fatigue*. Teachers' College Contributions to Education, 1912, No. 54, Pp. 115.
- Arrington, C. M., & Logan, G. D. (2004). The cost of a voluntary task switch. *Psychological science*, 15(9), 610-615.
- Bäckman, L., Nyberg, L., Soveri, A., Johansson, J., Andersson, M., Dahlin, E., ... & Rinne, J. O. (2011). Effects of working-memory training on striatal dopamine release. *Science*, 333(6043), 718-718.
- Baddeley, A. D., & Hitch, G. J. (1974). Working memory. *The psychology of learning and motivation*, 8, 47-89.
- Badre, D., & D'Esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical?. *Nature reviews. Neuroscience*, 10(9), 659.
- Badre, D., Lebrecht, S., Pagliaccio, D., Long, N. M., & Scimeca, J. M. (2014). Ventral striatum and the evaluation of memory retrieval strategies. *Journal of cognitive neuroscience*, 26(9), 1928-1948.
- Baker, G. P., Jensen, M. C., & Murphy, K. J. (1988). Compensation and incentives: Practice vs. theory. *The journal of Finance*, 43(3), 593-616.

- Bandalos, D. L., Yates, K., & Thorndike-Christ, T. (1995). Effects of math self-concept, perceived self-efficacy, and attributions for failure and success on test anxiety. *Journal of Educational Psychology*, 87(4), 611.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting linear mixed-effects models using lme4. *arXiv preprint arXiv:1406.5823*.
- Baumeister, R. E., Bratslavsky, E., Muraven, M., & Tice, D. M. (1998). Ego Depletion : Is the Active Self a Limited Resource ?, 74(5), 1252–1265.
- Baumeister, R. F., Heatherton, T. F., & Tice, D. M. (1994). *Losing control: How and why people fail at self-regulation*. Academic Press.
- Becker, G. S. (1965). A Theory of the Allocation of Time. *The economic journal*, 493-517.
- Behzadi, Y., Restom, K., Liao, J., & Liu, T. T. (2007). A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *Neuroimage*, 37(1), 90-101.
- Berridge, K. C., & Robinson, T. E. (1998). What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience?. *Brain Research Reviews*, 28(3), 309-369.
- 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., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological review*, 108(3), 624.
- Botvinick, M. M., Cohen, J. D., & Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: an update. *Trends in cognitive sciences*, 8(12), 539-546.

- Botvinick, M. M., Huffstetler, S., & McGuire, J. T. (2009). Effort discounting in human nucleus accumbens. *Cognitive, affective & behavioral neuroscience*, 9(1), 16–27. doi:10.3758/CABN.9.1.16
- Boureau, Y. L., Sokol-Hessner, P., & Daw, N. D. (2015). Deciding How To Decide: Self-Control and Meta-Decision Making. *Trends in cognitive sciences*, 19(11), 700-710.
- Bromberg-Martin, E. S., & Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron*, 63(1), 119-126.
- Brown, J. W., & Braver, T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science*, 307(5712), 1118-1121.
- Brzezicka, A., Kamiński, J., & Wróbel, A. (2013). Local resource depletion hypothesis as a mechanism for action selection in the brain. *Behavioral and Brain Sciences*, 36(06), 682-683.
- Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network. *Annals of the New York Academy of Sciences*, 1124(1), 1-38.
- Butler, K. M., Arrington, C. M., & Weywadt, C. (2011). Working memory capacity modulates task performance but has little influence on task choice. *Memory & Cognition*, 39(4), 708-724.
- Cavanagh, J. F., & Frank, M. J. (2014). Frontal theta as a mechanism for cognitive control. *Trends in cognitive sciences*, 18(8), 414-421.
- Cavanagh, J. F., Masters, S. E., Bath, K., & Frank, M. J. (2014). Conflict acts as an implicit cost in reinforcement learning. *Nature communications*, 5.

