

**From the holdings of the Department of Special Collections , Princeton University Library,
Princeton, NJ**

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyright law. We do not charge any permission or usage fees for the publication of images of material in our collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff. For more information regarding copyright, credit, and citation please review the information on the Copyright, Credit and Citation Guidelines page of our website.

<https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines>. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. <https://library.princeton.edu/special-collections/ask-us>

Doctoral Dissertations, 2011-2026

Psychology

2015

The intrinsic cost of cognitive control: Behavioral Phenomenology and Neural Correlates

Kool, Wouter

Princeton, NJ : Princeton University

<http://arks.princeton.edu/ark:/88435/dsp0108612q84b>

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. <https://www.copyright.gov/title17/title17.pdf>

THE COST OF COGNITIVE CONTROL:
BEHAVIORAL PHENOMENOLOGY AND NEURAL CORRELATES

WOUTER KOOL

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
PSYCHOLOGY

ADVISER: MATTHEW M. BOTVINICK

JUNE 2015

© Copyright by Wouter Kool, 2015. All rights reserved.

Abstract

Psychological theories have long maintained that tasks are chosen so as to minimize demands for cognitive exertion. This principle, known as the law of least mental effort, had, until recently, not been subjected to a direct test. However, a recent of studies has shown that people tend to avoid cognitive effort. The current work builds on this finding.

An initial set of experiments shows that demand avoidance does not merely reflect a tendency to maximize reward rate, or minimize the time needed to accurately reach task goals. In addition, they provide evidence for an important secondary prediction from the hypothesis that cognitive effort is associated with intrinsic disutility, namely that countervailing incentives increased the propensity to exert mental effort.

Daily life rarely offers a categorical choice between cognitive exertion and rest. Instead, most real life decisions involving effort are time-based, requiring an allocation decision between profitable but mentally demanding activities and activities that are undemanding but also unproductive. Results from four economic-choice experiments indicate that such decisions are guided by a motivation to strike an optimal balance between income and leisure, a principle derived from economic theories of labor supply.

Demands for cognitive control are ubiquitous in every day life, and therefore it is important to consider whether the intrinsic disutility of its exertion affects decisions studied in the broader field of psychological science. An individual difference experiment takes some of the first steps in this endeavor, probing a relationship between the avoidance of demands for mental effort and the capacity for self-controlled behavior.

Earlier research has revealed an important role for the prefrontal cortex in the registration of effort costs, but the role of neuromodulation in this process is less clear. We report a genotyping study that reveals a role for dopaminergic functioning in the representation of mental effort costs, analogous to findings in the animal literature on physical effort-based decision making.

This work reports initial progress in the understanding of the theoretical, behavioral, and neural implications of the intrinsic cost of cognitive control.

Acknowledgements

First and foremost, my deepest gratitude goes towards my parents Cobie van Gernerde and Peter Kool. Without their continual support in doing whatever I chose to do, whether that was to dye my hair green or fly overseas to live in the United States for almost a decade, I would have never had the audacity and sense of unlimited freedom that is necessary to start graduate school in another country. Our weekly Skype conversations have always kept me sane, and I am always grateful for your advice (even if it doesn't necessary look like it). Pap, mam, ik hou van jullie! The rest of my close family also deserves recognition: Jochen, Roland, Sjuul, Mellie, Walter, and Roald, I always look forward very much to seeing you when I come home. I love you all, and my deepest wish is that I could see all of you much more often than I do now.

Next, it should be noted that the work in this dissertation would not have been possible without the collaboration with my advisor Matt Botvinick. Without his wisdom, enthusiasm, and relentless guidance, I would not be the cognitive scientist I am today. His trust in bringing me on board as a visiting student, as a research assistant, and finally as a graduate student, has been pivotal in my career as a doctoral student. I have been very lucky to be a part of his team for six or so years, and will be forever grateful for the opportunities he gave me.

I have also learned tremendously from my collaboration with Nick Turk-Browne and Andy Conway. Thank you so much for your support and guidance in our scientific endeavors, and for your sound advice on all other aspects regarding life as an academic. I am very lucky to have had the opportunity to work closely with the two of you.

Life as a graduate student can be tough, and this makes the people that I got to meet along the way so important. Shane Blackman, thanks so much for the good times we had living together. I think of you as my brother. José Fernandes, life has not been the same since you left for Canada.

Your presence in the lab and in my life has been invaluable. Matt Johnson, I look back immensely fondly of discussing life over endless games of FIFA with you. Kyle, Drew, Peter, and Dan, I will always be grateful for the way in which you adopted me into the brohort. Grad school would have sucked without you. Mike and Anna, you guys are the best together, and I can't wait to continue our friendship in Boston and Cambridge. Kabir and Ben, thank you for providing the much-needed doses of rock, I wish the Benchwarmers could exist forever. Amitai, thanks so much for your support, advice, and friendship in the later stages of my graduate career. I can't wait to see the cool work that will come out of your lab. My cohort, I will always look back with great memories to our first year of graduate school, and our yearly fondue tradition, and I am very proud of all of you for making it all the way through. I am grateful to everyone in the Botvinick lab for their support and to so many others for their friendship.

Finally, I want to thank Blair Miller, for her enduring patience, support, and love. Your kindness has carried me through these last years of graduate school. I can't wait to see what life has in store for us!

Contents

Abstract	i
Acknowledgements	iii
Contents	iv
1 The law of least mental effort	1
1.1 Chapter summary	1
1.2 Introduction	2
1.3 The law of least work	4
1.4 From physical to cognitive work	5
1.5 Demand selection	7
1.6 Outstanding questions in the current research program	11
2 Rewards offset cognitive costs	14
2.1 Chapter summary	14
2.2 Introduction	15
2.3 Experiment 1a: RT minimization or demand avoidance?	15
2.4 Experiment 1b: Cost of effort trades off with reward	25
2.5 General discussion	28
3 A labor/leisure tradeoff in cognitive control	32
3.1 Chapter summary	32
3.2 Introduction	33
3.3 Experiment 2a: Income-compensated wage decrease	38
3.4 Experiment 2b: Income-compensated wage increase	43
3.5 Experiment 2a and 2b: Additional analyses	45
3.6 Experiment 2c: Prospective effort allocation	47
3.7 Experiment 2d: Mapping the utility surface	53
3.8 General discussion	65

4	The role of effort costs in self-control and intertemporal choice	69
4.1	Chapter summary	69
4.2	Introduction	70
4.3	The strength model of self control	71
4.4	A role for effort costs in self control	73
4.5	Experiment 3: Effort avoidance and self-control	74
4.6	General discussion	80
5	Neural underpinnings of control costs	83
5.1	Chapter summary	83
5.2	Introduction	84
5.3	Cognitive control and the brain	85
5.4	Neural structures representing mental effort costs	88
5.5	The neural correlates of self-control	91
5.6	An anatomical intersection between cognitive costs and self-control	94
5.7	Dopamine and effort	96
5.8	Experiment 4: Dopaminergic genes predict effort avoidance	101
5.9	General discussion	110
6	Conclusions	114
6.1	Chapter summary	114
6.2	Summary of the dissertation	115
6.3	The purpose of the cost of cognitive control	117
6.4	The neural implementation of avoidance decisions	121
6.5	Clinical implications	123
6.6	Concluding remarks	123
	Bibliography	124

Chapter 1

The law of least mental effort

1.1 Chapter summary

This chapter reviews recent research suggesting that demands for mental effort carry an intrinsic cost in decision making.

- Even though our ability for executive functioning plays an important role in many aspects of everyday life, attempts to exert cognitive control fail regularly.
- Many researchers have explained this observation by assuming that human decision makers carry an intrinsic bias to minimize cognitive demand.
- The effect of control demands on cognitive performance has been observed in research on decision making and executive control, without providing direct evidence for the notion of a cost of effort.
- Recently, a series of studies has provided the first direct evidence for cognitive costs, opening up a whole new area of research with many intriguing questions and issues.

1.2 Introduction

One of the most amazing functions of the brain is its capacity for *cognitive control*, the set of processes that reconfigure information processing so as to allow the implementation of effortful, non-routine tasks (Miller & Cohen, 2001). The ability to exert cognitive control plays a crucial role in a wide array of mental processes, such as working memory (Braver et al., 1997; Cohen et al., 1997), action representation (Badre, 2008; Badre & D'Esposito, 2009), attention (Rossi, Pessoa, Desimone, & Ungerleider, 2009), reasoning (Christoff et al., 2001), inhibition (Aron, Robbins, & Poldrack, 2004), task switching (Sohn, Ursu, Anderson, Stenger, & Carter, 2000), self-control (Hare, Camerer, & Rangel, 2009) and economic decision making (McClure, Laibson, Loewenstein, & Cohen, 2004). In other words, cognitive control is tremendously important in our daily lives, enabling people to transcend habit and instead plan future goals, suppress socially inappropriate behavior and make adaptive decisions.

Yet, despite its importance in many aspects of life, attempts to exert cognitive control regularly fail. In the lab, even the most able people show reduced performance on incongruent trials of the Stroop task (Stroop, 1935), and make errors on the straightforward problems in the Cognitive Reflection Test (Frederick, 2005), even though they possess the cognitive skills required to perform it accurately (see Figure 1.1).




Figure 1.1. Two simple psychological tasks on which people commit failures of control. A. Stroop stimulus. When required to name the color in which the word is displayed (“red”), people show increased response times and error rates when the word spells out a different color (“green”). B. In the Cognitive Reflection Test, people often respond to the question with an intuitive but wrong answer (\$0.10), in favor of inhibiting this impulse and responding correctly (\$0.05).

Failures of cognitive control are also observed outside of the laboratory. People rely on flawed stereotypes (Fiske & Taylor, 1991), are easily distracted by Facebook when there is homework to do, and consistently fail at attempts to resist eating unhealthy foods (Heatherton, Striepe, & Wittenberg, 1998).

Together, these observations lead to a deceptively simple question: Why don't people exert as much cognitive control as is needed to reach their goals? Why don't people always suppress their habit to read the word in the Stroop task? Why don't dieters exert more self-control in the face of tempting foods? Why do students succumb to procrastination when their ultimate goal is to succeed academically? Why don't people work as hard as possible *all the time*? How can cognitive control be so efficient, yet so fallible?

This is an important question, but it remains largely unknown why efforts to exert cognitive control are prone to failure. There is a wealth of evidence that relates the ability to exert cognitive control to many life outcomes, such as academic success, intelligence, social competence and the effectiveness of coping strategies (Casey et al., 2011; Mischel, 2011; Mischel, Shoda, & Rodriguez, 1989; Shamosh et al., 2008; Tangney, Baumeister, & Boone, 2004). In fact, lapses in cognitive control are associated with substance abuse (Baler & Volkow, 2006; van der Plas, Crone, van den Wildenberg, Tranel, & Bechara, 2008) and mental disorders such as depression (Murphy et al., 2001; Wenzlaff & Bates, 1998).

The current research program focuses on an account assuming there is an intrinsic cost attached to the exertion of mental effort. That is, failures of cognitive control might be a case of 'bounded rationality', or the normative result of a rational cost-benefit analysis, weighing the potential rewards of each option against anticipated costs (Anderson, 1990; Kahneman & Tversky, 1979; H. A. Simon, 1956; Stephens & Krebs, 1986). In this view, performance failures

are not caused by an inability to exert the necessary amount of control, but rather by a motivated unwillingness to do so.

1.3 The law of least work

A similar framework is widely accepted in research on physical effort-based decision making, where it is assumed that, all else being equal, actions are selected to minimize demands for physical labor. This idea was most famously formulated in Hull's (1943) law of less work:

If two or more behavioral sequences, each involving a different amount of energy consumption or work, have been equally well reinforced an equal number of times, the organism will gradually learn to choose the less laborious behavior sequence leading to the attainment of the reinforcing state of affairs. (p. 294)

This basic notion has held currency in psychology since at least the 1920s (see Solomon, 1948) and remains widely influential in modern studies on cost-benefit analyses involving physical costs (e.g., Salamone, Correa, Farrar, & Mingote, 2007; Walton, Kennerley, Bannerman, Phillips, & Rushworth, 2006). The law of least work is supported by extensive empirical evidence. The most frequent and direct approach has been to place subjects (animals or humans) in situations where they must choose between two courses of action associated with different exertional demands. When rewards are equated, a bias is typically observed toward the less demanding course of action (Solomon, 1948; Thompson, 1944).

1.4 From physical to cognitive work

Hull's (1943) least-effort principle primarily considered demands for physical effort, and experiments resulting from this line of work have continued this trend (Blough, 1966; Cuvo, Lerch, Leurquin, Gaffaney, & Poppen, 1998; Friedrich & Zentall, 2004). However, it has been consistently assumed that a similar law is at work in situations that involve demands for cognitive activity. Allport (1954), for example, famously explained social prejudice as a result of effort avoidance, because accurate social judgment requires controlled and effortful processing.

Similar assertions have been made in many other fields. For example, Baroody and Ginsburg (1986) accounted for strategy selection in arithmetic by invoking a “drive for cognitive economy”, W. J. McGuire (1969) characterized human beings as “lazy organisms,” seeking to spend as little mental energy as possible, and Taylor (1981) characterized humans as “cognitive misers”. According to Camerer and Hogarth (1999), “Economists instinctively assume thinking is a costly activity...mental effort is like physical effort—people dislike both.” Smith and Walker (1993) formalized this idea, proposing a theory of economic choice centering on the role of decision costs, linked to the cognitive or computational effort necessary for certain decision making strategies.

The idea that there exists a cost for cognitive activity has been particularly influential in two fields of psychological science, namely in research on judgment and decision making and in the literature on executive functioning.

In the field of judgment and decision making, it is common knowledge that humans tend to produce suboptimal outcomes in choice situations. This tendency has been explained as a result of a reliance on simplifying strategies for gathering and integrating information (Gigerenzer & Goldstein, 1996; H. A. Simon, 1955; Tversky & Kahneman, 1974). That is, more

effortful and accurate strategies are generally avoided because they carry cognitive costs, and are only employed when the incentives are high enough. Simplifying strategies might be favored for non-effort-related reasons; they might speed decisions or lend tractability to complex situations. However, an influential idea has been that decision-makers evaluate tradeoffs between the effort-related costs and the accuracy-related benefits of computationally intensive strategies (Payne, Bettman, & Johnson, 1993; Shah & Oppenheimer, 2008; Shugan, 1980; Smith & Walker, 1993). Here also, it has been argued that choosing simpler but less accurate strategies could be optimal when considering internal costs of effort (for relevant discussion see Anderson, 1990; H. A. Simon, 1956).

The idea that cognitive activity is undesirable has also surfaced as an explanatory principle in research on cognitive control or executive function (Posner & DiGirolamo, 1998; Shiffrin & Schneider, 1977). For example, in research on task switching the notion of processing costs has been a mainstay for many decades. Typically measured in terms of a difference in reaction times (RTs), the cost of task switching is presumed to reflect the control required to switch from one task set to another (Monsell, 2003). Importantly, this cost has been viewed, at least in part, as reflecting the outcome of a strategic decision involving effort (De Jong, 2000), so that increased switch costs reflect tradeoffs with perceived incentives for accurate task completion. Research on attention and target detection has revealed another case where rewards seem to tradeoff directly with the amount of cognitive control invested. In a recent study by Engelmann and Pessoa (2007), for example, sensitivity in a signal detection task increased as a function of incentive value (see also Engelmann, Damaraju, Padmala, & Pessoa, 2009), suggesting that increased attentional control was only employed as long it was matched with appropriate incentives. In a conceptually similar study, Westbrook, Kest and Braver (2013)

showed a direct tradeoff between reward and willingness to perform more difficult working memory tasks. Finally, a number of cognitive modeling enterprises have explicitly incorporated a principle of “minimal control” (Taatgen, 2007; Yeung & Monsell, 2003) or “least-effort” in executive control (Anderson, 1990; Gray, 2000). There is also direct evidence that human agents ‘offload’ control demands when possible, relying on information in the perceptual environment rather than internal working memory or cognitive control representations (Ballard, Hayhoe, Pook, & Rao, 1997; Droll & Hayhoe, 2007).

In sum, the ‘law of least mental effort’ has intuitive theoretical appeal and the potential to explain behavioral biases across a wide range of domains. It is therefore surprising that, despite its widespread application, the ‘law of least mental effort’ had never been subjected to a direct experimental test. Of course, as described above, a wide range of observations has been explained by yielding to the notion of effort avoidance, but in the vast majority of such cases, effort minimization has been proposed as an explanatory principle rather than a hypothesis to be tested in its own right.

1.5 Demand selection

Initial progress towards validating the idea of effort costs was made by Joseph McGuire, in collaboration with Matthew Botvinick. They developed the demand selection task (DST; Kool, McGuire, Rosen, & Botvinick, 2010) to directly test the hypothesis that people avoid the exertion of mental effort. In the DST, participants face a recurring choice between two alternative lines of action, associated with different levels of cognitive control, such as, for example, working memory or task switching demands. Here and in subsequent elaborations of

the DST paradigm, the general prediction is that participants would develop a tendency to select the course of action associated with the least cognitive demand.

The most often used version of the DST (see Figure 1.2) manipulates control demands by varying the frequency of shifts between tasks. In this version, participants respond to single numerical digits. Depending on the color of the numeral (blue or yellow), the participant uses a key-press to render either a parity judgment (odd/even) or magnitude judgment (less/greater than five). Switching between two tasks is generally understood to demand executive control (Monsell, 2003) and evidence suggests that when two task-sets are freely available, people tend to follow the same task repeatedly (Arrington & Logan, 2004).




Figure 1.2. Sample event sequence from the DST.

This simple number task is embedded in a series of choices. Participants draw each trial's numeral stimulus by selecting between two different 'patches', which vary in appearance and relative position from block to block. Crucially, the patches differ in their demands for task

switching. One patch yields numerals that tend to vary in color across trials (90% of trials), imposing executive demand through the requirement to switch stimulus-response mapping. The other target yields numerals that tend to maintain a consistent color (90% of trials).

The entry level prediction for this task is that participants favor courses of action that commit them to less-frequent task switching, selecting trials from the low-demand patch preferentially. This is exactly what was observed; participants in the DST favor the option that commits them to less frequent task switching (Experiment *a*; Figure 1.3). The results appear consistent with a law of least mental effort, the idea that, all else being equal, actions tend to be selected to minimize cognitive demand.




Figure 1.3. Distribution of low-demand selection rates across three different experiments in Kool et al. (2010).

However, there are some alternative explanations for the observed effect. For example, by consistently choosing the low-demand option, participants could finish the experiment slightly earlier, or be motivated to minimize errors. Therefore, a follow-up study was designed to

rule out these explanations, and to confirm that effort avoidance generalizes to other types of control demands. In this version, participants were given a fixed amount of time to spend performing two-digit mental subtraction problems. The patches manipulated whether the problem required carrying a digit, which increases computational complexity of a mental arithmetic problem (Hitch, 1978), and places cognitive demand on executive processes involved in working memory (Fürst & Hitch, 2000). Consistent with the demand avoidance principle, participants preferentially selected the patches with less demanding math problems (Experiment *b*, Figure 1.3). Because the duration of the session was fixed, the results speak against the possibility that participants simply aim to minimize time on task. The experiment also ruled out error avoidance as an alternative explanation, by showing that avoidance preferences formed before errors were committed.

An interesting ancillary finding in this experiment was a significant correlation between high-demand RT and preference for the low-demand alternative, suggesting that those individuals who found the high-demand task to be more cognitively demanding also showed stronger avoidance, just as one would anticipate based on a ‘law of least mental effort’. Varying levels of ability might influence the amount of cognitive demand experienced by individual participants in the same task. The experience of cognitive demand, in turn, could influence avoidance behavior. Because the direction of causality could not be established in this paradigm (individuals with less demand avoidance also gained more practice on the high-demand task), a third experiment was run to test this possibility more directly.

In this experiment, employing the original task-switching version of the DST, RT switch costs were measured in a preliminary period of task switching, involving isolated stimuli, before the choice paradigm was introduced. As before, results showed robust mental effort avoidance

(Experiment 4, Figure 1.3). More importantly, it was shown that the size of the switch cost correlated positively with the demand avoidance effect. Individuals showing a larger switch cost in the preliminary period showed higher levels of demand avoidance in the DST.

The results from the DST, reported in Kool et al. (2010), provided the first direct evidence in favor of the law of least mental effort. Perhaps unsurprisingly, but most certainly reassuringly, people seem to carry an intrinsic bias to avoid the exertion of mental effort or cognitive control. This initial empirical result opens up a whole new area of research. The goal of the present research program is to identify key questions in this new field of scientific inquiry, and shed some initial empirical light on each of these issues.

1.6 Open questions in the current research program

This chapter has considered the idea that cognitive demand weighs as a cost in the cost-benefit analyses underlying decision-making. Indeed, the experiments described above indicated that, *everything else being equal*, people tend to avoid demand. The idea that demand registers as a cost carries the very important prediction that effort costs trade off against incentives in cost-benefit analyses. In other words, avoidance should be reduced when incentives offset the cost of cognitive effort. **Chapter 2** will describe two experiments aimed at this particular prediction. Furthermore, they will also address one more alternative explanation for the demand avoidance effect in the DST, namely that instead of reflecting disutility of cognitive costs, avoidance is caused by participants' drive to maximize reward rate or to minimize RT per trial (Bogacz, Brown, Moehlis, Cohen, & Holmes, 2006).

Because cost-benefit analyses typically result in binary go/no-go decisions, it is tempting to think of the 'law of least mental effort' as implying that all effort-based decisions are

categorical. However, even the earlier studies on physical effort avoidance have almost universally reported graded rather than categorical effects (Solomon, 1948). **Chapter 3** will confront the fact that cognitive effort-related decisions are typically not categorical, but better characterized as time allocation problems between mentally demanding activities and cognitive rest. This chapter will explore whether a model from labor economics, according to which decision makers strive for an optimal balance between (rewards from) labor and leisure, can account for time allocations in cognitive control. It turns out that this framework does not only provide an insight into the nature of the utility function underlying cognitive effort cost-benefit analyses, but also allows for an explanation of self-control failures as a result of motivational cost-benefit tradeoffs.

The work discussed so far inspects the presence of effort costs in relative isolation. However, effort-based decisions are ubiquitous in everyday life. Think of a high school student deciding whether to do her calculus homework or to watch funny Youtube videos, or a doctoral student deciding whether to spend more time working on his dissertation or to just daydream for a while. Given the pervasiveness of demands for cognitive control, it is theoretically important to consider how demand avoidance affects decisions studied in the broader field of decision-making research. Since this dissertation aims to place effort costs in the larger framework of cost-benefit analyses, it appears to be particularly promising to investigate other forms of reward-based decision making. **Chapter 4** will consider self-control as an example of value-based choice that is commonly understood as involving effortful executive functioning.