- Cavanagh, J. F., Mueller, A. A., Brown, D. R., Janowich, J. R., Story-Remer, J. H., Wegele, A., & Richardson, S. P. (2017). Cognitive states influence dopamine-driven aberrant learning in Parkinson's disease. *Cortex*, *90*, 115-124.
- Chong, T. T. J., Bonnelle, V., Manohar, S., Veromann, K. R., Muhammed, K., Tofaris, G. K., ... & Husain, M. (2015). Dopamine enhances willingness to exert effort for reward in Parkinson's disease. *cortex*, *69*, 40-46.
- Cole, M. W., & Schneider, W. (2007). The cognitive control network: integrated cortical regions with dissociable functions. *Neuroimage*, *37*(1), 343-360.
- Cole, M. W., Yarkoni, T., Repovš, G., Anticevic, A., & Braver, T. S. (2012). Global connectivity of prefrontal cortex predicts cognitive control and intelligence. *Journal of Neuroscience*, *32*(26), 8988-8999.
- Cools, R., Frank, M. J., Gibbs, S. E., Miyakawa, A., Jagust, W., & D'Esposito, M. (2009). Striatal dopamine predicts outcome-specific reversal learning and its sensitivity to dopaminergic drug administration. *Journal of Neuroscience*, *29*(5), 1538-1543.
- Cools, R., Gibbs, S. E., Miyakawa, A., Jagust, W., & D'Esposito, M. (2008). Working memory capacity predicts dopamine synthesis capacity in the human striatum. *Journal of Neuroscience*, *28*(5), 1208-1212.
- Cools, R. (2015). The cost of dopamine for dynamic cognitive control. *Current Opinion in Behavioral Sciences*, *4*, 152-159.
- Czaczkes, T. J., Brandstetter, B., di Stefano, I., & Heinze, J. (2018). Greater effort increases perceived value in an invertebrate. *Journal of Comparative Psychology*.
- Cohen, J. (1988). Statistical power analysis for the behavioral sciences. 2nd.

- Dayan, P., & Balleine, B. W. (2002). Reward, motivation, and reinforcement learning. *Neuron*, 36(2), 285-298.
- De Houwer, J. (2003). On the role of stimulus-response and stimulus-stimulus compatibility in the Stroop effect. *Memory & Cognition*, 31(3), 353-359.
- de Manzano, Ö., Cervenka, S., Jucaite, A., Hellenäs, O., Farde, L., & Ullén, F. (2013). Individual differences in the proneness to have flow experiences are linked to dopamine D2-receptor availability in the dorsal striatum. *Neuroimage*, 67, 1-6.
- Deci, E. L. (1971). Effects of externally mediated rewards on intrinsic motivation. *Journal of personality and Social Psychology*, 18(1), 105.
- Decot, H. K., Namboodiri, V. M., Gao, W., McHenry, J. A., Jennings, J. H., Lee, S. H., ... & Deisseroth, K. (2017). Coordination of brain-wide activity dynamics by dopaminergic neurons. *Neuropsychopharmacology*, 42(3), 615-627.
- Dewey, J. (1897). The psychology of effort. *The Philosophical Review*, 6(1), 43-56.
- DeYoung, C. G. (2013). The neuromodulator of exploration: A unifying theory of the role of dopamine in personality. *Frontiers in Human Neuroscience*, 7, 762.
- Di Domenico, S. I., & Ryan, R. M. (2017). The emerging neuroscience of intrinsic motivation: a new frontier in self-determination research. *Frontiers in human neuroscience*, 11, 145.
- Dosenbach, N. U., Fair, D. A., Miezin, F. M., Cohen, A. L., Wenger, K. K., Dosenbach, R. A., ... & Schlaggar, B. L. (2007). Distinct brain networks for adaptive and stable task control in humans. *Proceedings of the National Academy of Sciences*, 104(26), 11073-11078.
- Dreisbach, G., & Fischer, R. (2011). If it's hard to read... try harder! Processing fluency as signal for effort adjustments. *Psychological research*, 75(5), 376-383.

- Dreisbach, G., & Fischer, R. (2015). Conflicts as Aversive Signals for Control Adaptation. *Current Directions in Psychological Science*, 24(4), 255-260.
- Dunn, T., & Gaspar, C. (2017). Cue Awareness in Avoiding Effortful Control.
- Dunn, T. L., Inzlicht, M., & Risko, E. F. (2017). Anticipating cognitive effort: roles of perceived error-likelihood and time demands. *Psychological research*, 1-24.
- Dweck, C. S., & Leggett, E. L. (1988). A social-cognitive approach to motivation and personality. *Psychological review*, 95(2), 256.
- Eklund, A., Nichols, T. E., & Knutsson, H. (2016). Cluster failure: why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences*, 201602413.
- Evans, A. H., Pavese, N., Lawrence, A. D., Tai, Y. F., Appel, S., Doder, M., ... & Piccini, P. (2006). Compulsive drug use linked to sensitized ventral striatal dopamine transmission. *Annals of neurology*, 59(5), 852-858.
- Fedorenko, E., Duncan, J., & Kanwisher, N. (2013). Broad domain generality in focal regions of frontal and parietal cortex. *Proceedings of the National Academy of Sciences*, 110(41), 16616-16621.
- Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science*, 299(5614), 1898-1902.
- Flandin, G., & Friston, K. J. (2016). Analysis of family-wise error rates in statistical parametric mapping using random field theory. arXiv preprint arXiv:1606.08199.

- Floresco, S. B., Onge, J. R. S., Ghods-Sharifi, S., & Winstanley, C. A. (2008). Cortico-limbic-striatal circuits subserving different forms of cost-benefit decision making. *Cognitive, Affective, & Behavioral Neuroscience*, 8(4), 375-389.
- Font, L., Mingote, S., Farrar, A. M., Pereira, M., Worden, L., Stopper, C., ... & Salamone, J. D. (2008). Intra-accumbens injections of the adenosine A2A agonist CGS 21680 affect effort-related choice behavior in rats. *Psychopharmacology*, 199(4), 515-526.
- Foster, E. R., Golden, L., Duncan, R. P., & Earhart, G. M. (2013). Community-based Argentine tango dance program is associated with increased activity participation among individuals with Parkinson's disease. *Archives of physical medicine and rehabilitation*, 94(2), 240-249.
- Frank, M. J., & O'reilly, R. C. (2006). A mechanistic account of striatal dopamine function in human cognition: psychopharmacological studies with cabergoline and haloperidol. *Behavioral neuroscience*, 120(3), 497.
- 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.
- Frank, M. J., Seeberger, L. C., & O'reilly, R. C. (2004). By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science (New York, N.Y.)*, 306(5703), 1940–3. doi:10.1126/science.1102941
- Frobose, M. I., Swart, J. C., Cook, J. L., Geurts, D. E., den Ouden, H. E., & Cools, R. (2017). Catecholaminergic modulation of the avoidance of cognitive control. *BioRxiv*, 191015.

- Gailliot, M. T., Baumeister, R. F., DeWall, C. N., Maner, J. K., Plant, E. A., Tice, D. M., ... & Schmeichel, B. J. (2007). Self-control relies on glucose as a limited energy source: willpower is more than a metaphor. *Journal of personality and social psychology*, 92(2), 325.
- Geana, A., Wilson, R. C., Daw, N. D., & Cohen, J. D. (2016). Boredom, information-seeking and exploration. In *Proc. 38th Annu. Conf. Cogn. Sci. Soc., Philadelphia* (pp. 1751-56). Austin, TX: Cogn. Sci. Soc.
- Gherman, S., & Philiastides, M. G. (2017). Human VMPFC encodes early signatures of confidence in perceptual decisions. *bioRxiv*, 224337.
- Giladi, N., & Hausdorff, J. M. (2006). The role of mental function in the pathogenesis of freezing of gait in Parkinson's disease. *Journal of the neurological sciences*, 248(1), 173-176.
- Gluth, S., Rieskamp, J., & Büchel, C. (2012). Deciding when to decide: Time-variant sequential sampling models explain the emergence of value-based decisions in the human brain. *Journal of Neuroscience*, 32(31), 10686-10698.
- Gold, J. M., Kool, W., Botvinick, M. M., Hubzin, L., August, S., & Waltz, J. A. (2015). Cognitive effort avoidance and detection in people with schizophrenia. *Cognitive, Affective, & Behavioral Neuroscience*, 15(1), 145-154.
- Gordon, E. M., Lee, P. S., Maisog, J. M., Foss-Feig, J., Billington, M. E., VanMeter, J., & Vaidya, C. J. (2011). Strength of default mode resting-state connectivity relates to white matter integrity in children. *Developmental science*, 14(4), 738-751.