Chapter 5 will focus on the neural mechanisms that are involved in cognitive effort-based decision making. Earlier work has described which neural structures are important for the registration of effort costs and their implementation in valuation and decision making. Looking

back at the research of self-control, the chapter will report a meta-analysis that reveals an interesting neural link with effort avoidance in a region of prefrontal cortex. Finally, this chapter will explore the role of neuromodulation in the representation of cognitive control costs. Earlier work in physical effort and recent neuroimaging research suggests a link between tonic levels of dopamine in the nucleus accumbens (NAcc) and the intrinsic cost of cognitive control. This prediction will be tested in a genotyping experiment.

Finally, **Chapter 6** will summarize the main results from the work in this dissertation, and speculate on the future directions that research might (or should) take.

Chapter 2

Rewards offset cognitive costs

2.1 Chapter summary

This chapter reports two behavioral experiments that provide further support for the hypothesis that demands for executive control weigh negatively in cost-benefit analyses.¹

- Experiment 1a demonstrates a demand avoidance bias in a novel paradigm. Results show that participants are biased against effortful task set reconfiguration, even in situations where task switching increases the time needed to achieve task objectives.
- Experiment 1b tests the hypothesis that effort avoidance is reduced when incentives are introduced for cognitively effortful lines of action.

¹ These studies comprise Experiments 6A and 6B in a published report to which the first two authors contributed equally (Kool, McGuire, Rosen, & Botvinick, 2010).

2.2 Introduction

The results from the DST reported in the previous chapter square well with a tendency toward demand avoidance. An important aspect of these results is that they provide evidence against error avoidance or minimization of time-on-task as full explanations for avoidance behavior; instead, the results are consistent with the idea that cognitive demand itself carries intrinsic costs. However, the results open a subtle alternative hypothesis. Bogacz et al. (2006) proposed that decision strategies are chosen so as to minimize the time required to achieve task objectives, or in other words, to increase the amount of reward per unit of time (the reward rate). In simple forced-choice decision tasks, like those employed in the DST (and those addressed by Bogacz and colleagues), this amounts to minimizing RT. Since the high-demand option in all the DSTs described in the previous chapter was associated with a larger mean RT, it is possible that participants' avoidance behavior reflected a motivation to minimize RTs on individual trials, thus minimizing the time required to achieve task goals. The current experiments aimed to address this alternative explanation by decoupling simple RTs from the time required to accomplish central task objectives.

2.3 Experiment 1a: RT minimization or demand avoidance?

The experiment employed a new DST, which will be referred to as the *fill/clear task*. The task involved a series of 'games.' At the outset of each game, an 8 by 11 grid (the 'board') appeared, with a random subset of cells filled, all in either green or blue (Figure 2.1). From here, the participant used two response keys to fill or clear cells ('add or remove pieces'), a few at a time, with the ultimate objective of either completely clearing or completely filling the board.

Participants were free to choose, on every step in the game, between adding and subtracting pieces, and between the goals of filling and clearing the board.

Importantly, the effects of the two response keys depended on the color of the pieces in the current display, which varied randomly across steps of the task. One of the keys (say, the left) added four pieces if the color was blue, but removed four pieces if the color was green. The other key (right) had the opposite pattern effects. Thus, if a participant were operating under a ‘fill’ strategy, it would be appropriate to respond left to blue and right to green. The ‘clear’ strategy would call for the opposite stimulus-response mapping. Note that this made it cognitively costly to switch between strategies (Monsell, 2003).




Figure 2.1. Sequence of events in the fill/clear task. At the outset of each game, an 8×11 board appeared, with a random subset of pieces filled in either green or blue. Participants filled or cleared pieces, with the ultimate objective of either completely clearing or completely filling the board. In the current example, the participant presses the right key to fill four pieces at the outset of the game. As the color changes after his response, the participant presses the left key to fill four subsequent pieces. Next, a jump occurs and only four pieces remain on the board. The participant decides to switch strategies and clears all remaining pieces in the grid, thereby winning the game.

This brings us to one final, crucial detail of the task. At some point during many (but not all) games, the participant's key-press yielded a sudden, unpredictable change in the number of pieces on the board. Following such 'jumps,' the game continued as before, with participants free as always to choose between fill and clear strategies. Our primary interest was in cases where the jump invited a change in strategy: cases where (1) the participant was filling the board and a jump yielded a relatively empty board, or (2) the participant was clearing the board and a jump yielded a relatively full board. In each of these scenarios, a motive to minimize the time to goal attainment would call for a task switch following the jump. In contrast, a motive to avoid cognitive demand would call for the less time-efficient strategy of sticking with the strategy in force before the jump.

Methods

Participants. Sixty-two members of the Princeton University and 22 members of the Leiden University communities (17-33 years of age; 50 females) participated in the study. Participation in the study was compensated for with course credit or a nominal payment. In this, and all following experiments, all participants provided informed consent, following procedures approved by the Princeton University Institutional Review Board and the Leiden University ethics committee.

Materials and procedure. The experiment was computer-based and programmed using the Psychophysics Toolbox extensions for Matlab (Brainard, 1997; Pelli, 1997). The protocol alternated between two tasks: the fill/clear task and a filler task involving trustworthiness judgments on face stimuli.

In the fill/clear task, the number of pieces at the outset of each ‘game’ was always a multiple of four, but was otherwise selected randomly without replacement. Participants responded using the “F” and “J” keys, with key-effect mappings (as characterized above) counterbalanced across participants. Except for when ‘jumps’ occurred, each response either added or subtracted four pieces, at randomly selected locations. The color of the pieces in the display (blue or green) was selected randomly following each response. The task was self-paced. When a board was successfully filled or cleared, the words “You win!” were briefly displayed.

Jumps in the state of the board, accompanied by a brief tone, occurred (only once) in a randomly selected 76% of games. On these trials, the timing of the jump was established probabilistically: The chance of a jump after a key press, given that no jump had yet occurred, was established as:

$$p(\text{jump}|n, \text{strategy}) = \begin{cases} 4 / (88 - n) & \text{if } \text{strategy} = \text{fill} \\ 4 / n & \text{if } \text{strategy} = \text{clear} \end{cases}$$

where n is the number of pieces before the jump, and *strategy* was inferred from the participant’s last response prior to the jump. This means that at each step of a game involving a jump, the jump was equally likely to occur on every subsequent step, given a fixed strategy, and also that the jump was guaranteed to occur before the end of the game. The number of pieces following the jump was selected randomly, with the constraint that it could not equal the number prior to the jump or the number that would have normally resulted from the participant’s last response.

Upon completion of each fill/clear game, participants were prompted to press the two response keys simultaneously. As a result, a face from the Productive Aging Lab Face Database

(Minear & Park, 2004) was presented for 3-5 seconds. Participants were instructed to verbally judge the trustworthiness of the face on a scale from 1 to 5, with 1 being lowest and 5 being highest. This filler task served to isolate rounds of the fill/clear task, minimizing carryover of strategy from one round to the next.

Midway through the study a minor modification to the paradigm was introduced. Initially, 57 participants each played a fixed total of 110 games; the remaining participants played a variable number of games for a fixed session duration of 30 minutes.

Analysis. Transitions from one strategy to another in the fill/clear task were predicted to carry switch costs. In order to confirm this, a paired Student's t-test was used to compare mean RTs immediately following jumps between cases where responses did or did not maintain the previously established strategy.

The strategy chosen at the outset of each game was predicted to vary depending on the number of pieces present. To confirm this, the 21 possible initial piece-counts were organized into seven bins (first bin: 4, 8 and 12 pieces, second: 16, 20, and 24 pieces, etc.). For each bin and each participant, we calculated the proportion of cases in which the fill strategy was adopted at game outset, labeling this OFP_i (Outset Fill Proportion in bin i). For illustration, see the blue trace in Figure 2.1.

A similar approach was adopted in analyzing strategy choice following jumps. Post-jump board states were binned as above, and in each bin we calculated the proportion of cases in which the fill strategy was adopted immediately following the jump (Jump Fill Proportion; JFP). This calculation was made separately for cases where the participant had been following the fill strategy immediately before the jump ($JFP_{i,stay}$), and cases where the participant had been following the clear strategy ($JFP_{i,switch}$). For illustration, see the green and red traces in Figure 2.1.

To evaluate whether participants were biased against switching strategies following jumps, we compared post-jump strategy selection to game-outset behavior. Post-jump strategy selection was compared to choice behavior at a game outset, since both situations reflect a state in which selection of the time-efficient strategy should be independent of previous strategy selection. For each participant we averaged OFP_i , $JFP_{i,stay}$ and $JFP_{i,switch}$ across bins, labeling the resulting means OFP , JFP_{stay} and JFP_{switch} . We then used paired-sample t -tests to perform pairwise comparisons, predicting first that JFP_{stay} would be significantly larger than JFP_{switch} , and at a more detailed level that JFP_{stay} would be significantly larger than OFP , while JFP_{switch} would be smaller than OFP .

A second analysis focused in on strategy choice in situations where switch avoidance was likely to delay game completion. This involved focusing on the slice of the data marked out by the gray areas in Figure 2.1. The highlighted points in the green data series derive from situations in which the fill strategy was being pursued just before a jump to a relatively empty board state. The highlighted points in the red data series derive from situations in which the clear strategy was being pursued just before a jump to a relatively full board state. In both of these situations, minimizing the average time to game completion required a switch to the opposite strategy. (Note that, given the presence of switch costs, it might sometimes have been more time-efficient to stay with the pre-jump strategy, even when the opposite strategy would allow game completion in fewer steps. That is, in such cases, the time-cost of the additional steps required would be outweighed by the time saved by avoiding switch costs. Preliminary analyses indicated that, across participants, this situation would only hold in board-state bin 4. This bin was therefore excluded from the relevant analyses.)

To quantify choice behavior in the relevant game situations, we calculated for each participant the proportion of trials on which the pre-jump strategy was maintained post-jump, despite it being time-inefficient, labeling it *JIP* (Jump Inefficiency Proportion):

$$JIP = \frac{1}{6} \left(\sum_{i \in \{1,2,3\}} JFP_{i,stay} \right) + \frac{1}{6} \left(\sum_{i \in \{5,6,7\}} 1 - JFP_{i,switch} \right).$$

We predicted that this value would be greater than *OIP* (Outset Inefficiency Proportion):

$$OIP = \frac{1}{6} \left(\sum_{i \in \{1,2,3\}} OFP_i \right) + \frac{1}{6} \left(\sum_{i \in \{5,6,7\}} 1 - OFP_i \right),$$

the proportion of cases in which the participant selected the time-inefficient strategy at game outset. The grey areas in Figure 2.1 mark out the portions of the choice data involved in the contrast.

Post-jump strategy maintenance might plausibly reflect participants' indifference or inattention when performing the task. To evaluate this possibility, we repeated our analyses, focusing on a subset of games involving what we termed *strategy coherence*. A game was judged to show strategy coherence if (1) the strategy selected at game outset was identical to the strategy selected in the pre-jump state, and (2) the strategy selected immediately post-jump matched the strategy on the final step of the game. We assumed that such consistency in strategy selection reflected a reasonable level of attention to the content of the task.




Figure 2.2. The *fill* proportions OFP_i , $JFP_{i,stay}$ and $JFP_{i,switch}$ in Experiment 1a are plotted for all bin numbers i (1-7). The overall pattern reveals that participants reasonably chose the fill strategy more often when the initial board state was nearer to full than nearer to empty. Postjump strategy choice revealed that participants tended to maintain their established strategy, instead of switching to the other strategy. The shaded areas delineate the contrast $JIP - OIP$. Thin lines show within-subject standard error bounds.

Results

Preliminary analyses confirmed that there were no statistically significant differences between the Leiden and Princeton groups, or between the group run with a fixed number of games and the group run for a fixed time, in the number of responses per game, number of responses per game in which a jump occurred, switch costs, the difference between *OIP* and *JIP*, or mean RT. Subsequent analyses therefore collapsed across these divisions.

RTs. The results showed that post-jump transitions from one strategy to another were associated with higher mean RTs (1421 ms, standard deviation = 366 ms) than when maintaining the established strategy (1014 ms, standard deviation = 269 ms), and this difference was statistically significant, $t(56) = 13.08, p < 0.0001$.

Strategy selection. The green trace in Figure 6 shows the mean values for OFP_i (as defined under Methods). The green and red traces in the figure show, respectively, mean values for $JFP_{i,stay}$ and $JFP_{i,switch}$. Mean values over bins were 0.53 for OFP , 0.38 for JFP_{switch} , and 0.66 for JFP_{stay} . In line with predictions, JFP_{stay} was significantly larger than JFP_{switch} , $t(83) = 11.01$, $p < 0.001$; JFP_{stay} was significantly larger than OFP , $t(83) = 7.79$, $p < 0.001$; and JFP_{switch} was significantly smaller than OFP , $t(83) = 10.61$, $p < 0.001$. Also in line with predictions, we found that JIP (mean = 0.43) was significantly greater than OIP (mean = 0.34), $t(83) = 5.07$, $p < 0.001$ (see Figure 2.5).




Figure 2.3. The Outset Inefficiency Proportion and Jump Inefficiency Proportion in Experiments 1a and 1b. Error bars indicate standard errors of the mean (SEM). * $p < 0.05$, ** $p < 0.001$

In this experiment, 72% of all games displayed strategy coherence, as defined under Methods. In this subset of games, an analogous pattern of results emerged. JFP_{stay} was significantly larger than JFP_{switch} , $t(83) = 8.04$, $p < 0.001$; JFP_{stay} was significantly larger than OFP ,

$t(83) = 5.65, p < 0.001$; and JFP_{switch} was significantly smaller than OFP , $t(83) = 7.36, p < 0.001$; JIP (mean = 0.27) was significantly greater than OIP (mean = 0.20), $t(83) = 3.33, p < 0.01$.

Discussion

Experiment 1a replicated in a new setting the finding that, absent compensating incentives, people tend to avoid cognitive demand. During performance of a multi-step task, participants tended to avoid switching task strategies, even when circumstances made this the fastest way to achieve task objectives. Participants were willing to delay goals in order to avoid a cognitively demanding task switch. This result goes some distance toward assuaging the concern that the bias observed in earlier experiments (Kool et al., 2010) reflected simply a motivation to meet task goals as quickly as possible.

One potential concern attaching to the results of the present experiment is that the task-switch avoidance observed might simply reflect priming. That is, the adoption of a particular strategy might prime associations between stimulus color and manual responses, so that after a jump these associations would bias responding toward the existing strategy (see Hommel, 2004). Note that this would constitute a non-motivational explanation of the avoidance effect. Thus, if priming entirely explained the results of Experiment 1a, the inclusion of incentives should not affect the magnitude of the switch-avoidance effect. On the other hand, if switch-avoidance in the fill/clear task is reflective at least in part of a motivation to avoid cognitive demand, then introducing incentives for early task completion should reduce the effect.

Additionally, this finding would lend support to another hypothesis coming from the idea that effort weighs negatively in cost-benefit decision making, namely. If mental activity

weights as a cost in decision making, then avoidance should be reduced when incentives are introduced that offset the cost of cognitive effort. Experiment 1b tested this prediction.

2.4 Experiment 1b: Cost of effort trades off with reward

Methods

Participants. Fifty-one participants from the Princeton University community (17-21 years of age, 39 females) participated.

Materials and procedure. The task and procedure were the same as those in Experiment 1a, with the important exception that participants were rewarded for each game they completed. Thirty-seven people received 10¢ for each completed game, and fourteen participants were rewarded with 1¢ per game.

Analysis. Choice behavior was characterized using the measures introduced in Experiment 1a. The central predictions, using the terminology established in Experiment 1a, were that the new set of rewarded participants group, when compared to the unrewarded group of Experiment 1a, would show (1) a smaller difference between JFP_{switch} and JFP_{stay} , and (2) more informatively, a smaller difference between JIP and OIP .

Results

Initial analyses revealed that there were no significant differences between the 1¢ and 10¢ groups in the number of steps per game, number of steps per game in which a jump occurred, switch costs or the difference between JIP and OIP ($ps > 0.31$). In the remaining analyses, we collapsed across the two groups.

RTs. As in Experiment 1a, rewarded participants responded more slowly when switching strategies (mean = 1446 ms, standard deviation = 375 ms) than when maintaining the established strategy (mean = 1101 ms, standard deviation = 311 ms) post-jump, and this difference was statistically significant, $t(50) = 11.80, p < 0.0001$.

Strategy selection. When the initial board was nearer to full than nearer to empty, participants chose the fill strategy more often. Mean values over bins were 0.55 for *OFP*, 0.45 for *JFP_{switch}*, and 0.60 for *JFP_{stay}*. Consistent with our earlier findings, *JFP_{stay}* was significantly larger than *JFP_{switch}*, $t(50) = 5.31, p < 0.001$; *JFP_{stay}* was significantly larger than *OFP*, $t(50) = 3.07, p < 0.01$; and *JFP_{switch}* was significantly smaller than *OFP*, $t(50) = 5.67, p < 0.001$. In contrast with Experiment 1a, *JIP* (mean = 0.32) was numerically but not statistically greater than *OIP* (mean = 0.30), $t < 1$ (see Figure 2.5).

In this study, 74% of all games were classified as involving strategy coherence, as defined earlier. In this subset of games, *JFP_{stay}* was significantly larger than *JFP_{switch}*, $t(50) = 2.80, p < 0.01$; *JFP_{stay}* was numerically but not statistically larger than *OFP*, $t(50) = 1.45, p = 0.15$; and *JFP_{switch}* was significantly smaller than *OFP*, $t(50) = 3.16, p < 0.01$; As before, *JIP* (mean = 0.15) was not statistically different from *OIP* (mean = 0.16), $t < 1$.

Paid vs. unpaid. Our central prediction was that the inclusion of incentives for early task completion would reduce the bias against strategy switching. This was tested by comparing whether payment for task completion significantly decreased the difference between *JIP* and *OIP*. A repeated measures ANOVA with Inefficiency Proportion Type (*JIP* vs *OIP*) as the within-subjects factor, Experiment as a between-subjects factor, and Inefficiency Proportion Type as the dependent variable revealed a main effect of Inefficiency Proportion Type, $F(1,133) = 15.78, p < 0.001$, a main effect of Experiment, $F(1,133) = 5.79, p < 0.05$, but, most importantly, an

interaction effect, $F(1,133) = 6.726, p < 0.05$. This effect indicated that, consistent with a motivational account, the paid group displayed a smaller difference between *JIP* and *OIP* when compared to the unpaid group (see Figure 2.5).

Overall, 73% of all games displayed strategy coherence. In this subset of games, a similar repeated measures ANOVA revealed a trending main effect of Inefficiency Proportion Type, $F(1,133) = 3.20, p = 0.08$, a main effect of Experiment, $F(1,133) = 7.07, p < 0.01$, and again an interaction effect, $F(1,133) = 6.88, p < 0.05$.

Discussion

The results from the DST described in the previous chapter and also the results from Experiment 1a, have indicated that, all else being equal, people tend to avoid demand. The idea that demand registers as a cost predicts, additionally, that avoidance should be reduced when incentives are introduced for the exertion of cognitive effort. The present experiment confirmed this prediction in the task setting introduced in Experiment 1a. When rewards were introduced for effortful lines of action, the avoidance tendency disappeared.

The results of the present experiment additionally rule out an alternative explanation for our findings in Experiment 1a, which was that switch avoidance might have simply reflected S-R priming. If this were the entire explanation, it is unclear why the effect would be altered by an incentive manipulation.

Even though the reward of our two incentive groups differed by a factor 10 (1¢ and 10¢), they did not display differential behavior on the fill/clear task. Although this result was not predicted, it may reflect a ceiling effect, since in both groups the difference between *JIP* and *OIP* was not significantly different from zero. Of course, however, despite the rather large sample

sizes our experiments involved, it is not possible to rule out insufficient power. In any event, though provocative, the absence of a difference between the two reward groups does not undermine the interpretability of our more central findings.

2.5 General Discussion

The law of less work, a time-honored principle in research on decision making, has been widely assumed to apply to mental effort. It has frequently been asserted that, all things being equal, people tend to avoid situations carrying a high demand for effortful cognitive processing. Until recently, no attempt had been made to test this assumption in a controlled and systematic fashion. Following up on results from the DST (Kool et al., 2010), we have provided another demonstration of human decision makers' intrinsic bias to avoid demands for cognitive control.

In the fill/clear task, participants are faced with a decision between effortful task-set reconfiguration, or avoiding a costly task switch and increasing the time it would take to complete a trial. In Experiment 1a, where there was no monetary incentive for completing each trial, participants displayed a clear bias toward the less demanding line of action. This result is important, because it rules out that possibility that behavior observed in the earlier DSTs reflected a drive to reach task objectives as quickly as possible. Previous work (Bogacz et al., 2006) has suggested that decision makers strive for an optimal trade off between decisions by maximizing reward rate, or the percentage correct response per unit time. In the DSTs described in the previous chapter, the relative absence of switch costs in the low-demand option decreased average RT when trials were sampled preferentially. This made the low-demand option the optimal line of action, according to Bogacz and colleagues. The results from Experiment 1a rule out this possibility, by showing demand avoidance even in a situation where it is associated with

reduced reward rate.

Like the experiments discussed in the previous chapter, the results presented here show clear support for a ‘law of least mental effort’, mirroring work on the ‘law of least work’ in the physical effort literature (Hull, 1943; Solomon, 1948). Importantly, this law predicts that agents will choose the line of action with the smallest effort provided all else is equal. In fact, if effort weighs negatively in cost-benefit analyses, then this idea predicts that effort avoidance should be reduced when compensating incentives are introduced. The results from Experiment 1b confirmed this hypothesis, lending more traction to the idea that cognitive performance should be viewed as a result of a cost-benefit analysis that considers a range of positive and negative determinants. In the physical effort literature, the behavioral and neural mechanism behind such tradeoffs have been investigated for a long time (Salamone, Cousins, & Bucher, 1994). Here, we show that cognitive effort participates in similar cost-benefit analyses. **Chapter 5** will describe a genotyping study that strengthens this similarity, by revealing a role for tonic striatal and prefrontal dopamine levels in the representation of cognitive control costs; a relationship that has strong empirical support in the physical effort literature (Salamone, Correa, Farrar, Nunes, & Pardo, 2009).

The results open a host of new questions about the nature of cognitive costs, some of which were already prefaced in the previous chapter. Another intriguing question considers the adaptive nature of mental effort costs. In what way would it be rational to avoid the exertion of cognitive control? According to one account, elaborated on in **Chapter 4**, withdrawal of control allows for the preservation of a limited willpower resource (Baumeister, Vohs, & Tice, 2007). Another theory (Kurzban, Duckworth, Kable, & Myers, 2013) suggests that the cost of control constitutes an opportunity cost, since the capacity to engage in parallel activities is reduced when

control demands are high. In this framework, a motivation to minimize cognitive costs would yield increased opportunity for engaging in additional tasks. Alternatively, a bias to minimize cognitive effort may steer people towards less error-prone or more time-efficient tasks (even if those are not enough to explain the bias *per se*). These and other possibilities will be discussed in more detail in **Chapter 6**.