- Gordon, E. M., Stollstorff, M., Devaney, J. M., Bean, S., & Vaidya, C. J. (2012). Effect of dopamine transporter genotype on intrinsic functional connectivity depends on cognitive state. *Cerebral Cortex*, 22(9), 2182-2196.
- Gu, B. M., Park, J. Y., Kang, D. H., Lee, S. J., Yoo, S. Y., Jo, H. J., ... & Kwon, J. S. (2007). Neural correlates of cognitive inflexibility during task-switching in obsessive-compulsive disorder. *Brain*, 131(1), 155-164.
- Guggenmos, M., Wilbertz, G., Hebart, M. N., & Sterzer, P. (2016). Mesolimbic confidence signals guide perceptual learning in the absence of external feedback. *Elife*, 5, e13388.
- Gyurkovics, M., Kotyuk, E., Katonai, E. R., Horvath, E. Z., Vereczkei, A., & Szekely, A. (2016). Individual differences in flow proneness are linked to a dopamine D2 receptor gene variant. *Consciousness and cognition*, 42, 1-8.
- Hagger, M. S., Wood, C., Stiff, C., & Chatzisarantis, N. L. D. (2010). Ego depletion and the strength model of self-control: a meta-analysis. *Psychological bulletin*, 136(4), 495–525. doi:10.1037/a0019486
- Harlow, Harry F. "Learning and satiation of response in intrinsically motivated complex puzzle performance by monkeys." *Journal of comparative and physiological psychology* 43.4 (1950): 289.
- Hartmann, M. N., Hager, O. M., Reimann, a. V., Chumbley, J. R., Kirschner, M., Seifritz, E., ... Kaiser, S. (2014). Apathy But Not Diminished Expression in Schizophrenia Is Associated With Discounting of Monetary Rewards by Physical Effort. *Schizophrenia Bulletin*, 41(2), 503–512. doi:10.1093/schbul/sbu102

- Hazy, T. E., Frank, M. J., & O'Reilly, R. C. (2010). Neural mechanisms of acquired phasic dopamine responses in learning. *Neuroscience & Biobehavioral Reviews*, 34(5), 701–720. doi:10.1016/j.neubiorev.2009.11.019
- Heckman, J. (1974). Shadow prices, market wages, and labor supply. *Econometrica: journal of the econometric society*, 679-694.
- Hewett, R., & Conway, N. (2016). The undermining effect revisited: The salience of everyday verbal rewards and self-determined motivation. *Journal of Organizational Behavior*, 37(3), 436-455.
- Hidi, S. (2016). Revisiting the role of rewards in motivation and learning: implications of neuroscientific research. *Educational Psychology Review*, 28(1), 61-93.
- Hirsh, J. B., Mar, R. A., & Peterson, J. B. (2012). Psychological entropy: a framework for understanding uncertainty-related anxiety. *Psychological review*, 119(2), 304.
- Hoffman, B. (2010). “I think I can, but I'm afraid to try”: The role of self-efficacy beliefs and mathematics anxiety in mathematics problem-solving efficiency. *Learning and individual differences*, 20(3), 276-283.
- Holroyd, C. B. (2016). The waste disposal problem of effortful control. *Motivation and cognitive control*, 235-260.
- Hughes, M. M., Linck, J. A., Bowles, A. R., Koeth, J. T., & Bunting, M. F. (2014). Alternatives to switch-cost scoring in the task-switching paradigm: Their reliability and increased validity. *Behavior research methods*, 46(3), 702-721.
- Hull, C. (1943). Principles of behavior.

- Hyman, S. E., Malenka, R. C., & Nestler, E. J. (2006). Neural mechanisms of addiction: the role of reward-related learning and memory. *Annu. Rev. Neurosci.*, 29, 565-598.
- 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.
- Kool, W., & Botvinick, M. (2014). A labor/leisure tradeoff in cognitive control.
- 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. doi:10.1371/journal.pone.0072626
- Koshino, H., Minamoto, T., Ikeda, T., Osaka, M., Otsuka, Y., & Osaka, N. (2011). Anterior medial prefrontal cortex exhibits activation during task preparation but deactivation during task execution. *PLoS One*, 6(8), e22909.
- Kray, J., & Eppinger, B. (2006). Effects of associative learning on age differences in task-set switching. *Acta Psychologica*, 123(3), 187-203.
- Kurzban, R. (2010). Does the Brain Consume Additional Glucose During Self-Control Tasks ?, 8(2), 244–259.
- Kurzban, R., Duckworth, A., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences*, 36(6), 661-679.
- Lemogne, C., Delaveau, P., Freton, M., Guionnet, S., & Fossati, P. (2012). Medial prefrontal cortex and the self in major depression. *Journal of affective disorders*, 136(1), e1-e11.