The results from Experiment 1b are consistent with the idea that demands for cognitive activity weigh as a cost that enters in to cost-benefit analysis. This result naturally begs the question of what the underlying utility function of such a tradeoff looks like. In confronting this issue, it should be noted that the DSTs discussed so far have treated effort-based decisions as a distinctly categorical choice between cognitive effort and rest. However, the law of least mental effort, in line with the law of less work before it, does not imply that human decision makers should categorically and uniformly avoid cognitive demand under any and all circumstances. Indeed, studies of physical effort avoidance have almost universally reported graded rather than categorical effects (see, e.g., Solomon, 1948). In fact, across the DSTs in this and the previous chapter, participants also displayed significant variability in the strength of their demand avoidance effect.

One possible explanation for this variability is the extent to which decisions are stochastically related to the outcome of cost-benefit analyses (as is the case in many models of human and animal decision making). If decisions were indeed the result of a non-deterministic evaluation mechanism, effort avoidance can be expected to assume a graded rather than categorical form.

Another, more interesting, possibility stems from the observation that real-life effort based decisions are rarely categorical. Consider, for example, an assistant professor writing a

grant proposal that is due within a couple of days. Here, the question of whether or not she *should* exert effort to finish her work is not very informative. Instead, the decision at hand can be more accurately cast as one that is more time-based. The next chapter will introduce a model from labor economics that is specifically developed to answer such questions about workers time allocations of labor over a fixed time interval, giving us a more detailed insight in the possible shape of the surface of preferences underlying cognitive effort-based decisions.

Chapter 3

A labor/leisure tradeoff in cognitive control

3.1 Chapter summary

This chapter reports four behavioral experiments examining how people allocate time between tasks that demand cognitive control and tasks that have minimal demands for cognitive effort.²

- Daily life frequently offers a choice between activities that are profitable but mentally demanding (cognitive labor) and activities that are undemanding but also unproductive (cognitive leisure).
- The previously discussed research has demonstrated that there is an intrinsic cost associated with the exertion of cognitive control, but very little is known about how cognitive labor/leisure decisions are made.

² Experiments 2a-c were reported in a published article (Kool & Botvinick, 2014).

- We argue that the motivation underlying cognitive labor/leisure decision making is to strike an optimal *balance* between income and leisure, as given by a joint utility function
- Experiments 2a-c demonstrate that decision makers are willing to trade cognitive rest for work, or work for rest, in cases where labor supply theory predicts that higher overall utility can be obtained.
- Experiment 2d provides independent evidence for the assumed shape of the utility surface underlying cognitive labor/leisure decisions, by fitting participants' preferences over a wide range of combinations to a production function derived from economic choice theory.

3.2 Introduction

Imagine a high-school student sitting at her home computer on a school night. Moment by moment, she faces a recurring decision: Should I focus on my calculus homework, or should I take a break to daydream or watch online videos? The student's decision is essentially one between cognitive labor and cognitive leisure, mental work and mental rest. In the language of cognitive psychology, the labor in question involves the effortful engagement of executive control: a set of functions, dependent on the prefrontal cortex, that configure information-processing resources for the execution of computationally intensive, non-routine tasks (Miller & Cohen, 2001). The student's moment-by-moment decision is thus between a gainful activity that demands robust cognitive control and alternatives that do not.

Cognitive labor/leisure decisions play an obvious role in determining academic and career performance. They may also have safety implications in settings such as air-traffic control or power-plant operation; neuroimaging research shows that lapses of attention and response errors are preceded by a shift in activation from dorsal cortical areas underlying cognitive

control to a default-mode network that is characteristically most active at rest (Eichele et al., 2008; Weissman, Roberts, Visscher, & Woldorff, 2006). Furthermore, cognitive labor/leisure decisions may go awry in clinical disorders like depression or addiction, leading to maladaptive decision-making (Baler & Volkow, 2006; Murphy et al., 2001; van der Plas et al., 2008; Wenzlaff & Bates, 1998). Given these and other considerations, it is important to understand how such decisions are made.

The previous chapters have presented evidence suggesting that the exertion of cognitive control is intrinsically costly or aversive, and that control is only robustly recruited in the presence of relevant incentives (Jimura, Locke, & Braver, 2010; Kool et al., 2010; J. T. McGuire & Botvinick, 2010). It thus seems to be the case that cognitive labor/leisure decisions involve a cost-benefit analysis, which weighs the payoff from cognitive work against its inherent cost.

Because cost-benefit analyses typically result in binary go/no-go decisions, it is tempting to think of cognitive labor/leisure decisions as involving a categorical, all-or-none choice between labor and leisure. The aim of the present work was to evaluate a more complex but also more interesting account, according to which decision-making is based on preferences for a particular balance or mixture of cognitive labor and leisure.

A formal framework for understanding labor/leisure tradeoffs is provided by economic models of labor supply (for an introduction see Nicholson & Snyder, 2008). Such models address the question of how workers choose their preferred hours: Given a particular hourly wage and a ‘budget’ of available hours per week, how does a worker choose to allocate time between work and leisure? Classical labor supply theory proposes that workers value both income and leisure, and that overall value or utility is a joint function of both. A graphical depiction of a representative utility function is presented in Figure 3.1A. Here, the x-axis

indicates leisure time, the y-axis total income from labor, and the z-axis utility. The utility function takes the form of a surface, which attaches a subjective value to each pairing of income with leisure time (the lower portion of the figure shows this as a contour plot). A central proposition of labor supply theory is that this utility function is quasi-convex to the origin, favoring combinations of labor and leisure that are relatively balanced.

Economic labor supply theory provides a simple and compelling account of how labor/leisure decisions are made, given this form of utility function. Note that the worker's available hours and wage define a range of attainable income-leisure combinations, referred to in labor supply theory as the budget constraint. An example budget constraint is shown in Figure 3.1A, appearing as a diagonal line in the lower part of the figure and as a grey plane in the surface plot. According to labor supply theory, the worker selects, from among all feasible income/leisure combinations, the one yielding the greatest utility. In Figure 3.1A, this optimum is marked in red. As the figure illustrates, the convex utility function leads to decisions that favor a balance between income and leisure, resisting a strong bias toward either one or the other

Labor supply theory predicts interesting, non-intuitive shifts in choice behavior with variations in wage. For example, Figure 3.1B diagrams the effect of an *income-compensated wage decrease*. Here, the worker is given an initial baseline wage (in blue) and selects the income/leisure combination associated with maximal utility (corresponding, in the diagram, to the point of contact between the budget constraint line and the highest isoutility curve with which it intersects). Next, a wage reduction is imposed (in red), but the worker is also given a 'freebie' payment up front. This unearned income precisely compensates for the wage reduction, in the sense that if the worker chooses the same work hours as she did before the wage change, her total income (including the freebie payment) will also remain unchanged (the blue and red

budget constraints cross at the initial allocation). Despite the availability of this *status quo ante* income/leisure combination, the geometry of the utility surface in Figure 3.1A dictates that the worker will reduce her hours, settling for a smaller total income. To see this, note that the new optimal point lies on an indifference curve associated with higher utility than the previous equilibrium.

An opposite but otherwise analogous effect is induced by an *income-compensated wage increase* (Figure 3.1C). Here, participants are initially given a freebie payment and baseline wage (in blue). In a second session, an increase in wage is paired with a reduction in unearned income (in green). Although the worker once again has access to the same income/leisure combination that she chose before the wage change, labor supply theory predicts that the worker will sacrifice leisure time in order to attain a larger total income.

When such effects are empirically observed, they constitute direct evidence for preferences that balance income against leisure. Income-compensated wage manipulations therefore offer a powerful tool for probing labor/leisure decision-making. In recognition of this, such wage manipulations have been applied in numerous experimental and field studies on labor markets (Charness & Kuhn, 2010; Dickinson, 1999; Fehr & Goette, 2007), as well as in animal conditioning research (Chen, Lakshminarayanan, & Santos, 2006; Kagel, Battalio, & Green, 1995). In the present work, we leveraged the same approach to test for balance-oriented preferences in human cognitive control. Our general hypothesis was that cognitive labor/leisure decisions operate like the labor/leisure decisions studied in labor economics, tending toward a balance between income and leisure. Based on this, our specific prediction was that cognitive labor/leisure decisions should respond to income-compensated wage manipulations as dictated by labor supply theory.




Figure 3.1. A. Above: A utility surface attaching a scalar value to each pairing of income with leisure time. Shown in black are convex isoutility or indifference curves. Below: A contour map, projected from the indifference curves above. The diagonal line represents a budget constraint, covering the range of income-leisure combinations open to a worker facing T available hours and a wage w . The grey rectangle above projects this constraint line back up through the three-dimensional utility surface, with the intersection forming an arc. In both top and bottom plots, the red marker indicates the income-leisure combination with the highest utility in the feasible set. B. An income-compensated wage decrease. The initial, baseline wage, together with the total available time T , defines the initial budget constraint line (blue). The worker chooses the point along this line that maximizes utility (blue point, with associated indifference curve shown in grey). The wage reduction results in a more horizontal budget constraint (diagonal red line). However, it comes along with unearned income (vertical red line), which causes the new budget constraint to intersect the income/leisure combination originally chosen (blue point). Although this combination is thus still available after the wage change, the worker also has access to a higher-utility combination (red point, with associated isoutility contour) involving increased leisure and reduced income. C. An income-compensated wage increase, illustrated in the same fashion. Here an initial unearned income (vertical blue line) plus an initial wage together define a budget constraint (diagonal blue line). When the wage is increased, the unearned income is reduced (vertical green line), making the originally chosen income-leisure combination (blue point) still available. However, the worker now has access to a higher-utility combination involving less leisure and greater income (green point).

In what follows, we report four experiments. In the first three, participants freely allocated time between a highly demanding cognitive task, which paid a wage ('labor'), and a non-demanding task, which did not ('leisure'). The 'leisure' task was closely matched to the 'labor' task on stimulus and response characteristics, demands for physical effort, and intrinsic interest; the primary factor that differentiated the two tasks was cognitive demand, the key ingredient distinguishing cognitive labor from cognitive leisure as we intend these terms. In each of these experiments, we measured changes in time allocation induced by particular wage manipulations, testing predictions derived from the labor supply theory model (Figures 3.1B and 3.1C).

3.3 Experiment 2a: Income-compensated wage decrease

Our first experiment imposed an income compensated wage reduction. As illustrated in Figure 3.1B, the labor supply theory account predicts that this manipulation should induce participants to forego a previously selected labor/leisure combination, choosing instead to sacrifice income by reducing time in cognitive labor.

Methods

Participants. Thirty-three participants (23 female) from the Princeton University community provided consent and received course credit or a nominal payment. All procedures were approved by Princeton University Institutional Review Board.

Materials and procedure. Participants were asked to spend forty minutes performing a combination of two computer-based tasks: a *high-demand* task that placed heavy demands on executive function and a very *low-demand* task that did not. Both tasks comprised of the

sequential presentation of photographs of human faces (one face every 1750 ms). In the low-demand task, if this showed a child, the participant was to depress a joystick button (20% of all trials). To minimize demands for cognitive control, a tone would always accompany the presentation of child images. The high-demand task involved identical stimuli and timing (including the tone accompanying child faces), but called for a button-press when any face matched the one presented three steps earlier (also 20% of all trials). Note that this task involves the active maintenance and updating of human faces in working memory, and therefore requires cognitive control (Braver et al., 1997).

Stimuli for the two tasks were presented on opposite sides of a computer screen, and participants chose between the tasks by directing a joystick to the relevant side. Instructions indicated that participants could allocate their time between the tasks however they pleased, and could switch from one to the other whenever and however often they wished. Importantly, participants were paid a wage, denominated in pieces of M&M's candy (Mars Incorporated, Hackettstown, NJ), for each trial dedicated to the high-demand task. The low-demand task, in contrast, yielded no such pay. A number indicating total accrued reward for the present session was always visible at the top of the computer monitor, and increased with each trial presented during the high-demand task according to the size of the wage. To minimize the possibility that participants had a non-monotonic preference for M&M's in the study, we made sure that the maximal number of candy pieces did not exceed the contents of medium sized bag of M&M's.

Participants were explicitly instructed that payment of the wage was not contingent on the accuracy of their responses on the high-demand task, but only on the number of trials dedicated to this task. To further discourage a special focus on error monitoring, no feedback on response accuracy was provided. Participants were simply instructed, at the outset of the

experiment, to do their best on both tasks. Each session of the experiment began with a short practice session, allowing familiarization with both tasks and with the payoff regime, and concluded with delivery of earned income. Because the tone accompanying child faces was also played during the high-demand task, it was possible that participants used this tone to more easily identify target trials with child faces. However, paired t -tests on the data of this and the following experiment showed that accuracy was not higher for target trials with child faces than for target trials with adult faces, $ps > 0.10$.

The experiment involved two sessions, separated by one to two weeks. In the first, all participants were accorded the same baseline wage (0.029 pieces of candy for each high-demand trial), and cumulative earnings (*income*) and the time allocated to the high- and low-demand tasks — time in *labor* and (relative) *leisure* — were recorded. In session two, an income-compensated wage reduction was imposed: Participants were explicitly told that the wage from session one was reduced by 50%, but that they also received a ‘freebie’ candy payment up front (though were asked not to consume any of it until after completion of the experiment). This unearned income was calibrated, separately for each participant, so as to make it possible for the participant to allocate exactly the same amount of time to labor as in session one and accumulate precisely the same total income (see Figure 3.1B). Our prediction, drawn from labor supply theory, was that the wage change would induce participants to forgo this option, settling instead for a smaller total income along with increased leisure time.

Because the wage level in session one was established *a priori*, without information about each participant’s baseline preferences, we anticipated that the wage might induce some participants to allocate the large majority of this session either to labor or to leisure. It is important to note that such events would not contradict any prediction from labor supply

theory. Despite the balance principle it implies, labor supply theory does allow that particular incentive regimes may yield an optimum strongly favoring either labor or leisure. Furthermore, some participants may have considered their time budget (as defined in the introduction) to extend beyond the forty minutes of the testing period, leading to choices that balanced within-session labor against pre- or post-session leisure. Notwithstanding these points, we were concerned that cases of nearly all-or-none time allocation might reduce the sensitivity of our design by introducing a ceiling or floor effect. For this reason, only participants who allocated more than 5% of each session, i.e., two minutes, to both labor and leisure ($n = 20$) were included in the final data set³. Only participants satisfying this criterion in session one were invited to complete session two.

Analysis. The central prediction was that leisure time would increase (and total income therefore decrease) from session one to session two. Time in leisure was compared between sessions one and two by way of a two-tailed paired t -test, excluding participants ($n = 2$) for whom the difference between sessions fell greater two standard deviations beyond the sample mean. Additional analyses, aimed at verifying task compliance, are described in conjunction with results⁴.

Results

To confirm compliance with task instructions on both high- and low-demand tasks, we computed response accuracies and d' indices for each. Results, summarized in Table 3.1, indicated that participants made an adequate effort to perform both tasks as instructed,

³ There were eleven participants who allocated more than 95% of the session to the high-demand task and two participants who allocated greater than 95% to the low-demand task.

⁴ All t -test analyses reported in the present paper were repeated using a non-parametric procedure (signed-ranks test), yielding similarly significant results in every case.

responding well above chance levels (single-sample t -test, $ps < 0.001$, two tailed). In both sessions, d' scores and accuracy were lower for the high-demand task than for the low-demand task (paired t -tests, $ps < 0.001$, two tailed), validating our demand manipulation.

Table 3.1. *Average Performance Scores (and Standard Errors) for the High-Demand and Low-Demand Tasks in Experiments 2a-c*

Measure	Experiment 2a		Experiment 2b		Experiment 2c	Experiment 2d
	Session 1	Session 2	Session 1	Session 2	Practice	Practice
d'						
Low demand	4.17 (0.10)	4.06 (0.10)	4.28 (0.15)	3.60 (0.26)	3.44 (0.04)	3.39 (0.10)
High demand	1.82 (0.08)	1.88 (0.13)	1.63 (0.10)	1.63 (0.11)	1.69 (0.11)	1.52 (0.06)
Accuracy (%)						
Low demand						
Target	98 (0.4)	96 (0.9)	95 (2.2)	85 (5.3)	100 (0)	92 (2.6)
Non-target	99 (0.2)	99 (0.1)	99 (0.1)	99 (0.3)	99 (0.2)	99 (0.3)
High demand						
Target	58 (3.2)	54 (4.5)	49 (4.0)	42 (4.7)	61 (4.9)	50 (2.1)
Non-target	94 (0.8)	96 (0.6)	95 (0.4)	96 (0.4)	92 (1.7)	94 (0.7)
RT (ms)						
Low demand	549 (24)	552 (23)	599 (32)	674 (51)		518 (8)
High demand	753 (20)	789 (30)	824 (28)	844 (38)		685 (10)

Labor/leisure allocations were positively correlated between the two sessions, $r = 0.72$, $p < 0.001$, indicating that preferences were relatively stable over the intersession interval. However, the wage manipulation had a significant effect. Mean leisure time and income for sessions one and two are shown in Figure 3.2A. On average across participants, the income-compensated wage decrease triggered a 29% increase in leisure time and a 6% decrease in income, a statistically significant shift in the predicted direction, $t(17) = 2.48$, $p < 0.05$, Cohen's $d = 0.58$ (average arc elasticity of labor supply was 0.24).




Figure 3.2. A. Results of Experiment 2a (compare Figure 3.1B). Leisure time is in minutes and income in pieces of candy. Red and blue lines indicate mean budget constraints. B. Results of Experiment 2b (compare Figure 3.1C). C. Results for wage-decrease trial-pairs in Experiment 2c (compare Figure 3.1B). D. Results for wage-increase trial-pairs in Experiment 2c (compare Figure 3.1C). Note that the higher baseline income in this condition reflects relatively high unearned income (see Table 3.2). Error bars indicate within-subject SEM (horizontal error bars are simply a scaled version of the vertical bars).

3.4 Experiment 2b: Income-compensated wage increase

The results of Experiment 2a supported the idea that effort-based decision making is governed by a preference for balance between control and leisure. As predicted by labor supply theory, an income-compensated wage reduction led participants to reduce time in labor, settling for less income, even though it was possible precisely to recreate an earlier-chosen labor/leisure combination. Experiment 2b tested the complementary prediction, by imposing an income-compensated wage increase. As articulated in Figure 3.1C, labor supply theory here predicts that

participants will here forego a previously selected labor/leisure combination, choosing to increase labor in order to accrue a larger income.

Methods

Participants. Twenty-one participants (15 female) members of the Princeton University community participated, providing informed consent and receiving course credit or a nominal payment.

Materials and procedure. The procedure was largely identical to Experiment 2a. However, here, in session one, participants were accorded a lower baseline wage and received an up-front freebie payment. In session two, the wage was doubled and participants received an up-front payment smaller than the payment in session one. This new freebie payment was calibrated so as to make it possible for each participant to allocate exactly the same amount of time to labor as in session one and to accumulate precisely the same total income (see Figure 3.1C). Again, our prediction from labor supply theory was that the wage change would induce participants to forgo this option, causing them instead to sacrifice leisure time in favor of a larger total income.

As in Experiment 2a, measures were taken to mitigate floor and ceiling effects. Here, this was accomplished by focusing analysis on participants ($n = 16$) who allocated at least 5% of each session to both labor and leisure tasks⁵. In contrast to Experiment 2a, all participants completed both sessions one and two.

Analyses. The central prediction was that mean leisure time would decrease (and total income thus increase) from session one to session two. This was tested by a two-tailed paired t -test, excluding participants ($n = 1$) for whom the difference between sessions fell greater two

⁵ There were four participants who allocated more than 95% of the session to the high-demand task and one participant who allocated greater than 95% of the time to the low-demand task.

standard deviations beyond the sample mean. Additional analyses are once again described in conjunction with results.

Results

As summarized in Table 3.1, response accuracies again indicated compliance with task instructions, far exceeding chance levels (single-sample t -test, $p < 0.001$, two-tailed). As in Experiment 2a, participants showed better performance for the low-demand task than for the high-demand task (paired t -tests, $ps < 0.001$, two tailed).

Again as in Experiment 2a, time-allocation was positively correlated across sessions one and two, $r = 0.81$, $p < 0.001$. But once again, the wage manipulation induced significant changes in task choice. Mean leisure time and income for sessions one and two are shown in Figure 3.2B. Across participants, the income-compensated wage increase triggered a 13% decrease in leisure time and a 9% increase in income, once again a statistically significant shift in the predicted direction. $t(14) = 2.36$, $p < 0.05$, Cohen's $d = 0.61$ (average arc elasticity of labor supply was 0.17⁶). This effect remained significant even when the analysis included participants who dedicated less than 5% of one session to either labor or leisure, $t(19) = 3.29$, $p < 0.004$, Cohen's $d = 1.51$.

3.5 Experiments 2a and 2b: Additional analyses

One potential concern in interpreting the time-allocation results from Experiments 1 and 2 is that participants might have based their choices exclusively on the proffered wage, ignoring

⁶ Elasticity is the degree to which supplied labor goes up or down as wages increase or decrease. Arc elasticity is defined as $[(x_1 - x_2)/(x_1 + x_2)]/[(w_1 - w_2)/(w_1 + w_2)]$, where x_1 and x_2 are the minutes spent on the high-demand/high-reward task in the baseline and compensated wage change conditions, and the w 's are the corresponding wage rates.

unearned income. As noted in the introduction, it is not necessarily the case that simple (uncompensated) increases in wage yield increases in work. Indeed, in the context of cognitive tasks, past empirical work has failed to reveal any simple relationship between incentive magnitude and effort investment, as reflected in task performance (Bonner & Sprinkle, 2002; Camerer & Hogarth, 1999). Nevertheless, it is worth asking whether the shifts in effort allocation observed in Experiments 2a and 2b might reflect a reaction to simple wage changes alone, in isolation from unearned income. To evaluate this, we compared leisure time between Experiment 2a, session one and Experiment 2b, session two, sessions that involved equal wages but different freebie payments (pair *a, d* in Figure 3.2). We also compared leisure time in Experiment 2a, session two against leisure time in Experiment 2b, session one, since here again wages were equal but unearned income differed (pair *b, c* in Figure 3.2). In both cases, leisure time differed across the relevant sessions (two sample *t*-tests, $ps < 0.05$). This indicates that time allocation was not based on wage alone. Rather, participants appear to have based their labor/leisure decisions jointly on wage and unearned income.

We interpret the results of Experiments 2a and 2b in terms of effort costs, or the intrinsic cost of cognitive control. It is partly relief from these costs, we propose, that underlies the inherent ‘utility of leisure,’ as it is termed in labor supply theory. However, one might speculate that what is aversive in the high-demand task is not its cognitive demand, *per se*, but rather its association with frequent response errors. Of course, our participants were given no explicit feedback on response accuracy, and knew that wage delivery was not contingent on performance. More importantly, data from another study (Kool et al., 2010) show that human decision makers avoid cognitive demand even in the absence of error commission. Error avoidance thus seems unlikely to be the only factor underlying the present results. However, to

evaluate the potential role of errors, we conducted two additional analyses. First, based on the recommendation of an anonymous reviewer, we tested whether the wage effects from Experiments 2a and 2b would remain significant after controlling for shifts in response accuracy on the high-demand task from session 1 to session 2. Repeated-measures ANCOVAs, using inter-session change in d' as a covariate, revealed intact effects of wage on leisure time for both experiments. Second, within each experimental session, we examined whether participants tended to switch away from the high-demand task after the occurrence of errors. We considered trials where the participant switched from the high- to the low-demand task (*switch* trials) and trials where the participant remained on the high-demand task (*stay* trials), comparing between these the mean accuracy over the preceding run of high-demand trials. Repeated-measures analyses of variance (ANOVAs) revealed no significant difference in either session of Experiments 2a and 2b for history lengths up to eleven trials. At longer history lengths (up to thirty trials), with no correction for multiple comparisons, an effect appeared in only one of four sessions (Experiment 2b, session 2, history lengths 12, 13 and 16-30). In sum, taken together with previous research, the data do not compellingly support an account based entirely on error avoidance. Further evidence along these same lines is presented in conjunction with the next experiment.

3.6 Experiment 2c: Prospective effort allocations

Experiments 2a and 2b provide support for an income/leisure tradeoff in cognitive control allocation, analogous to the one described in economic labor supply theory. However, one might suspect that the tradeoff evident in those experiments arises as an *emergent* feature of cognitive control dynamics. An influential line of research has suggested that cognitive control is subject

to fatigue, or resource depletion (Muraven, Tice, & Baumeister, 1998). If this is correct, then perhaps participants in Experiments 2a and 2b allocated time between tasks by simply responding, moment by moment, on the basis of their current level of fatigue, resting when highly ‘depleted,’ and returning to the high-demand task when ‘replenished.’

If this were the whole story, then the effects observed in our first two experiments, during on-line effort allocation, should disappear when allocation decisions must be made prospectively. The present experiment examined effort allocation in precisely this setting.

Methods

Participants. Nineteen participants (12 female) from the Princeton University community participated, providing informed consent and receiving either course credit or a nominal payment.

Materials and procedure. The experiment started with two minutes of practice on both the high- and low-demand tasks used in Experiments 2a and 2b. Next, prior to entering a twenty-minute task-performance period, participants used a graphical interface to allocate time prospectively between the high- and low-demand tasks (Figure 3.3). Each participant made 32 prospective allocation decisions, involving a variety of wages and up-front payments, on the understanding that one of their choices would be chosen randomly and used to program stimulus presentation during the later task-performance period. Embedded in the choice task was a set of randomly interleaved trial-pairs. One member of each pair served as a baseline, while the other implemented an income-compensated wage decrease or increase (see Table 3.2). On half of all trials, the awarded candy pieces were M&M’s, and on the other half Reese’s Pieces (The Hershey Company, Hershey, PA). Upon completion, participants received the amount of

candy selected on a randomly chosen time-allocation trial. Wage effects were quantified as the average change in work-time between a baseline trial and the associated wage-change trials. We predicted that both wage manipulations would again cause participants to forego their baseline time-allocation choices, leading them to sacrifice income in the case of wage decreases, and to sacrifice leisure in the case of wage increases.

Analysis. Analyses are described in conjunction with results. Tests on wage effects excluded participants ($n = 1$) falling more than two standard deviations from the sample mean, on either wage increase or decrease trials. All t -tests were two-tailed.




Figure 3.3. In the choice phase of Experiment 2c, the graphical interface showed a single pie chart indicating labor (red) and leisure (green) time, along with an associated total income (in pieces of candy). Participants explored the range of available income/leisure combinations, using arrow keys to increase or decrease leisure time in small increments (curved arrow, not shown in user interface). Panels A and C in the figure show the extreme allocations options on one choice trial, and B the combination actually chosen by the participant. In the lower row, panels D and G show extreme allocation options on another trial faced by the same participant, which implemented an income-compensated wage decrease relative to the trial diagrammed in the top row. Although the income-leisure combination from panel B was available to the participant on this trial (panel F), the participant chose a combination involving greater leisure and less income (panel E), consistent with the predictions of labor supply theory.

Table 3.2. *Trial-Type Parameters From Experiment 2c*

Group	Wage	Unearned Income
A	3.6	None
	$3.6 \times 1/2$	Compensatory (> 0)
	$3.6 \times 1/3$	Compensatory (> 0)
	$3.6 \times 1/4$	Compensatory (> 0)
B	1.8	None
	$1.8 \times 2/3$	Compensatory (> 0)
	$1.8 \times 1/2$	Compensatory (> 0)
	$1.8 \times 2/5$	Compensatory (> 0)
C	1.2	48
	1.2×1.5	Compensatory (< 48)
	1.2×2	Compensatory (< 48)
	1.2×3	Compensatory (< 48)
D	0.6	36
	0.6×2	Compensatory (< 36)
	0.6×3	Compensatory (< 36)
	0.6×4	Compensatory (< 36)

Note. Trials were divided into four groups, each containing four specific trial types. Each group included one parameterization that served as a baseline condition for the others (first row in each group). The remaining three trial types each implemented an income-compensated wage decrease (Groups A and B) or increase (Groups C and D) with respect to the relevant baseline. Each trial type formed the basis for four actual trials for each participant. Wages are designated in units of candy pieces per cumulative minute of labor.

Results and discussion

Table 3.1 lists mean response accuracies and d' scores for the low- and high-demand tasks, which as in previous studies were well above chance ($ps < 0.001$), indicating task compliance.

Mean leisure time and income before and after wage changes are shown in Figure 3.2C (for wage decreases) and 2D (increases). Across participants and trials, income-compensated wage decreases triggered a 29% increase in leisure and a 12% decrease in income, $t(17) = 3.41, p < 0.01$, Cohen's $d = 0.80$ (average arc elasticity of labor supply = 0.62), while income-compensated wage increases triggered a 15% decrease in leisure and a 11% increase in income, $t(17) = 2.79, p < 0.05$, Cohen's $d = 0.66$ (average arc elasticity of labor supply = 0.39).

Interestingly, the ‘elasticity’ of time allocation – that is, its sensitivity to a unit change in wage – appeared to differ reliably across participants. The degree to which each participant reduced work in the face of wage reductions correlated with the degree to which they increased work in response to a boost in wages, $r = 0.68, p < 0.001$ (Figure 3.4). Further analyses suggested that floor or ceiling effects were not wholly responsible for this correlation; the size of wage-change effects did not significantly correlate with mean leisure-time allocation, $r = 0.31, p = 0.20$, or with the absolute deviation of leisure-time allocation from 50%, $r = -0.39, p = 0.11$.




Figure 3.4. Individual-participant data from Experiment 3. The depicted correlation indicates that, when participants showed a large shift to leisure after a compensated wage decrease, they also showed a large shift to labor after a compensated wage increase. There was significant variability in this measure of elasticity, or the sensitivity to a unit change in wage across participants. Thin lines indicate the 95% confidence interval.

Beyond its primary goals, the present study also provided an opportunity to address a lingering concern from Experiments 2a and 2b. Under **Additional Analyses**, we compared cases where wage was matched but effort allocation differed, taking the difference as evidence that participants took unearned income into account. However, the sessions compared in that analysis differed not only in terms of unearned income. They also differed in terms of context: One of the sessions (the one drawn from Experiment 2b) was preceded by a session involving a lower wage, while the other (session 1 from Experiment 2a) was not. The difference in effort allocation could thus, arguably, have arisen from an effect of context on the evaluation of wage. Specifically, the wage offered in Experiment 2b, session 2, might have appeared more rich and motivating, given the contrast with the earlier, lower wage. In order to evaluate this possibility, as well as to further validate the role of unearned income, we conducted a linear regression for each participant in Experiment 2c, modeling the chosen leisure time on each trial as a function of (1) the present wage, (2) the present unearned income, and (3) the difference between present wage and the wage offered on the immediately preceding trial (Table 3.3). We tested whether each of these regressors, on average, explained variance in participants' choice behavior. Leisure time was, not surprisingly, predicted by present wage ($p < 0.01$). More importantly, it was also predicted by unearned income ($p < 0.05$), but not reliably predicted by wage difference ($p = 0.41$). This pattern of results further supports the integral role of income compensation in driving the main effects from all three studies.

In order to rule out a key role for error avoidance, we leveraged the data from the initial practice session. We reasoned that if effort allocation were based entirely on error rates rather than cognitive demand, then leisure time on baseline trials should correlate across participants with accuracy on the high-demand task during the practice block. However, no such correlation

was observed, regardless of whether accuracy was quantified as d' , proportion correct, or proportion correct only on target trials ($p > 0.80$ in all cases).

Table 3.3. *Regression Coefficients and Statistics for the Regression Analysis Performed in Experiment 2c*

Measure	Regressor current unearned income	Current wage	Wage _{current-previous}
Regression coefficient (<i>SE</i>)	-0.04 (0.02)	1.34 (0.42)	0.12 (0.15)
<i>t</i> (df)	-2.22 (17)	3.20 (17)	0.85(17)
<i>p</i>	< 0.05	< 0.01	0.41

Note. For each participant, we modeled the time allocated to labor on each trial as a function of the current wage, current unearned income, and the difference in wage between the current and the previous trials.

3.7 Experiment 2d: Mapping the utility surface

The correlational elasticity result from Experiment 2c hints at an interesting prediction from labor supply theory. Consistent with the results from all experiments, this correlation shows that people are willing to trade in some of leisure for labor, and labor for leisure (as long as the new combination is more valuable), and that the degree to which they are willing to substitute one for the other is consistent regardless of which the goods is traded in. In other words, these decisions are influenced by a parameter determining the degree of *substitutability* between labor and leisure.

The degree of substitutability, or elasticity of substitution, between goods is reflected in the shape of the indifference curves. In labor supply theory, these are assumed to be convex towards to the origin. As can be seen in Figure 3.5A, when indifference curves are convex towards the origin, utility maximization over a budget constraint yields optimal combinations that are relatively balanced between labor and leisure.




Figure 3.5. Different forms of substitutability between economic goods. A. In the case of partial substitutability, indifference curves are shaped convex the origin. Any budget constraint will yield optimal utility at a combination that has a balance between the goods, here labor and leisure (red dot). Partially substitutable goods require increasing amounts of one good to compensate for each additional unit decrease in the quantity of the other (see arrows). B. Some goods are perfect substitutes and have utility surfaces with linear indifference curves, here different brands of bottled water. In this case, budget constraints will always yield a corner solution (red dot), since the goods are perfectly equivalent and mixing is not necessary. In other words, each unit decrease in one good can be compensated with a constant quantity of the other (see arrows). C. Perfect complements are goods that have to be obtained in a perfect ratio, such as left and right shoes. Indifference curves for the goods are orthogonal lines, with the optimal combinations on budget constraints always lie on the line indicating the desired ratio. At each of these optimal points, a unit decrease in one good can never be compensated with a unit increase in the other (see arrows). D. Goods that do not mix well together have indifference curves that are concave towards the origin. Like perfect substitutes, any budget constraint will yield a corner solution, but tradeoffs *favor* extreme solutions. In other words, non-mixing goods require decreasing amounts of one good to compensate for each additional unit decrease in the quantity of the other (see arrows).

In the language of economics, pairs of goods that have a utility surface with convex indifference curves are called ‘partial’ or ‘imperfect’ substitutes (Kagel et al., 1995). For the current purpose, this means that labor and leisure are, to some degree, substitutable for each other, but cannot be traded off as completely equivalent. The defining trait of partial substitutes is that the willingness to substitute some amount of the first good against some amount of the second good depends on the particular combination of the goods that the decision maker currently owns. To see this, observe the indifference curves of labor and leisure in Figure 3.5A. At an initial, relatively balanced, combination, marked by the black dot, the red arrow indicates a reduction in labor, and the following black arrow indicates how much leisure should be supplied to *compensate* for this unit decrease in income, i.e., how much leisure needs to be given to return at the original utility level. Importantly, at this new combination with less income and more leisure, a larger amount of leisure needs to be supplied to compensate for another unit decrease in labor. In other words, when decision makers are at a tradeoff involving relatively high amounts of labor, they are more willing to trade this in for leisure, than at more balanced combinations.

If mental effort allocation decisions follow the laws of labor economics, preferences between cognitive control and cognitive rest should also be characterized by a utility surface with convex indifference curves. The results from Experiments 2a-c are consistent with this assumption; the predictions from the income-compensated wage changes are derived from the assumed shape of the utility function. However, with these income-compensated wage changes, we only get an insight into the rank order over two preferences among the entire space of labor/leisure combinations. To gain a more comprehensive perspective, we designed an experiment to assess participants’ preferences over a wide range of labor/leisure combinations

and test whether the shape of the resulting utility surface is consistent with the idea that cognitive work and rest are partial substitutes. In order to do this, we should first consider the other forms of substitutability that can characterize the shape of utility surfaces of pairs of goods.

Figure 3.5B shows the indifference curves associated with two goods that are ‘perfect substitutes’, such as, for example, different brands of bottled water. The indifference curves of perfect substitutes are always perfectly linear, so that the amount of one good needed to compensate for a decrease of one unit in the other good is constant across all combinations. This property of perfect substitutes results in ‘extreme’ allocation decisions over the budget constraint, with highest utility associated with a combination that contains no demand for the good with lesser marginal utility. To make this more concrete, one can think of perfect substitutes as equivalent, with no need for a balanced combination, such as would be the case when spending a budget on two different brands of bottled water.

Another form of substitutability occurs between goods that are called ‘perfect complements’. These goods always have to be obtained in a particular ratio, such as the number of left and right shoes. Perfect complements have indifference curves with an orthogonal angle, as can be seen in Figure 3.5C. This means that an increase in just one good never increases overall utility, and when already at the desired ratio there is no amount of one good to compensate for a unit decrease in the other good. To see this, think of the case of left and right shoes, where an excess amount in either would not be valuable without a commensurate increase in the other.

A final case of substitutability occurs between goods where mixing is not desirable. For example, one could imagine a bar where one has to determine the relatively balance of wine and

beer in the same glass. Most people will prefer a drink consisting solely of beer or solely of wine to many combinations of the two. Goods for which mixing yields lower utility have indifference curves that are concave towards the origin, as depicted in Figure 3.5D. Like the case of perfect substitutes, any linear budget constraint across utility surfaces for these goods will yield a corner solution, but unlike perfect substitutes, the rate of substitution is not constant across combinations. As can be seen in Figure 3.5D, the amount of one good that needs to be supplied to compensate for a unit decrement in the other decreases as the combination are closer towards one of the axes.

The utility surfaces of all the types of goods observed across the panels of Figure 3.5D are fully described by the Constant Elasticity of Substitution function (CES; Arrow, Chenery, Minhas, & Solow, 1961):

$$U(x_1, x_2) = \begin{cases} \alpha(\beta x_1^\rho + (1 - \beta)x_2^\rho)^{1/\rho} & \text{if } \rho \in (-\infty, 0) \cup (0, 1] \\ \alpha x_1^\beta x_2^{(1-\beta)\gamma} & \text{if } |\rho| \rightarrow 0 \end{cases}$$

where x_1 and x_2 are the input quantities of the two goods, and U is the output utility, and α , β , γ , and ρ are the parameters. The parameter $\alpha \in (0, \infty)$ acts as a scaling factor, $\beta \in (0, 1)$ determines the optimal distribution of the input quantities, $\rho \in (-\infty, \infty)$ determines the degree of substitutability between the goods and $\gamma \in (0, 1)$ dictates the rate of returns. When ρ approaches zero in the limit, the CES uses the Cobb-Douglas production function (Cobb & Douglas, 1928) as a special case. The CES is a standard neoclassical production function in the field of economics (R. H. Frank, 1999; Nicholson & Snyder, 2008), and has been successfully applied to labor/leisure decision making in the physical domain (Conover & Shizgal, 2005; Green &

Rachlin, 1991).

The parameter ρ is of particular importance to our current interest. That is, if $\rho = 1$ the resulting utility surface will describe perfect substitutes, if ρ approaches negative infinity the function yields the indifference curves of perfect complements, if $\rho > 1$ the goods do not mix well together, and if $-\infty < \rho < 1$ the goods are partial substitutes (see Figure 3.5). If cognitive labor and leisure are partial substitutes, they should have indifference curves that are convex to the origin, and its utility surface should have a substitutability parameter with $\rho < 1$. In order to test this hypothesis, we recorded participants' preferences over a wide range of labor/leisure combinations and tested whether the fits of the CES to these preferences yielded parameter estimates that were consistent with this prediction.

Methods

Participants. One hundred participants from the Princeton University community provided consent and received nominal payment for completion of the experiment.

Materials and procedure. The experiment consisted of three phases. In the first phase, participant practiced 100 trials of both the high- and low-demand tasks used in Experiments 2a-c. Next, participants were told that, in the third phase, they would perform a combination of both tasks for a period of thirty minutes and possibly earn some M&M candy. In between these two phases, participants evaluated 256 trials that spanned a range of combinations of time on the low-demand tasks (0-30 minutes) and a number of M&M candy (0-75 pieces).

The evaluation phase used the Becker-DeGroot-Marschak auction method (Becker, DeGroot, & Marschak, 1964) to elicit accurate estimates of subjective value for each of the different combinations of low-demand task time and M&M's. For each of the combinations,

participants were asked to offer a bid between \$0 and \$2 indicating the amount they would be willing to pay to perform this combination when compared to a default combination with 0 low-demand task time and 0 M&M's (Figure 3.6). After each bid was entered, a new, randomly selected, combination was presented until participants provided subjective value ratings for each of the 256 combinations of leisure and candy pieces (16 amounts of leisure time $\times$ 16 amounts of M&M pieces).

Participants were told that at the end of the second phase, the computer would randomly choose a combination and that their bid for this combination would be compared against a random bid by the computer. If the computer's bid was lower than the participants', they would pay this price, and continue to perform the randomly chosen combination, receiving the associated amount of M&M's and the difference between \$2 and the computer's bid. If the participants' bid were the lowest, the participant would perform the default combination and earn \$2. Before the bidding procedure started, participants performed a short practice block of the BDM procedure, including a simulated drawing and computer bid.

The third phase consisted of performance of either the default bid or the randomly chosen bid in the BDM procedure, depending on the bidding outcome. Participants completed their time in the low- and high-demand tasks in order, with a clock on screen reminding them of the time left in each of the blocks.




Figure 3.6. Design of the BDM phase of Experiment 2d. Participants provided subjective value ratings for each of 256 (16×16) combinations of pieces of M&M's and time on the low-demand task. A graphic display, depicted on the right, was presented to indicate the contents of the current combination. After each bid, a new random combination was drawn until the matrix with value ratings was completely filled.

Analysis. For each participant, we used the Matlab's *lsqcurvefit* function, to use a nonlinear least-squares method to fit the CES function to the bids as a function of the number of low-demand task minutes and M&M's in each combination. As in Experiments 2a-b, measures were taken to mitigate floor and ceiling effects. Here, this was accomplished by focusing analysis on participants ($n = 67$) who showed sensitivity to both the number of M&M's and leisure time in the combinations, as indicated by an optimal distribution parameter that was not heavily biased towards either good ($0.025 < \beta < 0.975$). We restricted our analyses to these participants, since when β approaches 0 or 1, the CES function reduces to

$$U(x,y) = \alpha x^\gamma$$

with x the amount of the good that is favored by the weighting parameter. Note that in this case

the substitutability parameter ρ is eliminated from the expression.

The central prediction was that the substitutability parameter was smaller than one, indicating partial or imperfect substitutability between leisure and income. Therefore, we used a one-tailed one-sample t -test to compare the average estimate against this predicted value. Other analyses, aimed at verifying task compliance and robustness of the curve fitting procedure, are described in conjunction with results.

Results

Task performance. We started by confirming compliance with task instructions on both high- and low-demand tasks during the practice blocks by computing response accuracies and d' indices. The results, summarized in Table 3.1, showed that these scores were well above chance ($ps < 0.001$), indicating task compliance.

Parameter estimates. Table 3.4 shows the average estimates of the nonlinear least-squares fitting procedure that was performed to capture each participant's bids to the different combinations of leisure time and number of M&M's as a function of the CES expression. Consistent with our predictions, we found that the substitutability parameter ρ was significantly smaller than 1 (mean = 0.85), $t(66) = 2.49$, $p < 0.01$, indicating partial substitutability between the goods.

Table 3.4. *Parameter Estimates of the CES for the Fitting Procedure Performed in Experiment 2d*

Parameter	Estimate (SEM)
α	0.19 (0.04)
β	0.77 (0.03)
γ	1.15 (0.10)
ρ	0.85 (0.06)

When we visually inspected participants' spatially smoothed surface map of preferences, it became clear that the large majority of them showed a positive response to both the amount of M&M's and leisure time. This was confirmed by a series of linear regressions. For each participant and good, we predicted the average bid as a function of the quantity of the good. These analyses yielded positive and significant regression slopes in the majority of the cases for both leisure (96%) and income (97%). However, upon closer scrutiny, we found that some participants' surface maps revealed a non-monotonic or negative preference for one or two of the goods. To quantify this tendency, we ran another set of linear regressions, now concentrating on the last quarter of the range for both goods. A small subset of participants ($n = 5$), showed a significant negative relationship between leisure time and bid value over this range with relatively high amounts of leisure time, $ps < 0.05$ (see Figure 3.7). Interestingly, four of these five participants showed a significant positive relationship between leisure time and bid value over the first half of the leisure time range, $ps < 0.05$ (the remaining participant showed no significant relationship). One of these participants also showed a significantly negative relationship over the range with relatively high income, but a positive slope for the first half, $ps < 0.05$.

The existence of these participants, even though they were a minor subset, surprised us. Why would people show a preference towards a mixture between leisure and labor, independent of the income? One possibility, elaborated on in the Discussion, is that participants considered the two different tasks partial substitutes, maximizing utility over other features of the low- and high-demand tasks. Because the assumption of the CES function that utility increases for both goods does not hold for these participants, it was important to confirm that the partial substitutability effect remained when these cases were excluded from analysis, $t(61) = 2.00$, $p < 0.05$.




Figure 3.7. Average bid values over leisure time for participants ($n=5$) with non-monotonic preferences.