- Levy, D. J., & Glimcher, P. W. (2011). Comparing apples and oranges: using reward-specific and reward-general subjective value representation in the brain. *Journal of Neuroscience*, 31(41), 14693-14707.
- Lewis, M. (1965). PSYCHOLOGICAL EFFECT OF EFFORT 1, 64(3), 183–190.
- Libedinsky, C., Massar, S. A., Ling, A., Chee, W., Huettel, S. A., & Chee, M. W. (2013). Sleep deprivation alters effort discounting but not delay discounting of monetary rewards. *Sleep*, 36(6), 899.
- Liston, C., Matalon, S., Hare, T. A., Davidson, M. C., & Casey, B. J. (2006). Anterior cingulate and posterior parietal cortices are sensitive to dissociable forms of conflict in a task-switching paradigm. *Neuron*, 50(4), 643-653.
- Massar, S. A., Libedinsky, C., Weiyan, C., Huettel, S. A., & Chee, M. W. (2015). Separate and overlapping brain areas encode subjective value during delay and effort discounting. *NeuroImage*, 120, 104-113.
- McCool, M. F., Patel, S., Talati, R., & Ragozzino, M. E. (2008). Differential involvement of M1-type and M4-type muscarinic cholinergic receptors in the dorsomedial striatum in task switching. *Neurobiology of learning and memory*, 89(2), 114-124.
- Musslick, S., Jang, S. J., Shvartsman, M., Shenhav, A., & Cohen, J. D. (2018). Constraints associated with cognitive control and the stability-flexibility dilemma.
- Musslick, S., Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2015). A computational model of control allocation based on the expected value of control. In *The 2nd Multidisciplinary Conference on Reinforcement Learning and Decision Making*.

- Musslick, S., Dey, B., Ozcimder, K., Patwary, M. M. A., Willke, T. L., & Cohen, J. D. (2016). Parallel processing capability versus efficiency of representation in neural networks. *Network*, 8, 7.
- McGuire, J. T., & Botvinick, M. M. (2010). Prefrontal cortex, cognitive control, and the registration of decision costs. *Proceedings of the National Academy of Sciences*, 107(17), 7922-7926. PubMed PMID: 20385798.
- Milyavskaya, M., Inzlicht, M., Johnson, T., & Larson, M. (2017). Reward sensitivity following boredom and cognitive effort: A high-powered neurophysiological investigation. *bioRxiv*, 177220.
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, 7(3), 134–140. doi:10.1016/S1364-6613(03)00028-7
- Muraven, M., Tice, D. M., & Baumeister, R. F. (1998). Self-Control as Limited Resource : Regulatory Depletion Patterns, 74(3), 774–789.
- Naccache, L., Dehaene, S., Cohen, L., Habert, M.-O., Guichart-Gomez, E., Galanaud, D., & Willer, J.-C. (2005). Effortless control: executive attention and conscious feeling of mental effort are dissociable. *Neuropsychologia*, 43(9), 1318–1328. doi:10.1016/j.neuropsychologia.2004.11.024
- Nagase, A. M., Onoda, K., Foo, J. C., Haji, T., Akaishi, R., Yamaguchi, S., ... & Morita, K. (2018). Neural Mechanisms for Adaptive Learned Avoidance of Mental Effort. *Journal of Neuroscience*, 1995-17.
- Nagel, I. E., Preuschhof, C., Li, S. C., Nyberg, L., Bäckman, L., Lindenberger, U., & Heekeren, H. R. (2011). Load modulation of BOLD response and connectivity predicts working

- memory performance in younger and older adults. *Journal of Cognitive Neuroscience*, 23(8), 2030-2045.
- Nakamura, J., & Csikszentmihalyi, M. (2009). Flow theory and research. *Handbook of positive psychology*, 195-206.
- Nejad, A. B., Fossati, P., & Lemogne, C. (2013). Self-referential processing, rumination, and cortical midline structures in major depression. *Frontiers in human neuroscience*, 7, 666.
- Niendam, T. A., Laird, A. R., Ray, K. L., Dean, Y. M., Glahn, D. C., & Carter, C. S. (2012). Meta-analytic evidence for a superordinate cognitive control network subserving diverse executive functions. *Cognitive, Affective, & Behavioral Neuroscience*, 12(2), 241-268.
- Nieuwenhuis, S., Schweizer, T. S., Mars, R. B., Botvinick, M. M., & Hajcak, G. (2007). Error-likelihood prediction in the medial frontal cortex: A critical evaluation. *Cerebral Cortex*, 17(7), 1570-1581.
- Niv, Y., Daw, N. D., Joel, D., & Dayan, P. (2007). Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology*, 191(3), 507–20. doi:10.1007/s00213-006-0502-4
- O'doherty, J. P., Hampton, A., & Kim, H. (2007). Model-based fMRI and its application to reward learning and decision making. *Annals of the New York Academy of sciences*, 1104(1), 35-53.
- Oswald, L. M., Wand, G. S., Zhu, S., & Selby, V. (2013). Volunteerism and self-selection bias in human positron emission tomography neuroimaging research. *Brain imaging and behavior*, 7(2), 163-176.

- Oudeyer, P. Y., Kaplan, F., & Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, 11(2), 265-286.
- Oudeyer, P. Y., Gottlieb, J., & Lopes, M. (2016). Intrinsic motivation, curiosity, and learning: Theory and applications in educational technologies. In *Progress in brain research* (Vol. 229, pp. 257-284). Elsevier.
- Petzinger, G. M., Fisher, B. E., McEwen, S., Beeler, J. A., Walsh, J. P., & Jakowec, M. W. (2013). Exercise-enhanced neuroplasticity targeting motor and cognitive circuitry in Parkinson's disease. *The Lancet Neurology*, 12(7), 716-726.
- Piguet, C., Sterpenich, V., Desseilles, M., Cojan, Y., Bertschy, G., & Vuilleumier, P. (2013). Neural substrates of cognitive switching and inhibition in a face processing task. *NeuroImage*, 82, 489-499.
- Ploghaus, A., Tracey, I., Gati, J. S., Clare, S., Menon, R. S., Matthews, P. M., & Rawlins, J. N. P. (1999). Dissociating pain from its anticipation in the human brain. *Science*, 284(5422), 1979-1981.
- Pompeu, J. E., dos Santos Mendes, F. A., da Silva, K. G., Lobo, A. M., de Paula Oliveira, T., Zomignani, A. P., & Piemonte, M. E. P. (2012). Effect of Nintendo Wii™-based motor and cognitive training on activities of daily living in patients with Parkinson's disease: A randomised clinical trial. *Physiotherapy*, 98(3), 196-204.
- Pooresmaeili, A., Wannig, A., & Dolan, R. J. (2015). Receipt of reward leads to altered estimation of effort. *Proceedings of the National Academy of Sciences*, 112(43), 13407-13410.

- 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), 14080–90. doi:10.1523/JNEUROSCI.2752-10.2010
- Ray, K. L., Lesh, T. A., Howell, A. M., Salo, T. P., Ragland, J. D., MacDonald, A. W., ... & Carter, C. S. (2017). Functional network changes and cognitive control in schizophrenia. *NeuroImage: Clinical*, 15, 161-170.
- Rosenberg, M. D., Finn, E. S., Scheinost, D., Papademetris, X., Shen, X., Constable, R. T., & Chun, M. M. (2016). A neuromarker of sustained attention from whole-brain functional connectivity. *Nature neuroscience*, 19(1), 165-171.
- Rushworth, M. F., Noonan, M. P., Boorman, E. D., Walton, M. E., & Behrens, T. E. (2011). Frontal cortex and reward-guided learning and decision-making. *Neuron*, 70(6), 1054-1069.
- Ryan, R. M., & Deci, E. L. (2017). *Self-determination theory: Basic psychological needs in motivation, development, and wellness*. Guilford Publications.
- Sakai, K., & Passingham, R. E. (2003). Prefrontal interactions reflect future task operations. *Nature neuroscience*, 6(1), 75-81.
- Salamone, J D, Correa, M., Farrar, a, & Mingote, S. M. (2007). Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. *Psychopharmacology*, 191(3), 461–82. doi:10.1007/s00213-006-0668-9