It is possible that this result was driven by a form of attenuation bias, with the nonlinear least-squares fitting procedure biasing the parameter to be lower than the true substitutability of labor and leisure. If this were the case, it is possible that even though average ρ was one, the partial substitutability effect was merely a reflection of a bias in the fitting procedure.

To rule out this alternative explanation, we generated new synthetic data for each participant using the CES function with all individual parameters as estimated previously, except for the substitutability parameter ρ , which was set to 1. Next, we distorted these individual synthetic evaluations by applying random Gaussian noise to each cell in the matrix, with the standard deviation equal to the average absolute error of the participants' original fit. Finally, we fit these new synthetic data to the CES with the same nonlinear least-squares procedure as used before. This process was repeated 100 times, and from these fits we extracted the average substitutability parameter value for each participant. Since the distribution of these fits did not

resemble a normal distribution, subsequent analysis used a non-parametric approach. The median of the estimates for the substitutability parameter was significantly higher than 1 (median = 1.02), $\hat{\rho} = 4.1$, $p < 0.001$. A subsequent paired-samples test revealed a significant difference between the original and synthetic fits for ρ , $\hat{\rho} = 2.97$, $p < 0.05$, suggesting that any bias in the fitting procedure was not sufficient to account for the partial substitutability effect.

Discussion

The results from this study show converging evidence that the utility function underlying mental effort-based decisions favors balanced combinations between cognitive leisure and the rewards from cognitive labor. The average estimate for the elasticity of substitution of between economic goods in the CES function (Arrow et al., 1961) suggested that people consider cognitive leisure a good that is partially substitutable against the rewards from mental work. Note that we found this pattern of results without directly asking participants to substitute work for leisure. Instead, this feature of participants' preferences was derived from fitting a mathematical expression for utility surfaces to their isolated evaluations of combinations of leisure and income. The result supports the findings from Experiments 2a-c, and indicate that cognitive labor/leisure operate like those in labor economics.

A minor subset of our participants (5/67) displayed a preference for a balance between work and leisure, independent of income. At first glance, this observation seems to be at odd with the law of least mental effort. However, it is likely that there are many other factors such as a desire for change, or variety, that affect decision making in the labor/leisure task. Evidence for such a tradeoff comes from work on achievement motivation (Atkinson, 1964). Specifically, Atkinson and Birch (1978) proposed that all activities are subject to a 'consummatory' force,

which reduces the satisfaction associated with engaging in tasks, leading to a natural propensity towards variety in tasks independent of the incentives associated with task execution.

For these participants, one important prediction is that increasing the control demands of the high-effort option would shift the peak inverted-u shaped curve towards the combinations with relatively more cognitive leisure. For the other group of participants, increasing the length of the interval should reveal more non-monotonic preference curves, depending on the size of these consummatory forces associated with both the high- and the low-demand tasks.

3.8 General Discussion

To summarize, we found across the first three experiments that income compensated wage-changes impacted cognitive labor/leisure decisions in a fashion consistent with economic labor supply theory: Even when participants had the option to recreate a previously selected combination of income and cognitive leisure, wage reductions led them to work less and settle for a smaller income, and wage increases led them to give up leisure in favor of a larger income. In a final experiment, participants provided subjective value ratings over a range of labor/leisure combinations, and we found that fits of the CES function to these preferences indicated people treat cognitive rest and work as partial substitutes. The observed pattern of behavior and evaluation suggests that cognitive labor/leisure decisions are guided by preferences that jointly weigh income and leisure.

Our findings directly parallel results from both experimental labor economics (Charness & Kuhn, 2010; Fehr, Goette, & Zehnder, 2009) and animal conditioning research (Chen et al., 2006; Conover & Shizgal, 2005; Kagel et al., 1995), indicating that the principles that guide effort

allocation both in labor markets and in animal foraging may also apply in the more rarified domain of human cognitive control.

On the broader stage, effects of incentives on cognitive effort and performance have been of interest both in basic-science settings (Brehm, Wright, Solomon, Silka, & Greenberg, 1983; Kool et al., 2010; Ryan & Deci, 2000c; Sarter, Gehring, & Kozak, 2006), and in applied settings including education (Ryan & Deci, 2000a; Weiner, 1985). Our work suggests that, within such work, it may be fruitful to apply formal tools from labor economics. For example, economic research has demonstrated non-monotonic effects of incentives on worker effort (resulting in what is known as the ‘backward-bending’ labor supply curve), and interactions between wage incentives and social preferences on worker behavior. Camerer, Babcock, Loewenstein, and Thaler (1997), for example, found that New York taxi drivers supply smaller amounts of labor on busy days, i.e., when their ‘wage’ was high (but, see Farber, 2005). The present findings open up the question of whether effects such as these might also characterize decision-making in cognitive control.

The work presented here relates closely to the experiments discussed in **Chapters 1 and 2**, which lent empirical support the long-standing proposition that the exertion of executive control is intrinsically costly (Botvinick, 2007; Kool et al., 2010). As we saw before, this principle becomes evident in situations where participants choose between two lines of actions associated with different levels of cognitive demand. The demand-avoidance effect fits into the framework we have been developing here via a simple change of sign: In cognitive labor/leisure decisions, the utility of leisure derives, in important part, from the relief it offers from costly control. Equivalently, one may view the cost of control as an opportunity cost, relating to foregone

leisure (Kurzban et al., 2013). The beauty of the labor supply model is that it renders the choice between these two perspectives moot. The shape of the utility function says it all.

Additional studies of the demand-avoidance effect have shown that it is reduced in the presence of countervailing incentives, indicating that the inherent cost of control is weighed against potential rewards (Kool et al., 2010). The present work provides a richer picture of this cost-benefit analysis. By default, one might assume a simple linear model, under which the net value of a course of action is computed by subtracting the cost of control from associated rewards. The present findings suggest, instead, that the marginal cost of control varies as a function of context: A unit increment in effort carries a greater subjective cost when one is already working hard than when one is hardly working.

We began this chapter with an imagined ‘real-life’ example of labor/leisure decision making. The question may now be asked, How do effort costs relate to decisions made in everyday life, or in the broader field of psychological research? Two aspects of our experimental approach call for comment in this connection. First, the decisions made by participants in these experiments were deliberate and explicit, a feature that allowed us to interpret them as reflecting participants’ actual preferences. The deliberate, motivational, cost-benefit framework that has been discussed in this chapter suggests that other cases of reward-based decision making would be a fruitful direction to answer this question. Second, the labor/leisure decision in the experiments involved selecting between two active tasks, a measure that allowed us to match ‘labor’ and ‘leisure’ conditions in terms of stimulus and response characteristics, physical effort demands, and intrinsic task interest. One example of a cognitive labor/leisure decision in everyday life that involves a deliberate choice between two tasks and has been understood to involve reward-based decision making is the case of self-control (Hare et al., 2009; Mischel et al.,

1989; Muraven & Slessareva, 2003). Our earlier homework-versus-video example fits this description, where one throws off the burden of self-discipline in order to indulge in a formerly resisted temptation. The next chapter will be aimed specifically at this question of whether the cost for cognitive activity, studied in this and the previous chapter, has any bearing on decisions made in the broad field of psychological research.

Chapter 4

The role of effort costs in self-control and intertemporal choice

4.1 Chapter summary

This chapter reports an individual differences experiment investigating the relationship between self-control and the cost for cognitive activity.⁷

- The capacity for self-control is critical to adaptive functioning, yet our knowledge of the underlying processes and mechanisms is only incipient.
- Theoretical models of self-control in economics rely on the assumption that self-control is intrinsically costly, but this idea remains in need of empirical validation.
- We report an experiment demonstrating a correlation between demand avoidance and two measures of self-control, providing clear support for the idea of self-control costs.

⁷ A version of this chapter was reported as a published article (Kool, McGuire, Wang, & Botvinick, 2013).

4.2 Introduction

Chapters 2 and 3 demonstrated that human decision makers avoid demands for cognitive control and have shown that this cost can be offset by relevant incentives, suggesting that demands for executive functioning carry intrinsic disutility. Given the important role for cognitive control in every day life and other fields of psychological science (Casey et al., 2011; Otto, Raio, Chiang, Phelps, & Daw, 2013; Paxton, Ungar, & Greene, 2011; Shamosh et al., 2008), it is important to ask whether demands for cognitive control affect choices studied in the broader field of decision making research.

The current chapter will focus on one case of reward-based decision making that has been understood as involving top-down executive functioning, namely the case of self-control: an important, though fallible, ability to resist immediate pleasures in favor of longer-term goals. The importance of this faculty and the consequences of its occasional failure are evident from everyday life. Scientific investigation has linked individual differences in self-control to significant life outcomes, including obesity, academic performance, and mental health (Mischel, 2011; Moffitt et al., 2011; Tangney et al., 2004). As such findings have emerged, the goal of understanding the principles and mechanisms underlying self-control has come increasingly to the fore.

Over recent years, behavioral economics has generated several formal theoretical models of self-control. In what is now arguably the modal model (Berns, Laibson, & Loewenstein, 2007; Fudenberg & Levine, 2006; Gul & Pesendorfer, 2001; Thaler & Shefrin, 1981), the exertion of self-control involves the overriding of a ‘short-term self,’ fixated on immediate rewards, by a ‘long-term self,’ which seeks to maximize reward over the long-term. Recent work in cognitive neuroscience has revealed independent evidence for the dual selves component of this economic

perspective. The prime candidate for the neural analogue of the patient system is the dorsolateral prefrontal cortex (dlPFC) – a region known to be important for executive functioning and cognitive control (Miller & Cohen, 2001). In a study on choice behavior in self-proclaimed dieters, for example, Hare et al. (2009) observed increased dlPFC activity when participants exerted successful self-control by choosing a healthy food over an unhealthy option, whereas reward-related areas exhibited the opposite pattern (see **Chapter 5** for a more detailed review of this work).

Such evidence speaks to the most salient assumption of the prevailing economic framework, the idea of dual ‘selves.’ However the economic model also depends upon a second key assumption: It supposes that self-control is *costly*. The exertion of self-control is assumed, within the dual-self framework, to carry inherent disutility. Therefore, in order for self-control to be imposed, its expected payoffs must surpass this intrinsic cost, constituting a cost-benefit analysis in the exertion self-control (Berns et al., 2007; Fudenberg & Levine, 2006; Gul & Pesendorfer, 2001; Thaler & Shefrin, 1981).

The cost of self-control plays a pivotal role in the dual-self framework, empowering it to account for a range of important behavioral phenomena (Fudenberg & Levine, 2006; Gul & Pesendorfer, 2001; Thaler & Shefrin, 1981). However, in contrast to the assumption of dual selves, no independent experimental evidence has yet been brought to bear on this tenet of the standard model.

4.3 The strength model of self-control

Research in social psychology has recently documented failures of control that can be interpreted to be reflective of the existence of self-control costs. The dominant account in this

literature centers on the notion that self-control is dependent on a limited and depletable ‘resource’. This idea is most famously proposed in the strength model of self-control, developed by Baumeister and colleagues (Baumeister, 2002; Baumeister et al., 2007; Muraven, 2011; Muraven & Baumeister, 2000). Based on the observation that attempts at self-control are reduced if they follow periods of extended regulation (Fry, 1975; Glass, Siger, & Friedman, 1969; Knapp & Clark, 1991), the strength model posits that the exertion of self-control expends a resource, which is limited in its capacity and globally shared across operations with demands for control. Prolonged exertion of control may thus deplete this resource and render any subsequent exertion impossible, a state that has been termed ‘ego depletion’. In recent work, the proponents of the model have even argued that resource depletion is “more than a metaphor” and that glucose levels in the brain directly represent ego strength (Gailliot & Baumeister, 2007). Evidence for this idea comes from a series of studies in which blood glucose levels were decreased following a depletion manipulation, and in which participants showed no ego depletion after drinking a lemonade drink that contained glucose (Gailliot et al., 2007).

The strength model has received widespread attention in the scientific community and in the popular press (Tierney, 2011, August 17). One reason for its popularity is that it draws on common-sense plausibility and intuition. However, the most important reason for its large influence is undoubtedly the impressive array of studies that confirm the ego depletion effect, the finding that willpower strength is weakened following periods of forced exertion (for a review see Hagger, Wood, Stiff, & Chatzisarantis, 2010).

Despite its popularity in the scientific community, the strength model of control has been repeatedly challenged by new, inconsistent, findings. For example, several research groups have argued that it is highly unlikely that the brain depletes glucose resources when exerting

mental effort (Clarke & Sokoloff, 1998; Gibson, 2007; Madsen et al., 1995; Messier, 2004). It is of course possible that, even if control does not draw on a biological resource, willpower resource is represented in a more abstract fashion. Regardless of the neural implementation of such a resource, though, the strength model predicts that in a state of depletion the exertion of control is impossible. This also turns out not to be true. Ego depletion effects vanish in the face of motivation, whether it be in the form of explicit reward (Muraven & Slessareva, 2003; Stewart, Wright, Hui, & Simmons, 2009) or positive affect (Tice, Baumeister, Shmueli, & Muraven, 2007). It has even been argued that the counteracting effect of glucose drinks can be attributed to increased motivation to continue exerting effort, as triggered by neural reward signals to glucose (even when it is not swallowed, see Chambers, Bridge, & Jones, 2009; Kurzban, 2010).

These observations are worrisome for the strength model and suggest that declines in voluntary self-control may reflect reduced motivation rather than an inability to exert control. In fact, they have recently triggered a trend away from resource-based theories and towards accounts centering instead on motivation or value (Hagger et al., 2010; Inzlicht & Schmeichel, 2012; Job, Dweck, & Walton, 2010).

4.4 A role for effort costs in self-control

The research presented in the previous chapters may be useful in this recent theoretical shift in the nature of costs in self-control, especially when considering the dual-self model. In this model, self-control is implemented as an inhibitory influence of the long-term self over the short-term self. From a psychological perspective, this inhibitory function can be easily interpreted as a case of executive functioning, or the exertion of cognitive control (Aron, 2011).

Chapter 1 reviewed evidence from the DST, suggesting that such demands on executive functioning register as subjectively costly or aversive. In the DST, participants choose repeatedly between task options associated with differing levels of executive demand (Figure 4.1). While there is important variability across individuals, overall behavior in the DST displays a pattern of demand avoidance, a bias away from the choice option that carries greater executive demand (Kool et al., 2010).

These studies examined the cost of cognitive effort in settings quite different from those involved in the self-control research introduced earlier. However, it seems natural to consider whether the cost of cognitive control, as demonstrated in this research on executive function, might bear a connection with the cost of self-control hypothesized in the economic dual-self model. The hypothesized inhibitory function of self-control suggests that its exertion might be costly because it carries cognitive effort costs.

However, there is no empirical evidence that speaks directly to the prediction that failures in self-control are a result of the avoidance of cognitive control. Here, we aimed to this idea by examining the relationship between individual differences in cost processing and self-control. Specifically, individuals who are particularly sensitive to cognitive costs, as reflected in strong demand avoidance, should display relatively weak self-control.

4.5 Experiment 3: Effort avoidance and self-control

We conducted an experiment in which participants completed three tasks, in counterbalanced order. Each participant's propensity to avoid cognitive demand was measured using the DST illustrated in Figure 4.1. Self-control was measured in two ways. First, participants completed the Self-Control Scale (Tangney et al., 2004), a standardized self-report measure involving 36

questions concerning self-regulation in everyday life. Second, participants performed an ITC task closely modeled after the one used in the fMRI study of McClure et al. (2004). Based on previous findings, we anticipated significant variability across participants within each of the three tasks. More importantly, we predicted that the strength of demand avoidance in the DST would correlate inversely with both measures of self-control.




Figure 4.1. Sample event sequence from the DST. On each trial, the participant selects freely between two patterned targets, which vary in appearance and relative position every 75 trials. Selection of a target reveals an Arabic numeral. Depending on the color of the numeral (blue or yellow), the participant uses a key-press to render either a parity (odd/even) judgment or magnitude (less/greater than five) judgment. Unannounced to the participant, one target yields numerals that tend to vary in color across trials (90% of selections), imposing executive demand through the requirement to switch stimulus-response mapping. The other target yields numerals that tend to maintain a consistent color (90% of trials). Demand avoidance is quantified as the proportion of trials on which the low-demand target is selected, across the entire experiment.

Methods

Participants. Fifty students from the Princeton University (31 females, 18-24 years) participated, providing informed consent and receiving \$10 plus whatever bonus was received in the ITC task.

Materials and procedure. All participants performed an ITC task and the DST, with order counterbalanced across participants, followed by completion of the Self-Control Scale (Tangney et al., 2004). Both the ITC task and the DST were programmed using Matlab and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997).

In the ITC task, all participants were presented with an identical sequence of 100 unique choice trials. In each trial, participants chose between two monetary offers, one involving a smaller sum to be delivered immediately following the experiment, the other a larger sum to be conferred after a specified delay. To construct the offer sequence, immediate rewards were sampled from a normal distribution with mean \$8 and standard deviation of \$2 (range \$3.01-\$12.85). The delayed option was between 5% and 50% larger than the immediate option (range \$3.93-\$16.69) and was available after a period ranging between two and ten weeks, in one-week increments. Participants were (truthfully) informed that one of their selections, from a randomly chosen trial, would be awarded as an Amazon gift card at the selected time-point. Participants' intertemporal choice score was computed as the proportion of trials on which the delayed option was chosen. In an additional analysis, we fitted each participant's choice behavior to a hyperbolic discounting curve (Mazur, 1987), according to which the subjective value of the larger later option is computed as:

$$V = \frac{A}{1 + kD}$$

where V is the subjective value of the delayed reward A at delay D , and k is a parameter that determines to which degree the reward is discounted. This discounting parameter is an alternative way to characterize self-control in an intertemporal choice task. We used maximum-

likelihood estimation, using Matlab's *fminunc* function, to fit this model to choice behavior. The model implements choice as a result from a logistic transformation of the difference between the immediate value and the hyperbolically discounted delayed option. Thus, the maximum-likelihood procedure produced fits for two parameters: the discounting parameter k and the parameter determining the noise, or gain, of the logistic function.⁸

The DST was drawn without modification from Kool et al. (2010). The task was divided into eight blocks of 75 trials, each featuring a visually contrasting pair of choice targets (Figure 4.1). The position and appearance of the targets remained fixed within each block but varied across blocks, always appearing along the perimeter of a virtual circle, and separated by an angular distance of 45 degrees. Participants were told they were free to sample from either target, but that if they developed a preference they should feel free select one more than the other. Selection of a target revealed an Arabic numeral. Depending on the color of the numeral (blue or yellow), the participant used a key-press to render either a parity (odd/even) judgment or magnitude (less/greater than five) judgment. During each block, the numerals in one (*high-demand*) target switched color relative to the previous trial — requiring an effortful stimulus-response remapping — with a probability of 0.9. In the other (*low-demand*) target, colors switched with a probability of 0.1. Participants' demand-avoidance score was computed as the proportion of trials on which the low-demand cue was selected.

The Self-Control Scale (Tangney et al., 2004) is a 36-item questionnaire that measures self-regulatory behavior throughout four domains: thoughts, emotions, impulses and performance. It comprises a list of statements (e.g., 'I am always on time'), whose self-relevance

⁸ The fitting procedure produced k values for two participants that were extreme outliers (more than 93 and 85000 standard deviations from the mean of the rest of the participants), so these were not considered in the correlation analyses.

participants rate using a five-point scale. This scale has been previously found to correlate with self-reported grades, interpersonal success, and self-discipline (Duckworth & Seligman, 2005; Tangney et al., 2004)

Based on previous findings, we anticipated significant variability across participants within each of the three tasks. More importantly, we predicted that the strength of demand avoidance in the DST would correlate inversely with both measures of self-control.

Results

Consistent with earlier studies, participants as a group displayed an overall demand avoidance effect in the DST, tending to favor the low-demand option, $t(49) = 4.60, p < 0.0001$, but also showed considerable cross-participant variability (see Figure 4.2A-B). As predicted, this variability tracked individual differences on the two self-control measures. First, demand avoidance on the DST correlated negatively with scores on the Self-Control Scale, with a higher proportion of low-demand choices in the DST predicting a lower quantity of self-regulatory behavior reported in the Self-Control Scale, $r(48) = -.38, p < 0.01$ (Figure 4.2A). Second, DST performance correlated with choice behavior in the ITC task, with a greater proportion of low-demand choices in the DST predicting a smaller proportion of delayed-option ITC choices, $r(48) = -.49, p < 0.001$ (Figure 4.2B), and larger hyperbolic discounting parameters, $r(46) = 0.29, p < 0.05$.




Figure 4.2. Results of individual differences experiment. A. Relationship between demand avoidance in the DST and Self-Control Scale score. Each point corresponds to a single participant. B. Relationship between demand avoidance in the DST and the proportion of delayed-option choices in the ITC task. *Thin lines* indicate the 95% confidence interval.

Results indicated a positive correlation between our two main measures of self-control: Participants with high point-scores on the Self-Control Scale also tended more often to select the delayed option in the ITC task, $r(48) = 0.39$, $p < 0.01$. However, a mediation analysis indicated that DST performance continued to predict each self-control measure even after the other self-control measure was covaried out (Figure 4.3). Factoring out the relationship between each self-control measure and demand avoidance reduced the correlation between the two self-control measures below the threshold for statistical significance (see Figure 4.3), suggesting that the DST tapped a factor common to both.




Figure 4.3. Results from a mediation analysis. This figure shows the relationship between demand avoidance and ITC, as mediated by the Self-Control Scale measure (“self-control”). Numeric labels indicate standardized regression coefficients deriving from an analysis regressing self-control onto demand avoidance (upper left), an analysis regressing ITC onto demand avoidance (coefficient in parentheses), and an analysis regressing ITC onto both self-control (upper right) and demand avoidance (bottom). Demand avoidance explains variance in both ITC and in Self-Control Scale scores. Although Self-Control Scale score predicts ITC, this effect falls below statistical significance when demand avoidance is included as an additional regressor. Thus, demand avoidance appears to reflect a common factor underlying both Self-Control Scale responses and ITC behavior.

4.6 General Discussion

In sum, the results of this individual-differences study confirmed an inverse relationship between cognitive demand avoidance and the efficacy of self-control. This result lends considerable support to the idea that the exertion of self-control carries intrinsic subjective costs.

As discussed in the Introduction, the cost of self-control plays a pivotal role in the influential dual-self model that has emerged from economics, empowering that model to account for a wide range of behavioral effects (Fudenberg & Levine, 2006; Gul & Pesendorfer, 2001; Thaler & Shefrin, 1981). The present findings bolster the psychological plausibility of the

dual-self model, providing empirical confirmation for one of its key stipulations, that self-control carries an intrinsic cost.