- Salamone, J. D., Cousins, M. S., & Bucher, S. (1994). Anhedonia or anergia? Effects of haloperidol and nucleus accumbens dopamine depletion on instrumental response selection in a T-maze cost/benefit procedure. *Behavioural brain research*, 65(2), 221-229.
- Samavatyan, H., & Leth-Steensen, C. (2009). The time course of task switching: A speed—accuracy trade-off analysis. *Memory & cognition*, 37(7), 1051-1058.
- Sayali, C., Badre, D. (2018). Neural Systems of Cognitive Effort Avoidance. *Neuropsychologia*
- Schouppe, N., Demanet, J., Boehler, C. N., Ridderinkhof, K. R., & Notebaert, W. (2014a). The Role of the Striatum in Effort-Based Decision-Making in the Absence of Reward. *Journal of Neuroscience*, 34(6), 2148–2154. doi:10.1523/JNEUROSCI.1214-13.2014
- Schouppe, Nathalie, Ridderinkhof, K. R., Verguts, T., & Notebaert, W. (2014b). Context-specific control and context selection in conflict tasks. *Acta psychologica*, 146, 63–6. doi:10.1016/j.actpsy.2013.11.010
- Schuller, I. S. (1999). Procrastination, need for cognition and sensation seeking. *Studia psychologica*, 41(1), 73-85.
- Schultz, W., & Dickinson, A. (2000). Neuronal coding of prediction errors. *Annual review of neuroscience*, 23(1), 473-500.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593-1599.
- Shen, W., Flajolet, M., Greengard, P., & Surmeier, D. J. (2008). Dichotomous dopaminergic control of striatal synaptic plasticity. *Science*, 321(5890), 848-851.

- 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), 217–40. doi:10.1016/j.neuron.2013.07.007
- Shenhav, A., Straccia, M. A., Cohen, J. D., & Botvinick, M. M. (2014a). Anterior cingulate engagement in a foraging context reflects choice difficulty, not foraging value. *Nature neuroscience*, 17(9), 1249.
- Shenhav, A., Straccia, M. A., Botvinick, M. M., & Cohen, J. D. (2016). Dorsal anterior cingulate and ventromedial prefrontal cortex have inverse roles in both foraging and economic choice. *Cognitive, Affective, & Behavioral Neuroscience*, 16(6), 1127-1139.
- Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D., & Botvinick, M. M. (2017). Toward a rational and mechanistic account of mental effort. *Annual review of neuroscience*, 40, 99-124.
- Spreng, R. N. (2012). The fallacy of a “task-negative” network. *Frontiers in psychology*, 3, 145.
- Thura, D., & Cisek, P. (2014). Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron*, 81(6), 1401-1416.
- Ulrich, M., Keller, J., Hoenig, K., Waller, C., & Grön, G. (2014). Neural correlates of experimentally induced flow experiences. *Neuroimage*, 86, 194-202.
- Ulrich, M., Keller, J., & Grön, G. (2015). Neural signatures of experimentally induced flow experiences identified in a typical fMRI block design with BOLD imaging. *Social cognitive and affective neuroscience*, 11(3), 496-507.
- Vallesi, A., Arbula, S., Capizzi, M., Causin, F., & D'Avella, D. (2015). Domain-independent neural underpinning of task-switching: an fMRI investigation. *Cortex*, 65, 173-183.

- van Lieshout, L. L., Vandenbroucke, A. R., Müller, N. C., Cools, R., & de Lange, F. P. (2018). Induction and relief of curiosity elicit parietal and frontal activity. *Journal of Neuroscience*, 38(10), 2579-2588.
- van Veen, V., & Carter, C. S. (2005). Separating semantic conflict and response conflict in the Stroop task: a functional MRI study. *Neuroimage*, 27(3), 497-504.
- Vincent, J. L., Kahn, I., Snyder, A. Z., Raichle, M. E., & Buckner, R. L. (2008). Evidence for a frontoparietal control system revealed by intrinsic functional connectivity. *Journal of neurophysiology*, 100(6), 3328-3342.
- Westbrook, A., Kester, D., & Braver, T. S. (2013). What is the subjective cost of cognitive effort? Load, trait, and aging effects revealed by economic preference. *PLoS One*, 8(7), e68210.
- Whitfield-Gabrieli, S., & Nieto-Castanon, A. (2012). Conn: a functional connectivity toolbox for correlated and anticorrelated brain networks. *Brain connectivity*, 2(3), 125-141.
- Wu, Y., Li, L., Yuan, B., & Tian, X. (2016). Individual differences in resting-state functional connectivity predict procrastination. *Personality and Individual Differences*, 95, 62-67.
- Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., ... & Fischl, B. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology*, 106(3), 1125-1165.
- Yin, S., Wang, T., Pan, W., Liu, Y., & Chen, A. (2015). Task-switching cost and intrinsic functional connectivity in the human brain: Toward understanding individual differences in cognitive flexibility. *PloS one*, 10(12), e0145826.