The precise characterization of control costs has in fact taken two subtly different forms in economic dual-self models. In some models, a cost attaches directly to the exertion of top-down control (Fudenberg & Levine, 2006); others frame the cost of control as an opportunity cost, arising when self-control requires the short-term ‘self’ to forego tempting immediate reward (Kurzban et al., 2013). These two possibilities are difficult to differentiate empirically, since control demands will generally increase with temptation (Kool & Botvinick, 2013). However, the present results offer differential support for the idea that self-control exertion carries an inherent cost, since this view (but not the opportunity-cost alternative) provides an explanation for why self-control should correlate with demand avoidance in the DST.

By validating the notion of self-control costs, our findings also indirectly support the other key tenet of the economic model, the idea that choice is governed by two ‘selves’ with differing preferences, and that self-control reflects the ascendancy of one of these selves — the one with more patient preferences — over the other. This notion, and the dual modes of valuation that it implies, is not universal among formal models of self-control. Indeed, neuroscientific theories of self-control involving a single, fixed utility function remain widely considered⁹, especially in work on ITC (Kable & Glimcher, 2010; Kable & Glimcher, 2007). However, in contrast to the dual-self framework, such a perspective provides no obvious

⁹ The single-utility view is commonly attributed to Kable and Glimcher (2010; 2007), and their proposals can be so interpreted. However, as Hare et al. (2009) noted, the model advanced by Kable and Glimcher (2007) does not explicitly rule out top-down modulation of value representations. In fact, Kable and Glimcher (2010) explicitly left open the possibility that top-down modulation, perhaps driven by dlPFC, might play a role. Some caution is thus required in framing the debate.

entrypoint for effort costs, since it includes no distinct self-control function to which such costs might attach.

As described in the introduction, the notion of self-control costs has very recently begun to appear in psychological theories of self-control failure and ‘ego depletion.’ For many years, work in this area has been dominated by the idea that self-control draws on a limited resource — possibly glucose (Baumeister, 2002; Gailliot & Baumeister, 2007) — and that impulsive behavior arises when this resource is depleted, making the exertion of self-control impossible (Baumeister et al., 2007; Muraven, 2011). Over time, however, accumulating empirical observations have placed an increasing strain on the resource account (Hagger et al., 2010; Inzlicht & Schmeichel, 2012; Kool & Botvinick, 2013; Kurzban, 2010; Kurzban et al., 2013), contributing to an emerging trend toward motivation-based theories of self-control failure. Under this emerging perspective, self-control failures arise not from an inability to self-regulate, but from a *decision* not to do so, based on a cost-benefit analysis that takes into account the intrinsic cost of self-control (Inzlicht & Schmeichel, 2012; Kurzban, 2010).

Together with its implications for psychological and economic models, the present findings are also relevant for neuroscientific studies on self-control costs and ITC. Based on the observation that the dlPFC is activated during the exertion of self-control (Hare et al., 2009; Wagner & Heatherton, 2012), and that this region plays a crucial role for the implementation of cognitive control (Miller & Cohen, 2001), an intuitive prediction is that activation of the dlPFC is correlated with the intrinsic cost of cognitive control. **Chapter 5** will evaluate this idea, by reviewing the current neuroscientific evidence on cognitive control, demand avoidance, and self-control.

Chapter 5

Neural underpinnings of control costs

5.1 Chapter summary

This chapter discusses the neural implementation of cognitive control costs, both from the perspective of the relevant structures and of the neuromodulatory processes involved.¹⁰

- Previous research indicates that the dlPFC and the anterior cingulate cortex (ACC) play important and complementary roles in the regulation of top-down control.
- Activity in both these areas has been shown to correlate positively with different measures of cognitive effort costs.
- The exertion of self-control and representation of cognitive costs recruit a common neural substrate in dlPFC, providing further evidence that self regulation carries intrinsic effort

¹⁰ The meta-analysis in section 5.5 is reported in a published article (Kool et al., 2013).

costs.

- In the physical effort literature, the neurotransmitter dopamine appears to discount the subjective cost of physical demands.
- Experiment 4 provides some of the first evidence for a similar role of dopaminergic neuromodulation in mental effort-based decision making.

5.2 Introduction

The previous chapters have considered several pieces of evidence suggesting that demands for cognitive control factor into decision making (Kool & Botvinick, 2014; Kool et al., 2010; Kool et al., 2013). Results from these studies suggest that human decision makers have intrinsic drive to minimize demands for cognitive control, choosing lines of action with the smallest mental effort demands. **Chapters 1-3** described this tendency in experimental situations that were primarily focused on isolated effort-based decisions, whereas **Chapter 4** provided evidence that demands for cognitive control play a role in the broader field of decision making research.

The current chapter will shift focus from the behavioral implications of control costs and towards the neural computations underlying mental effort-based decision making. First, we will discuss what is known about the neural structures involved in this process, and uncover converging evidence that links self-control to effort costs through a common neural substrate in the prefrontal cortex. This research has provided several pieces of evidence about the circuitry underlying the representation of control costs. However, very little is known about the role of neuromodulation in representing control costs. Therefore, this chapter will conclude with an experiment investigating variation in demand avoidance as a function of differences in genotypes related to dopaminergic functioning.

5.3 Cognitive control and the brain

Over recent decades, it has become clear that the functioning of cognitive control depends critically on neural activity in a set of regions of the frontal and parietal cortices. Neuroimaging studies of cognitive control tasks have revealed a consistent pattern of increased activation of the dlPFC (Braver et al., 1997; Cohen et al., 1997) and the ACC (Botvinick, Braver, Barch, Carter, & Cohen, 2001; Ridderinkhof, Ullsperger, Crone, & Nieuwenhuis, 2004; Shenhav, Botvinick, & Cohen, 2013), but also in non-frontal areas such as the anterior insula and the intraparietal sulcus (Duncan, 2010; Duncan & Owen, 2000). Together, these areas are known as the “multiple-demand network” or “task-positive” network (Duncan, 2010; Fox et al., 2005). Research on the specific roles of each of these areas in cognitive control has primarily focused on the areas in the frontal lobes: the dlPFC on the lateral side of the frontal cortex and the ACC on the ventromedial wall of the brain.

The guided activation principle

The most prominent theory of function of the prefrontal cortex in cognitive control has been formulated by Miller and Cohen (2001). They propose that the role of lateral prefrontal cortex in executive functioning can best understood as ‘guiding activation’ between other, more posterior areas of the brain, establishing correct mappings between input signals and outputs in order to implement goal-directed behavior.

The elegance of the guided activation principle comes prominently to the fore in research on the Stroop task (Stroop, 1935). In the Stroop task, participants are asked to name the color in which a color word is displayed. The crucial finding, also known as the Stroop interference effect, is that RTs are larger when the word spells out a different color than the

color in which the word is displayed. The guided activation principle explains this effect in terms of automaticity: when the word and the color are incongruent, cognitive control is needed to guide activation in the color naming pathway in order to overcome the more habitual response of reading the word. Early work by Cohen, Dunbar and McClelland (1990) gave computational traction to this idea by showing that a neural network model of Stroop performance with a set of ‘prefrontal’ task set nodes, captured many aspects of performance on the task, including the basic interference effect. In addition, the model captures the finding that damage to the prefrontal cortex is associated with increased response conflict on incongruent trials (Cohen & Servan-Schreiber, 1992). Consistent with the guided activation hypothesis, activity of the dlPFC has been observed across many different cognitive control functions, including overriding of response interference in the Stroop task (Zysset, Müller, Lohmann, & von Cramon, 2001), but also working memory maintenance (Cohen et al., 1997), action representation (Badre, 2008; Botvinick, 2008), planning (D. A. Simon & Daw, 2011), and task switching (Sohn et al., 2000).

The conflict monitoring hypothesis

In the previous chapters, we have repeatedly observed that the implementation of cognitive control is intrinsically costly. When considering the neural implementation of the demand avoidance effect, the guided activation principle predicts that this should manifest as a motivation to minimize activity in the multiple demand network (Duncan, 2010).

Evidence for this prediction comes from work on the conflict monitoring model (Botvinick et al., 2001). This model is based on the connectionist model of prefrontal cortex functioning described above, but adds a complementary mechanism that adjusts the amount of top-down control signal necessary to maintain task performance. Specifically, the model adjusts

this top-down signal as a function of the current demand for cognitive control. This demand-for-control signal is operationalized as the amount of response conflict, i.e., the degree to which incompatible response tendencies are simultaneously activated. In the model, the amount of conflict between responses is monitored by a conflict monitoring unit, which projects this signal back to the control nodes (Botvinick et al., 2001). When conflict is high, the conflict monitoring unit signals the need for increased top-down activation of the correct task representation. This adjustment, in turn, leads to reduction in conflict, since increased task activation will bias activation towards the correct response representation.

The conflict monitoring model explains important behavioral findings and makes crucial predictions about the way in which the brain aims to minimize control demands. It has been repeatedly observed that performance on interference tasks, such as the Stroop task (Stroop, 1935) or the Eriksen flanker task (Eriksen & Eriksen, 1974), is increased when response conflict is high, for example immediately after an incongruent trial (Gratton, Coles, & Donchin, 1992; Ullsperger, Bylsma, & Botvinick, 2005). The computational implementation of the conflict monitoring model accounts for this finding in both RT and error rate simulations.

The conflict monitoring unit in the model has been proposed to map to the ACC in the medial prefrontal cortex. Neuroscientific data shows that the ACC shows increased activity when response conflict is high, and also predicts adjustments in control in the lateral prefrontal cortex, as measured by reduction in interference effects (Botvinick, Cohen, & Carter, 2004; Botvinick, Nystrom, Fissell, Carter, & Cohen, 1999; Kerns et al., 2004; Ridderinkhof et al., 2004). More recently, these data and other reports showing that the ACC plays a role in valuation and decision making processes (e.g., Eisenberger, Lieberman, & Williams, 2003; Kolling, Behrens, Mars, & Rushworth, 2012; Rushworth & Behrens, 2008; Yeung, Botvinick, &

Cohen, 2004) have been integrated into a proposal that the ACC computes value signals for different lines of actions in decision making situations (Shenhav et al., 2013). This signal, a combination of expected rewards and control costs, is then used to guide action selection in terms of both the amount of control allocated and where it should be applied.

In summary, neuroscientific research has identified the prefrontal cortex and the ACC as two key neural structures for the implementation of cognitive control. The ACC is proposed to monitor ongoing cognitive demands during task performance, whereas the prefrontal cortex is proposed to engage in top-down guiding of activation to implement goal-relevant task-set representations. As detailed above, the two systems are linked in a circular feedback loop, so that a detection of an increase in cognitive demands leads to increased top-down control to reduce the need for reactive control adjustments.

In the next section, I will discuss recent research that suggests that that activity in both areas has been linked to an increase in mental effort costs (Botvinick, Huffstetler, & McGuire, 2009; J. T. McGuire & Botvinick, 2010). From a theoretical standpoint, a drive to minimize activation of the ACC might be reflective of a drive towards engaging in goals that minimize the need for effortful reactive responding (Braver, 2012). An additional possibility is that activity in the ACC registers as aversive, serving as a teaching signal that attaches negative value to tasks with high demands for cognitive control and conflict monitoring (Botvinick, 2007).

5.4 Neural structures representing mental effort costs

The first neuroimaging study that directly probed the neural basis for cognitive costs focused on the phenomenon of cognitive effort discounting. This term refers to the finding that rewards are estimated to be more valuable when obtained after relatively little amounts of exerted effort

(Phillips, Walton, & Jhou, 2007; Rudebeck, Walton, Smyth, Bannerman, & Rushworth, 2006). This effect suggests that the valuation of reward is the result of a cost-benefit tradeoff, with the value of the incentive discounted by the disutility of effort.

Botvinick et al. (2009) tested the hypothesis that cognitive effort discounting would be reflected in the NAcc in the ventral striatum. This area of the brain has long been known to be responsive to reward outcomes (Montague, Dayan, & Sejnowski, 1996; Schultz, Dayan, & Montague, 1997), presumably as a result of being a prime recipient of dopaminergic neurons in the midbrain. In the experiment, participants completed blocks of the same task switching trials in the DST (Kool et al., 2010), but without the choice component of the original task. Instead, cognitive demand was manipulated across blocks. On high-demand blocks, participants were to switch tasks on every trial, whereas no switching was required on the low-demand blocks. After each block, participants either received or did not receive a reward presented as payment for the task, which was not contingent on any aspects of task performance. Results indicated that the NAcc response to the outcome was attenuated when preceded by a high-demand block, consistent with the interpretation that this part of the striatum receives effort costs and integrates it with the value of the outcome. In the initial report of this study, an region-of-interest (ROI) analysis localized the source of these effort costs to the ACC, since activity in this area during task performance predicted the size of the subsequent decrease in the NAcc (Botvinick et al., 2009). A subsequent whole brain re-analysis of the same dataset tested the relationship between task-phase activity in each voxel and the subsequent reward response in the NAcc ROI. This new analysis revealed an area in bilateral dlPFC (Figure 5.1; top row) which negatively correlated with the reward response in the NAcc, implicating a role for this area in effort discounting (Kool et al., 2013). These results suggest a neural manifestation of a mental

effort discounting effect: an increase in cognitive costs, as registered by the ACC and the dlPFC, decreased the reward response in the NAcc.

In a related study, performed by McGuire and Botvinick (2010), participants also completed blocks of the same task switching trials from the DST, again with the choice component omitted. Participants were told that the trials in each block would be drawn from one of four different algorithms or “sources”, even though the task switching stimuli were generated from the same random process on each block. After completion of each block of trials, participants rated the degree to which they wished to avoid similar task blocks in the future. Consistent with the effort discounting results, within-subject analyses revealed that these avoidance ratings were correlated with the degree to which the ACC and the dlPFC were active during task performance. The more these regions were recruited while participants performed the task switching trials, the more likely they were to indicate that they wanted to avoid similar blocks in the future. Interestingly, the dlPFC remained the only area correlated with avoidance ratings after factoring out the effects of RT and error (Figure 5.1; top row).

In summary, these neuroimaging studies link the demand avoidance effect to neural activity in the dlPFC and ACC, regions known to be critical for different aspects of cognitive control (Botvinick et al., 2001; Miller & Cohen, 2001). These results suggest that perceived effort costs are proportional to the degree with which executive control regions are activated during performance. With these findings in mind, it may be useful to return to the research domain of self-control, where activity in the frontal control regions has been found in situations when people resist immediate pleasures in favor of longer-term goals.

5.5 The neural correlates of self-control

Chapter 4 reported an individual differences experiment providing evidence that the exertion of self-control carries intrinsic effort costs. As described in more detail above, this result supports one of the two key assumptions of economic models of self-control, namely that its exertion carries disutility (Fudenberg & Levine, 2006; Thaler & Shefrin, 1981). The other assumption relates to the notion of dual-selves, the idea that the exertion of self-control involves the overriding of a ‘short-term self’ by a ‘long-term self’, in order to maximize long-term reward.

Apparent support for this dual-self view has come from a set of neuroscientific studies focusing on the dlPFC. In one such study, Hare et al. (2009) used functional magnetic resonance imaging (fMRI) to measure brain activity as dieters made decisions about which food to eat. In an initial phase, they provided evaluations about the food items, rating each on their tastiness and healthfulness. Next, participants made repeated binary choices involving all food options. The key finding was that, when participant displayed self-control, choosing foods that were consistent with their diets, fMRI revealed accompanying activation in the dlPFC (Figure 5.1; bottom row). Activation in this area was reduced when participants failed to exert self-control, choosing more items impulsively based on their taste ignoring the healthfulness.

A similar finding has come from studies of ITC, where decisions are made between tempting immediate rewards and larger delayed rewards. In two studies of ITC, participants performed an ITC task similar to the one described in Section 4.5 (McClure, Ericson, Laibson, Loewenstein, & Cohen, 2007; McClure et al., 2004). In one experiment, participants chose between Amazon gift cards with different dollar values and delays in the scanner. The results revealed two dissociable networks of brain regions that responded to different types of decisions (McClure et al., 2004). A network of brain regions including the dlPFC was activated when

participants selected the larger later gift card – behavior understood to reflect self-control (Ainslie, 1975; Frederick, Loewenstein, & O'Donoghue, 2002). Consistent with the dual-selves framework, reward-related areas, including the ventral striatum and ventromedial prefrontal cortex, showed an increase in activation when the immediate option was selected.

In a second study (McClure et al., 2007), participants repeatedly chose between a small number of squirts of juice or water administered immediately and a larger number of squirts given later during the scanning session. Again, the authors found that a set of regions, including dlPFC, increased activation when the delayed larger reward was selected¹¹ (see Figure 5.1; bottom row). Transient inactivation of dlPFC has subsequently been shown to yield more impatient behavior in ITC (Figner et al., 2010).

These studies have given rise to competing views on the precise role of dlPFC in self-control and the data are not free of inconsistency (Kable & Glimcher, 2010; Kable & Glimcher, 2007). According to one view, stemming from McClure et al. (2004), the dlPFC participates in one of two competing systems, each of which carries its own representation of choice value. Under a contrasting account stemming from Hare et al. (2009), the brain carries only a single representation of value, but one that is subject to top-down modulation by the dlPFC. Despite the important differences between these theories as accounts of neural implementation, it is important to note that they are both entirely consistent with the more abstract dual-self framework. Under both neuroscientific theories, self-control depends upon the activity of a distinct mechanism, which overrides the behavioral preferences arising from a second, more basic, system. Taken together, the available findings do appear to align with the dual-self model,

¹¹ This effect appeared in McClure et al. (2007) as a statistical trend. Nevertheless, that paper concluded that activity within a network including the dlPFC “is associated with choice, such that lesser activity ... predict[s] a greater likelihood of choosing the sooner, lesser option” (p. 5801).

by supplying evidence for an isolable neural system whose activity produces patient, far-sighted choices, and whose inactivity releases more impulsive behavior (Heatherton & Wagner, 2011), regardless of whether the override operation occurs through competition or through modulation.




Figure 5.1. dlPFC regions identified in five neuroimaging studies. The upper tier shows areas displaying effects related to self-control in multi-attribute choice (Hare et al., 2009) and ITC (McClure et al., 2007; McClure et al., 2004). The lower tier shows areas displaying effort cost effects in the form of demand avoidance (J. T. McGuire & Botvinick, 2010) and effort discounting (Botvinick et al., 2009). The images were rendered in three-dimensional space using AFNI's *Render dataset* function (Cox, 1996).

5.6 An anatomical intersection between cognitive costs and self-control

The results of Experiment 3 provided evidence for a connection between the intrinsic disutility of self-control assumed in the dual-self model and the cognitive effort avoidance effect measured in the DST. The existence of such a connection is strengthened by the fact that neuroscience research on cognitive costs and separate research on self-control have both converged on the dlPFC. Indeed, the specific areas implicated in the two domains of research bear a striking resemblance to one another, as can be seen in Figure 5.1. In this Figure, we present dlPFC regions identified in five neuroimaging studies. The top of the Figure depicts areas in dlPFC that display effects related to self-control from both the dietary choice study by Hare et al. (2009) and the ITC studies by McClure and colleagues (2007; 2004), both of which are described above. The lower part shows areas displaying effort costs effects in dlPFC in the form of demand avoidance (J. T. McGuire & Botvinick, 2010) and effort discounting (Botvinick et al., 2009).

To test for genuine overlap, we conducted a set of ROI analyses, returning to the fMRI datasets that demonstrated demand avoidance (J. T. McGuire & Botvinick, 2010) and effort-discounting (Botvinick et al., 2009) effects in dlPFC, but testing for these effects within the dlPFC regions identified in the dieter study by Hare et al. (2009), and the ITC studies by McClure and colleagues (2007; 2004). In order to perform our ROI analyses, we took the whole-brain maps of correlations between task-phase activity and subsequent avoidance ratings in the McGuire and Botvinick (2010) demand avoidance study ($n = 10$) and of correlations between task-phase activity and the subsequent reward-related response in the NAcc in the Botvinick et al. (2009) effort discounting study ($n = 23$). For each participant, we computed the average correlation in each of the three dlPFC clusters from the self-control datasets (Talarach-

transformed for the McGuire and Botvinick dataset). For both studies and all three dlPFC clusters, group-level t -tests revealed that the task-phase signal showed a significant correlation with avoidance ratings (Table 5.1, top row) and activity in the NAcc (Table 5.1, bottom row). This anatomical intersection between self-control and effort-cost effects suggests a functional connection between the two, providing further support for the finding that self-control itself carries intrinsic effort costs.

Table 5.1. *Results from a Set of ROI Analyses*

Dataset	Effect	McClure et al., 2004		McClure et al., 2007		Hare et al., 2009	
		ITC – Monetary rewards		ITC – Primary rewards		Self control – dieter study	
		t (df)	p	t (df)	p	t (df)	p
McGuire & Botvinick, 2010	Demand avoidance	1.90 (9)	< 0.05	3.03 (9)	< 0.01	3.11 (9)	< 0.05
Botvinick et al., 2009	Effort discounting	-2.68 (22)	< 0.05	-2.99 (22)	< 0.005	-1.81 (22)	< 0.05

Note. The dlPFC regions-of-interest (ROI) were drawn from fMRI studies on self-control in dietary choice and in ITC (Hare et al., 2009; McClure et al., 2007; McClure et al., 2004). Using data from two previous studies (Botvinick et al., 2009; J. T. McGuire & Botvinick, 2010), we tested whether task-induced activity in these regions predicted a reduction in reward-receipt responses in ventral striatum, i.e. effort discounting (Botvinick et al., 2009), and such activity predicted subsequent task-avoidance behavior (J. T. McGuire & Botvinick, 2010). All tests yielded significant effects, based on one-tailed t -tests.

This result adds to the behavioral finding from Experiment 3, implicating a role for cognitive control costs, as registered by the dlPFC, in self-control and ITC. Despite the findings reviewed earlier, some important negative results have left room for uncertainty about the dual-self model of self-control, especially in the case of ITC (Kable & Glimcher, 2010; Kable & Glimcher, 2007). Our results indirectly support the relevance of dlPFC, by providing evidence that effort avoidance and self-control and ITC are associated with costs, costs that are presumably indexed by a shared neural substrate in the dlPFC (Botvinick et al., 2009; J. T. McGuire & Botvinick, 2010). At a broader level, the present findings establish a new bridge between neuroscientific research on self-control and parallel research on effort costs and

demand avoidance, prompting further investigations into the relationship between these two domains and their intersection in neural structures.

5.7 Dopamine and effort

Throughout the previous sections, it has become clear there is already a fair amount of knowledge about which neural structures are important for the representation of mental effort costs. In short, increased activity in areas that are important for the direct implementation of control and the monitoring of cognitive demand, such as the dlPFC and the ACC, is positively correlated with the size of effort costs. These neural representations of effort costs in turn influence evaluation and decision making processes, presumably through value representation in the NAcc and ACC.

In contrast to the evidence pertaining the neural structures involved in cognitive effort-based decision making, very little is known about the role of neuromodulation in this process. However, the neural structures identified above are all projection sites of several important neurotransmitter pathways, including the dopamine (Schultz, 1999), norepinephrine (Aston-Jones, Foote, & Bloom, 1984), and serotonin (Geyer, Puerto, Dawsey, Knapp, & Bullard, 1976) systems. Since each of these neurotransmitters have been identified as key players in the regulation of cognitive functioning and decision making (Boureau & Dayan, 2010; Nieuwenhuis, Aston-Jones, & Cohen, 2005; Wunderlich, Smittenaar, & Dolan, 2012), it is important to understand how neuromodulation influences the computational processes underlying cognitive costs.

Here, we will start this line of investigation by exploring the role of dopamine in the avoidance of cognitive control demands. The rationale for focusing on this neurotransmitter was

twofold. First, research on physical effort-based decision making has discovered that dopamine plays an important role in this process. Second, tonic levels of dopamine have been consistently linked to performance on tasks that require executive functioning.

Dopamine and physical effort

Over the last decade, dopamine has been of prime focus in research on reward-based decision making. This neurotransmitter is most famously known for its role in signaling rewarding experiences, especially because of its involvement in encoding reward prediction errors, the difference between expected and obtained reward, in the NAcc (Montague et al., 1996; Schultz et al., 1997), and other forms of internal reward signals (Schultz, 2002; Spanagel & Weiss, 1999). It turns out, however, that dopamine has been interpreted to play a different role in reward-based decisions involving physical effort.

Salamone, Cousins, and Bucher (1994), for example, observed that rats with NAcc dopamine depletions showed a decreased willingness to spend physical effort to obtain a reward. In this experiment, rats at baseline chose to exert effort to climb over barrier to get a large reward, even if they could have gotten a small reward without exertional demands. However, after 6-OHDA was administered to their NAcc, depleting dopamine levels, the rats changed their preference to the small reward over the larger more effortful option (Salamone et al., 1994). Importantly, this preference reversal was only observed in the presence of effort costs for the high-reward arm; when effort was equated, dopamine did not affect the decision. This effect has been replicated several times with several different paradigms. For example, rats become less willing to increase effortful lever pressing in response to an increase in ratio requirements (the number of times the animal has to press a lever to receive some reward) when their NAcc

dopamine levels are depleted (Aberman & Salamone, 1999; Salamone, Wisniecki, Carlson, & Correa, 2001). Subsequent research has shown that the depleted rats' unwillingness to exert effort can not be explained by altered time discounting parameters (Floresco, Tse, & Ghods-Sharifi, 2008), and that their preference reverts to the high-reward option if the low-effort arm contains no reward (Cousins, Atherton, & Turner, 1996). Furthermore, dopamine agonists, such as D-amphetamine, biases rats to the high-reward/high-effort option (Bardgett, Depenbrock, Downs, Points, & Green, 2009; Cagniard, Balsam, Brunner, & Zhuang, 2005). Finally, in monkeys, dopamine neurons in the substantia nigra, which project dopamine to the NAcc, signal increased effort costs by reducing tonic activity (Pasquereau & Turner, 2013). In these animal studies, physical effort costs have most often been found to be tracked by dopamine levels in the NAcc (Denk et al., 2004; Salamone, Correa, Mingote, Weber, & Farrar, 2006; Salamone, Correa, & Weber, 2003), but a similar effect has been reported for the cingulate cortex (Schweimer & Hauber, 2006; Walton et al., 2006).

Much less work has been done on the neuromodulatory influences on human physical effort-based decision making, but a few studies should be mentioned. In one fMRI study, Croxson, Walton, O'Reilly, Behrens, and Rushworth (2009) examined how the brain integrates physical effort costs with monetary benefits into a composite net value. Similar to the effort discounting study discussed earlier, increased effort costs yielded reductions in reward-related signal in the (dopamine-rich) ventral striatum. Complementary evidence for dopaminergic involvement in this process comes from a positron emission tomography (PET) study by Treadway et al. (2012), who found that increased dopamine responsivity was correlated with a willingness to expend greater effort for large rewards. A pharmacological study from the same group showed that a dopamine agonist increased willingness to exert physical effort (Wardle,

Treadway, Mayo, Zald, & de Wit, 2011).

In summary, the results show that the level of dopamine in the NAcc inversely correlates with sensitivity to physical effort requirements, i.e., that increased dopamine biases cost-benefit analyses by discounting physical effort costs (Phillips et al., 2007). The current aim is to explore whether dopamine has a similar influence on cognitive effort-based decision making.

Dopamine and cognitive control

This possibility is supported by evidence linking performance on cognitive control tasks to functioning of the dopaminergic system in the striatum.

Several studies have suggested that increased striatal dopamine is associated with enhanced performance on cognitive control tasks. For example, a series of PET studies has suggested that people with increased presynaptic synthesis in the striatum show better performance on working memory span tasks (Cools, Gibbs, Miyakawa, Jagust, & D'Esposito, 2008; Landau, Lal, O'Neil, Baker, & Jagust, 2009). The dopamine agonist amphetamine decreases switch costs in human participants, especially in those with increased availability of striatal D2 receptors (Samanez-Larkin et al., 2013). Consistent with these findings, global dopamine reduction, by depleting a dopamine precursor, has been found to decrease performance on mentally demanding working memory tasks (de Wit et al., 2012), whereas administration of the dopamine agonist bromocriptine resulted in the reverse pattern (Cools, Sheridan, Jacobs, & D'Esposito, 2007).

A common approach for studying the effects of dopamine on cognitive control is by investigating the effect of genetic variations in dopaminergic functioning. For instance, a genetic polymorphism related to reduced striatal dopamine has been linked to worse performance on

the Wisconsin Card Sorting Task (Rodriguez-Jimenez et al., 2006), a task that requires flexible task switching and performance monitoring. Bolton et al. (2010) found that genetic variations associated with higher striatal dopamine were correlated with increased fluid intelligence. In a study that investigated variations in another striatal gene, Aarts et al. (2010) found that increased dopamine functioning led to reduced switch costs in the face of reward, consistent with a role for dopamine in mental-effort based decision making.

It is important to note that the role of NAcc dopamine in physical effort-based decision making has been cast in very different term than its influence on cognitive performance. Dopamine seems to play a motivational role in the physical effort literature, influencing animals' willingness to exert effort, whereas it has been interpreted to play a role in setting hard capacity limits for performance on executive functioning tasks (cf. Aarts et al., 2010). An interesting alternative possibility is that dopamine plays a similar part in cognitive performance as it does in physical effort-based decision making, discounting the subjective cost of cognitive control and thereby yielding more effortful behavior. The results from the physical effort literature are consistent with this prediction. Furthermore, indirect support for this hypothesis comes from the cognitive effort discounting study discussed in Section 5.5 (Botvinick et al., 2009), where control costs were reflected in reward signals in the NAcc. However, to our best knowledge, so far there has not been a direct demonstration of a role for dopamine in the representation of effort cost. The hypothesis that striatal dopamine discounts mental effort remains in need of a direct empirical test.

5.8 Experiment 4: Dopaminergic genes predict effort avoidance

The approach in the current experiment was straightforward. We collected DNA from participants who completed the DST and tested for a relationship between demand avoidance and genetic profile, focusing on a set of seven genetic polymorphisms known to affect dopaminergic function. Following the physical effort literature, our main focus was on five polymorphisms associated with dopamine function in the striatum.

Striatal polymorphisms

We collected genotyping data for five genetic variations related to dopaminergic functioning in the striatum.

First, we assessed a single nucleotide polymorphism (SNP) in the PPP1R1B gene coding for DARPP-32 (rs909074). This intracellular protein has been shown to modulate dopamine receptor D₁-dependent synaptic plasticity in the striatum, with the number of T alleles indicative of enhanced function (Meyer-Lindenberg et al., 2007).

Next, we investigated three SNPs in the DRD2 gene, which codes for the dopamine receptor D₂. The C957T SNP (rs6277) in this gene affects postsynaptic D₂ receptor density in the striatum, with the number of T alleles associated with increased availability (Hirvonen et al., 2004). The other SNP, rs1076560, has been shown to decrease expression of the presynaptic D₂ receptors for carriers of the T allele (Zhang et al., 2007). For the last DRD2 SNP, rs12364283, carriers of the C allele have increased expression of D2 mRNA (Zhang et al., 2007).

Finally, we also considered individual differences in the 40 base pair variable number tandem repeats polymorphism in the 3' untranslated region of the dopamine active transporter (DAT) gene (DAT1 or SLC6A3; rs28363170). This gene encodes for a membrane protein that

clears dopamine out of the synaptic space in the striatum. The most common alleles for this polymorphism are the 9-repeat and 10-repeat versions, with the latter associated with increased expression of DAT and therefore lower levels of striatal synaptic dopamine (Heinz et al., 2000).

Prefrontal polymorphisms

In addition, we also collected genotype data on two polymorphisms associated with dopamine function in the prefrontal cortex, given the crucial role that frontal regions play in cognitive control.

The first of these was the functional Val158Met polymorphism within the COMT gene (rs4680). This gene encodes for the enzyme catechyl-O-methyl transferase, an enzyme that degrades dopamine, primarily in the prefrontal cortex. Specifically, the Met allele in this gene is associated with lower activity of the COMT enzyme, and therefore higher levels of dopamine in the prefrontal cortex (Meyer-Lindenberg et al., 2005).

Second, we collected genotyping data for C-526T, a SNP in the DRD4 gene (rs1800955; Okuyama, Ishiguro, Toru, & Arinami, 1999) for which it has been shown that the T allele is associated with a reduction in the expression of the D4 receptor in the prefrontal cortex, up to 40% compared with the C allele.

Methods

Participants. Ninety-seven participants from the Princeton University community participated in the study. Twenty-six of these participants were a part of the sample reported by in **Chapter 4**. Since DNA collection was not part of the original study, these participants were contacted after completion with a request for saliva sample collection (51% of the total sample

provided a sample). The other seventy-one participants were collected as a part of this experiment and all provided saliva samples at the start of the experiment. The DST was programmed using Matlab and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997).

Materials and procedure. All participants completed the DST described in **Chapters 1 and 4**. The task was divided into blocks of 75 trials, each featuring a visually contrasting pair of choice targets. The first group of participants ($n = 26$) completed 8 blocks of trials, and the second group completed 4 blocks of trials ($n = 71$). The position and appearance of the targets remained fixed within each block but varied across blocks, always appearing along the perimeter of a virtual circle, and separated by an angular distance of 45 degrees. Participants were told they were free to sample from either target, but that if they developed a preference they should feel free select one more than the other. Selection of a target revealed an Arabic numeral. Depending on the color of the numeral (blue or yellow), the participant used a key-press to render either a parity (odd/even) judgment or magnitude (less/greater than five) judgment. During each block, the numerals in the high-demand target switched color relative to the previous trial — requiring an effortful stimulus-response remapping — with a probability of 0.9. In the low-demand target, colors switched with a probability of 0.1. Participants’ demand-avoidance score was computed as the proportion of trials on which the low-demand cue was selected.

Genotyping method. We collected DNA was extracted from saliva samples collected with Oragene OG-500 DNA Collection Kits (DNA Genotek), as per manufacturer’s instructions. Each participant’s sample was analyzed at the DNA Sequencing Facility at the Department of Genetics of the University of Pennsylvania. The regions of genomic DNA that contained the genetic variations of interest were amplified using Applied Biosystems TaqMan primers and probe pairs. Analyses were performed using Applied Biosystems TaqMan Technology.

Fluorescence was analyzed with the ABI GeneMapper software. All genotypes were scored by comparison to sequence-verified standards.

Analysis. For all genotype analyses, we compared the largest homozygote group to the combination of the other homozygote group and the heterozygotes. This produced the following groups: DARPP-32 – T/T, T/C+C/C, C957T – C/C, T/C+T/T, rs1076560 – C/C, T/C+T/T, rs12364283 – T/T, C/T, DAT1 – 10/10, 9/9+(6,8,9)/10, COMT – Val/Val, Met/Val+Met/Met, C-521T – T/T, C/C+C/T. For these genotypes, the largest group comprised at most 67% and at least 53% of all participants. However, for the DRD2 rs12364283 SNP, the largest group (T/T) comprised 89% of all participants, so we did not compare behavior for this SNP. All six remaining SNPs were in Hardy–Weinberg equilibrium ($p > 0.38$). For each of the genotypes, we designated one group as associated with increased dopaminergic functioning and the other with decreased dopaminergic functioning.

For each genotype, we compared low-demand selection rate between the two groups using standard two-sample t -tests. To correct for multiple comparisons, we also ran two univariate ANOVAs with gene groups as categorical variables, separately for the striatal and prefrontal genes. Finally, we ran a regression analysis, predicting demand avoidance from two continuous predictors, which were computed as the sum of the number of alleles associated with high dopamine levels, again separately for the striatal and prefrontal genes.

Results

Demand avoidance. Average low-demand selection rate for both groups of participants is depicted in Figure 5.2. Participants displayed a significant demand avoidance effect, both in the subset of participants from Experiment 3 (average low-demand choice = 61%), $t(25) = 2.68, p <$

0.05, and the larger new group (average low-demand choice = 63%), $t(70) = 6.90, p < 0.001$. There was no significant difference in demand avoidance and accuracy between the two groups, $t < 1$, but the smaller group had significantly smaller RTs on correctly answered trials, $t(95) = 2.41, p < 0.05$. This difference probably reflected a training effect, since RTs did not differ when comparing performance across the first four blocks, $t(95) = 1.58, p = 0.12$.




Figure 5.2. Distribution of participants' low-demand selection rate in Experiment 4 ($n = 97$).

Striatal genes. The results for the analysis of the effect of striatal genotypes on demand avoidance are depicted in Figure 5.3. There was no significant difference in the proportion low-demand choice between the groups for DARPP-32, $t < 1$, and DAT1, $t < 1$, but there was a significant effect for C957T, $t(95) = 3.06, p < 0.01$, and a marginal effect for rs1076560, $t(95) = 1.78, p = 0.08$. In a four-way ANOVA with the different groups for the striatal genotypes as factors and the proportion low-demand choice as the independent measure, the effect for

C957T remained significant, $F(1,81) = 5.78$, $p < 0.05$, with other main effects or interaction effects not reaching statistical significance.




Figure 5.3. Proportion low-demand choice on the DST is affected by the DRD2 C957T genotype and marginally by the DRD2 rs1076560 genotype (uncorrected). Error bars indicate across-subject SEM. ** $p < 0.01$, † $p < 0.08$

To probe the relationship between the C957T SNP and demand avoidance, we examined low-demand selection rate for all three different groups. As can be seen in Figure 5.4, for this genotype the presence of two T alleles was associated with increased demand avoidance.




Figure 5.4. The C957T genotype in the DRD2 gene predicts low-demand selection rate in the DST. Error bars indicate across-subject SEM.

Prefrontal genes. The two SNPs related to prefrontal dopaminergic functioning did not, by themselves, influence low-demand selection rate, COMT, $t < 1$, C-521T, $t(94) = 1.45$, $p = 0.15$ (see Figure 5.5). A subsequent two-way ANOVA showed no significant demand avoidance effect for COMT, $F < 1$, C-521T, $F(1,93) = 2.67$, $p = 0.10$, or their interaction, $F < 1$.




Figure 5.5. Prefrontal gene effects on demand avoidance. By themselves, the prefrontal genes showed no influence on behavior. Error bars indicate across-subject SEM.

Composite scores. Even though we only found one significant difference in demand avoidance across our genotypes, it should be noted that all striatal genes showed the same numerical effect, with increased dopamine functioning related to lower demand avoidance. Similarly, both prefrontal SNPs had a similar numerical trend in their effect, with increased dopamine related increased demand avoidance. To further probe this possibility, we increased power in our genotyping analyses by computing *composite gene dose* scores separately for the polymorphisms related to dopaminergic functioning in striatum and prefrontal cortex. These composite scores were computed as the sum of alleles associated with increased dopaminergic function. Specifically, we computed the striatum dopamine score as the sum of the number of T alleles for DARPP-32 and C957T, the number of C alleles for rs1076560 and the number of 9-repeats for DAT1. The prefrontal dopamine score was computed as the sum of the number of Met alleles for COMT and the number of C alleles for DRD4.

We submitted these scores (centered around their mean), and their interaction, as predictors to a multiple linear regression with the proportion low-demand choice as the dependent measure (Table 5.2). This regression analysis again confirmed that increased dopaminergic functioning of the striatum was associated with decreased demand avoidance ($p < 0.01$). Moreover, we now also found evidence for an inverse relationship for prefrontal dopaminergic functioning: increased prefrontal dopaminergic functioning was associated with increased demand avoidance ($p < 0.05$). The interaction term remained insignificant ($p = 0.86$).

Table 5.2. *Results of the Multiple Regression Investigating the Relationship of Genotypes Associated with Dopaminergic Functioning in the Striatum and Prefrontal Cortex on Demand Avoidance*

Predictor	Estimate (SEM)	<i>p</i>
Intercept	.625	-
Striatum score	-.028 (.010)	< 0.01
Prefrontal score	.039 (.019)	< 0.05
Interaction	-.003 (.012)	0.86

Discussion

The study establishes the first direct evidence for a role for dopamine in effort-based decision making. The results show that participants with increased functioning of dopaminergic system in the striatum, as indexed by individual differences in a set of relevant genotypes, have reduced aversion for cognitive demand. In contrast, enhanced dopamine levels in the prefrontal cortex were associated with increased demand avoidance.

The striatal results mirror the findings on neuromodulation in physical effort-based decision making (for a review see Salamone & Correa, 2012). Here, dopamine is proposed to affect cost-benefit tradeoffs by decreasing animals' sensitivity to physical effort. The link between these results and the current investigation is clear. We replicated our basic demand avoidance effect, but found that the strength of this bias, a proxy of the cost of mental effort, was significantly reduced by genotypes related to increased dopaminergic functioning, mirroring the effect observed in the physical effort literature.

Since we were primarily focused on following the predictions from the physical effort literature, which only pertain to NAcc dopamine, we did not have a strong hypothesis about the direction of prefrontal dopaminergic modulation on cognitive costs. Our results indicated that

prefrontal dopaminergic functioning was positively related to demand avoidance. Interestingly, recent genotyping studies on cognitive control have found that increased prefrontal dopaminergic functioning is associated with decreased performance on task-switching paradigms (Colzato, Waszak, Nieuwenhuis, Posthuma, & Hommel, 2010), decreased patience in an ITC experiment (Boettiger et al., 2007) and worse performance on a “frontal functioning” task (Mitaki et al., 2012), perhaps hinting at a reverse relationship between dopamine and performance on effortful tasks in the frontal cortex when compared to the striatum.

5.9 General Discussion

This chapter reviewed the current knowledge on the neural correlates of mental effort costs. Previous neuroimaging studies have revealed that the dlPFC and ACC perform important and complementary functions for cognitive control. Activity in both these areas correlates with a measure of mental effort in the form of cognitive effort discounting (Botvinick et al., 2009). Self-reported task avoidance correlated with both areas as well, but only with dlPFC in a behavior-independent fashion (J. T. McGuire & Botvinick, 2010). A meta-analysis showed that this same frontal region is recruited during the exertion of self-control, providing further evidence for the idea that self-regulation carries intrinsic effort costs. Finally, Experiment 4 showed that individual differences in genotypes associated with dopaminergic functioning predicted the cost of mental effort, mirroring results from the physical effort literature.

At a broader level, the findings from the meta-analysis establish a new bridge between neuroscientific research on self-control and parallel research on effort costs and demand avoidance, prompting further investigations into the relationship between these two domains. For example, future work could employ fMRI or transcranial magnetic stimulation methods to

more directly test for the role of dlPFC in representing effort costs during self-control. One possibility would be to measure individual differences in dlPFC sensitivity to cognitive effort and predict individual differences in behavior and prefrontal activity during self-control (and vice versa). In addition, the current results suggest that other forms of decision making that depend on activity in dlPFC may show similar sensitivities to individual differences in effort costs. For example, one might predict that an aversion to cognitive effort predicts less utilitarian moral reasoning (Greene, Nystrom, Engell, Darley, & Cohen, 2004) and increased reliance on habit or model-free reinforcement learning (Gläscher, Daw, Dayan, & O'Doherty, 2010), since these cognitive functions are dependent on computation implemented by the dlPFC.

Cognitive demand is typically associated with activation across a set of regions throughout the frontal and parietal lobes, including the dlPFC and ACC, but also the anterior insula and the intraparietal sulcus (Duncan, 2010; Duncan & Owen, 2000). The imaging studies of effort costs and self-control, however, showed that the ACC but primarily the dlPFC registered effort costs, and did not provide evidence for a role of the insula and the parietal cortex. This dichotomy is interesting and worthy of empirical attention. Why do people aim to minimize activity in the dlPFC, but not in other parts of the multiple demand network? One possibility, is that the cost of cognitive control stems from a drive to minimize cross-talk between different task representations and thus to limit the number of concurrent control demanding tasks (Feng, Schwemmer, Gershman, & Cohen, 2014). Task set representation depends critically on the frontal cortex, which projects to more posterior, associative, regions (Miller & Cohen, 2001) and the ACC, which has been shown to receive many different inputs relating to costs and benefits (Shenhav et al., 2013), which would explain why they play a prime part in representing effort costs. To shed more light on this issue, future research should address

the functional specialization of the other regions in the multiple demand network, possibly pulling out neural characteristics that reveal which task demands would correlate with activity in these areas.

Experiment 4 revealed a role for striatal and prefrontal dopaminergic functioning in the avoidance of demands for mental effort, especially for a polymorphism relating to post-synaptic D2 receptors. Specifically, increased post-synaptic density of these receptors was associated with reduced demand avoidance. The C957T genotype has been related to another type of avoidance learning, namely the avoidance of choice options with low-reward probability. In a probabilistic reinforcement learning task, Frank and Hutchison (2009) found that T carriers (associated with increased density of D2 receptors) displayed more optimal choice behavior when the choice pair included an option that was statistically inferior (see also Doll, Hutchison, & Frank, 2011). Can we reconcile these findings? One possibility is that learning of negative reward associations is dependent on working of the prefrontal cortex, so that increased effort avoidance reduces optimal behavior in this probabilistic reinforcement learning task. Frank and Hutchison (2009) also report that people with increased availability of D1 receptors, assessed through the DARPP-32 genotype, were more optimal when facing a choice with a statistically superior option. This pattern of results suggests that dopamine might discount effort costs associated with dissociable forms of cognitive control, related to different modes of learning (but see Baker, Stockwell, & Holroyd, 2013).

The results from Experiment 4 also suggest that dopamine levels in prefrontal cortex and striatum have an opposite relationship with the cost of cognitive activity. In light of this finding, it is interesting to consider that dopamine levels in the prefrontal cortex and NAcc seem to be negatively correlated (Del Arco & Mora, 2008; Tzschentke, 2001; Wilkinson, 1997). This

observation may provide useful for understanding the inverse relationship between prefrontal and striatal dopamine levels on demand avoidance. Another set of findings suggests that this relationship is more complex than a simple inverse association. Cools and D'Esposito (2011), for example, argue that there exists an inverted-U-shaped pattern between prefrontal dopamine and cognitive control, with optimal performance at moderate tonic dopamine levels (Cools & Robbins, 2004; Vijayraghavan, Wang, Birnbaum, Williams, & Arnsten, 2007). To gain a better hold on the implications of this insight on effort costs, future research should employ a comprehensive design using genotyping, fMRI, and pharmacological interventions. If this inverted-U relationship speaks to cost-benefit analyses involving effort costs, the administration of dopamine agonist bromocriptine should have differing effects on demand avoidance, based on baseline dopamine function as assessed by variations in genomic DNA.

Chapter 6

Conclusions

6.1 Chapter summary

The final chapter will summarize the work reported in this dissertation and discuss implications for future research.¹²

- The current work has explored the implications of mental effort costs for decision making, with a specific focus on (1) its role in cost/benefit analyses, (2) its underlying utility function, (3) its relationship to self-control, and (4) its neuromodulatory underpinnings.
- Our labor/leisure model may provide useful in understanding why there is a cost associated with the exertion of cognitive effort
- The intrinsic-cost framework puts several constraints for formulations of precise neural theories of action selection.

¹² Sections of this chapter have previously been reported as discussions in published articles and a commentary (Kool & Botvinick, 2013; Kool & Botvinick, 2014; Kool et al., 2010; Kool et al., 2013).

- The notion of control costs has the potential to contribute to knowledge in the clinical domain.

6.2 Summary of the dissertation

Behavioral and economic theories have long maintained that actions are chosen so as to minimize demands for exertion or work, a principle sometimes referred to as the *law of less work*. The data supporting this idea pertained almost entirely to demands for physical effort. However, the same minimization principle has often been assumed also to apply to cognitive demand. **Chapter 1** and **Chapter 2** reported two new behavioral paradigms, validating the notion of cognitive control costs, ruling out the possibility that this tendency reflects a maximization of the rate of goal achievement. Consistent with a motivational account, avoidance of demand displayed sensitivity to task incentives.

The findings from these chapters investigate effort costs in categorical all-or-none choices between cognitive work and cognitive rest. Daily life frequently offers a choice between these activities, but here the critical decision lies not in *whether* to work, but rather *how much time* to allocate to either activity. In **Chapter 3**, we identified a central principle, derived from the field of labor economics, that guides such decisions. Results from four economic-choice experiments indicated that the motivation underlying cognitive labor/leisure decision making is to strike an optimal balance between income and leisure, as given by a joint utility function.

Our ability to exert cognitive control, arguably one of the most powerful capacities of the human mind, pervades every day life. Since cognitive costs are tied to the working of executive functioning, mental effort avoidance should also affect decisions studied in the broader field of psychological science. **Chapter 4** focused on the capacity for self-control, our ability to inhibit

obtaining immediate pleasures in favor of longer-term goals, since it has long been understood to require effortful executive functioning. A behavioral experiment showed that the degree of demand avoidance correlated with two measures of self-control. This finding constitutes the first independent evidence for an intrinsic disutility for self-control costs in the form of mental effort costs.

Finally, in **Chapter 5**, we turned to the neural correlates underlying demand avoidance. Previous research suggests that effort costs are primarily registered by an area in lateral prefrontal cortex, and the anterior cingulate cortex, with the former being tied to the a behavior-independent subjective appraisal of cognitive costs. Interestingly, this region in dlPFC turns out to be a common neural substrate underlying cognitive costs and the exertion of self-control, further validating the notion that the cost of self-control might be driven by effort avoidance. Research has identified the key neural structures relevant for the computation of control costs, yet our knowledge of the neuromodulatory underpinnings in this process is presently only inchoate. Results from a genotyping study revealed a role for striatal and prefrontal dopamine in the discounting of mental effort costs, analogous to this neurotransmitter's function in physical effort-based cost/benefit analyses.

The notion of mental effort costs has long been invoked as an explanation for errors and suboptimal behavior in human decision making. However, cognitive scientists have only recent started to employ tasks to measure this tendency (Kool et al., 2010; Schouppe, Ridderinkhof, Verguts, & Notebaert, 2014; Westbrook et al., 2013), starting with some of the work described here. In this dissertation, I have discussed a number of key questions that come up in this new area of research. Of course, the field is very much in its infancy, and there still many open

questions to be addressed. The remainder of this chapter will cover some of the most promising issues for future investigation.

6.3 The purpose of a cost of cognitive control

The work in this dissertation, as a descriptive theory, does not center on (or depend upon) a normative argument. However, the question of why cognitive effort carries a cost is very deserving of empirical attention. Over the last decades, there have been several theories that address why control is prone to failure. A review of these theories suggests that all of them explain the failure as an opportunity cost, the loss of a potential gain from alternatives when an action is chosen. However, the nature of this opportunity cost varies per proposal.

In the strength model of self-control (Baumeister et al., 2007), for example, control poses an opportunity cost in the sense that expended willpower resource could have been allocated to alternative actions. In fact, the model was developed to explain why voluntary cognitive effort tends to decline after bouts of forced cognitive exertion (Baumeister, 2002), i.e., which would be a direct observation of how this opportunity cost extends over time. According to some researchers, this effect reflects a depletion of blood glucose, a metabolic resource that is hypothesized to be necessary for the exertion of cognitive control (Gailliot & Baumeister, 2007; Gailliot et al., 2007). Therefore, refraining from cognitive control is rational, because it allows for preservation or replenishment of this metabolic resource, and therefore a minimization of the incurred opportunity costs.

According to Kurzban and colleagues (2013), the perception of mental effort is directly related to another form of opportunity costs, namely the foregone reward from other options. In their model, the cost of control is directly attached to the presence of task options that are

more valuable than the current task, or the felt output of ongoing cost/benefit computations that monitor for more profitable lines of actions. Even though this framework relies intrinsically on the concept of opportunity costs, it is significantly different from the ego depletion framework. First, this model does not rely on the existence of a depletable resource. Empirical evidence for the role of glucose (or any other metabolic resource) in the depletion effect is, at best, mixed (Kurzban, 2010), undermining its force as a mechanistic argument. Second, the model puts motivation in the center of focus, by making perceived effort, and resulting allocation decisions, a result of cost-benefit analyses, weighing the opportunity cost against the current incentive structure. Recently, a trend has emerged in research on self-control, moving away from resource-based theories and toward accounts centering instead on motivation or value (Inzlicht & Schmeichel, 2012; Job et al., 2010). Evidence for such tradeoffs in depletion comes a range of findings that show that depletion effects disappear in the face of increased incentives (Muraven & Slessareva, 2003). This has led even proponents of the ego-depletion theory to acknowledge the need for an additional layer, at which effort expenditure is shaped by motivational factors (Hagger et al., 2010).

In contrast to the ego depletion and opportunity cost models, we favor the simpler hypothesis that effort reflects an intrinsic cost attaching directly to the exertion of cognitive control. This cost can be also be seen a instantiation of an opportunity cost, relating its cost directly to forgone mental leisure. The labor/leisure model suggests an alternative view of the depletion effect, according to which it arises from a particular form of cost–benefit analysis: In the context of prolonged, obligatory mental effort, the marginal cost of further effort is elevated, leading in some cases to a subsequent withdrawal from cognitively challenging activity. By

characterizing the depletion effect in this way, the present work situates this phenomenon within a framework for understanding labor/leisure decisions.

Even though our and Kurzban's frameworks can be cast in terms of opportunity costs, it should be noted that the nature of these costs is fundamentally different. Kurzban et al. (2013) propose that the opportunity cost is directly dependent on the next-best use to which the executive functioning system can be deployed, whereas the opportunity cost in the intrinsic-cost framework is dependent on the degree to which cognitive is, and has been, exerted. There are at least two issues that favor the intrinsic-cost framework over Kurzban and colleagues' proposal. The first problem involves the question of sufficiency: One can think of many situations that feature salient opportunity costs, but that seem unlikely to involve any sense of subjective effort. Imagine, for example, sitting in a restaurant with a friend who is enjoying a dish you wish you had ordered. Of course, this scenario involves no obvious role for effort, despite the opportunity costs. Second, according to the opportunity cost model, ego depletion occurs because, over time, the expected utilities of alternative mental activities rise through learning, ultimately triggering a shift in focus. This account relies on the unfounded assumption that initial value estimates for alternative activities will generally display a negative bias, and implausibly predicts that depletion effects should be isolated to novel task circumstances.

A view of effort based on the intrinsic cost of cognitive control appears to avoid some of the difficulties of the opportunity cost model. The restaurant scenario introduced above is no longer problematic, since it features no demand for cognitive control, and therefore predicts no sense of effort. The intrinsic-cost perspective also fares better with depletion effects, as the context-sensitivity of control costs predicts that the sustained exertion of control will trigger eventual cost-driven disengagement, even in contexts involving no learning.

Normative reasons for an intrinsic cost of control

However, as noted before, the intrinsic-cost view does not speak to the question of why mental effort should be intrinsically costly. Having said this, it is possible, if one is willing to speculate, to fit the present descriptive theory within a normative framework. From a computational point of view, there are at least three reasons why the occasional withdrawal of control may be rational, and why it might thus be rational for decision-making to place a cost on control.

First, a cost of control would lead cognition towards tasks that require less executive functioning, and therefore would steer to more efficient and less error-prone processing (Botvinick, 2007). For example, people select strategies that are consistent with their cognitive strengths (MacLeod, Hunt, & Mathews, 1978; Mathews, Hunt, & MacLeod, 1980; Reichle, Carpenter, & Just, 2000), consistent with the idea that effort avoidance leads to efficient task performance. Related to this idea, it has recently been proposed that people select efficient task representations by aiming to minimize the need for cross-talk among processing pathways (Feng et al., 2014). One possibility is that the cost of control plays a direct role in this process by limiting the deployment of control signals as to minimize such cross-talk between different task representations.

Second, a wealth of research indicates that controlled information processing is highly capacity-limited (Engle & Kane, 2003; Luck & Vogel, 1997). Thus, a motivational bias in favor of automatic processing would have the effect of reserving this limited computational bandwidth for operations that are likely to have a high payoff.

Third, it seems plausible that the ‘leisure’ state arising when control is disengaged may allow the decision-maker to survey or explore available activities, preventing a myopic focus on one potentially suboptimal line of behavior. In this sense, the labor/leisure tradeoff

demonstrated in our experiments may address a high-level instantiation of the exploration-exploitation tradeoff that arises in reinforcement-learning models of reward-based decision-making (Cohen, McClure, & Yu, 2007) and of foraging (Kolling et al., 2012). Recent research has gives traction to this idea by providing evidence that people attach intrinsic value to information derived from exploring, which trades off with the control costs incurred in exploration behavior (Wilson, Geana, White, Ludvig, & Cohen, 2014). In this sense, the cost of control might be likened to an opportunity cost of information.

Currently, these are only speculations about the purpose of the cost of cognitive control. Further experimentation will be required to establish whether the intrinsic-cost perspective can indeed be informative for a normative account of mental effort avoidance.

6.4 The neural implementation of avoidance decisions

As discussed previously, the exertion of cognitive control is well known to engage a specific set of areas within dorsolateral prefrontal cortex, dorsomedial frontal cortex, the insula and parietal cortex (Duncan, 2010). Recent evidence indicates that portions of this network also index, in tandem with related subcortical structures, both the subjective cost (Botvinick et al., 2009; J. T. McGuire & Botvinick, 2010), and the potential payoffs of control (Kouneiher, Charron, & Koechlin, 2009; Locke & Braver, 2008). But how is the decision whether to exert control implemented in the brain? A recent proposal by Shenhav et al. (2013) positions the ACC at the center of this function, integrating benefits and costs of different lines of actions into “an expected value of control” (EVC), i.e., a value signal used to determine whether and how much cognitive effort to allocate.

How does the labor/leisure framework relate to this neuroscientific theory of effort allocation? One possibility is derived from a parallel literature pointing to a neural correlate of cognitive ‘leisure,’ linking the disengagement of cognitive control with activation of a default mode network, centering on ventral prefrontal cortex (Kelly, Uddin, Biswal, Castellanos, & Milham, 2008; Raichle et al., 2001). We speculate that the labor/leisure tradeoff demonstrated in the present behavioral studies reflects the operation of a neural mechanism regulating the balance of activity between these dorsal ‘task-positive’ and ventromedial ‘task-negative’ systems (Menon & Uddin, 2010; Sridharan, Levitin, & Menon, 2008). Preliminary support for this hypothesis comes from a recent study by Pyka et al. (2009), in which resting-state default mode activity was greater following performance of a cognitively demanding task, when compared to a less demanding task. If this is correct, then activity in the task-positive and default networks, and their balance, should influence the EVC by modulating the perceived cost of control.

The present data also suggest that tonic levels of dopamine in the striatum and prefrontal cortex should be taken into account when formulating a computational account of the EVC framework. One possibility is that levels of striatal dopamine influence the system’s ‘willingness-to-pay’, increasing the maximal intensity of control that will still be determined to be worth applying, given some kind of expected payoff. One promising avenue for research would be to investigate EVC signals in the ACC as a function of dopaminergic functioning, assessed by either genotyping differences or PET, or modulated by pharmacological interventions, under task situations with varying levels of cognitive control demands.

6.5 Clinical implications

The current work has some potential to inform research in clinical psychology and neuroscience. People with attention deficit hyperactivity disorder (ADHD), for example, show impaired performance in academic settings and other learning situations. Research has shown that this deficit can be explained by the employment of less effortful learning strategies (Egeland, Nordby Johansen, & Ueland, 2010), or decreased effort expenditure (Gaultney, Kipp, Weinstein, & McNeill, 1999; Harrison, Flaro, & Armstrong, 2014). Furthermore, recent findings show that ADHD is associated with reduced levels of dopamine in the striatum and that dopamine agonists reduce symptoms (for a review see Gold, Blum, Oscar-Berman, & Braverman, 2014). Taken together, these findings suggest that ADHD might, at least in part, reflect an increase in subjective effort costs. Future work should explore to which degree the notion of cognitive costs is able to contribute knowledge in this domain.

6.6 Concluding remarks

The studies in this dissertation have made some initial progress in understanding the theoretical consequences of people's drive to avoid demands for mental effort. However, we currently only have a rudimentary understanding of the precise mechanisms underlying this behavioral tendency. The current work provides scaffolding for future research aimed at furthering our understanding of the behavioral, neural, and computational implications of the intrinsic cost of cognitive control.

Bibliography

- Aarts, E., Roelofs, A., Franke, B., Rijpkema, M., Fernandez, G., Helmich, R. C., & Cools, R. (2010). Striatal dopamine mediates the interface between motivational and cognitive control in humans: Evidence from genetic imaging. *Neuropsychopharmacology*, 35, 1943-1951.
- Aberman, J. E., & Salamone, J. D. (1999). Nucleus accumbens dopamine depletions make rats more sensitive to high ratio requirements but do not impair primary food reinforcement. *Neuroscience*, 92, 545-552.
- Ainslie, G. (1975). Specious reward: A behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, 82, 436-496.
- Allport, G. W. (1954). *The nature of prejudice*. New York, NY: Addison Wesley.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- Aron, A. R. (2011). From reactive to proactive and selective control: Developing a richer model for stopping inappropriate responses. *Biological Psychiatry*, 69, e55-e68.
- Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2004). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences*, 8, 170-177.
- Arrington, C., & Logan, G. (2004). The cost of a voluntary task switch. *Psychological Science*, 15, 610-615.
- Arrow, K., Chenery, H., Minhas, B., & Solow, R. (1961). Capital-labor substitution and economic efficiency. *The Review of Economics and Statistics*, 43, 225-250.
- Aston-Jones, G., Foote, S. L., & Bloom, F. E. (1984). Anatomy and physiology of locus coeruleus neurons: Functional implications. In M. Ziegler & C. R. Lake (Eds.), *Norepinephrine: Frontiers of clinical neuroscience* Baltimore, MD: Williams & Wilkins.
- Atkinson, J. W. (1964). *An introduction to motivation*. New York, NY: Van Nostrand.

- Atkinson, J. W., & Birch, D. (1978). *An introduction to motivation* (2nd ed.). New York, NY: Van Nostrand.
- Badre, D. (2008). Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends in Cognitive Sciences*, *12*, 193-200.
- Badre, D., & D'Esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical? *Nature Reviews Neuroscience*, *10*, 659-669.
- Baker, T., Stockwell, T., & Holroyd, C. B. (2013). Constraints on decision making: Implications from genetics, personality, and addiction. *Cognitive Affective, & Behavioral Neuroscience*, *13*, 417-436.
- Baler, R. D., & Volkow, N. D. (2006). Drug addiction: the neurobiology of disrupted self-control. *Trends in Molecular Medicine*, *12*, 559-566.
- Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. N. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, *20*, 723-742.
- Bardgett, M. E., Depenbrock, M., Downs, N., Points, M., & Green, L. (2009). Dopamine modulates effort-based decision making in rats. *Behavioral Neuroscience*, *123*, 242-251.
- Baroody, A. J., & Ginsburg, H. P. (1986). The relationship between initial meaningful and mechanical knowledge of arithmetic. In J. Hiebert (Ed.), *Conceptual and procedural knowledge: The case of mathematics*. (pp. 75-112). Hillsdale, NJ: Erlbaum.
- Baumeister, R. F. (2002). Ego depletion and self-control failure: An energy model of the self's executive function. *Self and Identity*, *1*, 129-136.
- Baumeister, R. F., Vohs, K. D., & Tice, D. M. (2007). The strength model of self-control. *Current Directions in Psychological Science*, *16*, 351-355.
- Becker, G. M., DeGroot, M. H., & Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavior Science*, *9*, 226-232.

- Berns, G., Laibson, D., & Loewenstein, G. (2007). Intertemporal choice-toward an integrative framework. *Trends in Cognitive Sciences*, 11, 482-488.
- Blough, D. S. (1966). The study of animal sensory processes by operant methods. In W. K. Honig (Ed.), *Operant behavior: Areas of research and application*. New York, NY: Appleton-Century-Crofts.
- Boettiger, C., Mitchell, J., Tavares, V. C., Robertson, M., Joslyn, G., D'Esposito, M., & Fields, H. L. (2007). Immediate reward bias in humans: fronto-parietal networks and a role for the catechol-O-methyltransferase 158Val/Val genotype. *The Journal of Neuroscience*, 52, 14383-14391.
- Bogacz, R., Brown, E., Moehlis, J., Cohen, J., & Holmes, P. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113, 700-765.
- Bolton, J. L., Marioni, R. E., Deary, I. J., Harris, S. E., Stewart, M. C., Murray, G. D., Fowkes, F. G. R., & Price, J. F. (2010). Association between polymorphisms of the dopamine receptor D2 and catechol-o-methyl transferase genes and cognitive function. *Behavior Genetics*, 40, 630-638.
- Bonner, S. E., & Sprinkle, G. B. (2002). The effects of monetary incentives on effort and task performance: theories, evidence, and a framework for research. *Accounting, Organizations and Society*, 27, 303-345.
- Botvinick, M. M. (2007). Conflict monitoring and decision making: Reconciling two perspectives on anterior cingulate function. *Cognitive, Affective, & Behavioral Neuroscience*, 7, 356-366.
- Botvinick, M. M. (2008). Hierarchical models of behavior and prefrontal function. *Trends in Cognitive Sciences*, 12, 201-208.

- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108, 624-652.
- Botvinick, M. M., Cohen, J. D., & Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: an update. *Trends in Cognitive Sciences*, 8, 539-546.
- Botvinick, M. M., Huffstetler, S., & McGuire, J. T. (2009). Effort discounting in human nucleus accumbens. *Cognitive, Affective, & Behavioral Neuroscience*, 9, 16-27.
- Botvinick, M. M., Nystrom, L. E., Fissell, K., Carter, C. S., & Cohen, J. D. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature*, 402, 179-181.
- Boureau, Y.-L., & Dayan, P. (2010). Opponency Revisited: Competition and Cooperation Between Dopamine and Serotonin. *Neuropsychopharmacology*, 36, 74-97.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433-436.
- Braver, T. S. (2012). The variable nature of cognitive control: A dual-mechanisms framework. *Trends in Cognitive Sciences*, 16, 106-113.
- Braver, T. S., Cohen, J. D., Nystrom, L. E., Jonides, J., Smith, E. E., & Noll, D. C. (1997). A parametric study of prefrontal cortex involvement in human working memory. *NeuroImage*, 5, 49-62.
- Brehm, J. W., Wright, R. A., Solomon, S., Silka, L., & Greenberg, J. (1983). Perceived difficulty, energization, and the magnitude of goal valence. *Journal of Experimental Social Psychology*, 19, 21-48.
- Cagniard, B., Balsam, P. D., Brunner, D., & Zhuang, X. (2005). Mice with Chronically Elevated Dopamine Exhibit Enhanced Motivation, but not Learning, for a Food Reward. *Neuropsychopharmacology*, 31, 1362-1370.
- Camerer, C. F., Babcock, L., Loewenstein, G., & Thaler, R. (1997). Labor supply of New York City cabdrivers: One day at a time. *The Quarterly Journal of Economics*, 112, 407-441.

- Camerer, C. F., & Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty*, 19, 7-42.
- Casey, B. J., Somerville, L. H., Gotlib, I. H., Ayduk, O., Franklin, N. T., Askren, M. K., Jonides, J., Berman, M. G., Wilson, N. L., Teslovich, T., Glover, G., Zayas, V., Mischel, W., & Shoda, Y. (2011). Behavioral and neural correlates of delay of gratification 40 years later. *Proceedings of the National Academy of Sciences USA*, 108, 14998-15003.
- Chambers, E. S., Bridge, M. W., & Jones, D. A. (2009). Carbohydrate sensing in the human mouth: effects on exercise performance and brain activity. *The Journal of Physiology*, 587, 1779-1794.
- Charness, G., & Kuhn, P. (2010). What can labor economists learn from the lab? *IZA Discussion Paper no. 4941*.
- Chen, M. K., Lakshminarayanan, V., & Santos, L. R. (2006). How basic are behavioral biases? Evidence from capuchin monkey trading behavior. *Journal of Political Economy*, 114, 517-537.
- Christoff, K., Prabhakaran, V., Dorfman, J., Zhao, Z., Kroger, J. K., Holyoak, K. J., & Gabrieli, J. D. E. (2001). Rostrolateral prefrontal cortex involvement in relational integration during reasoning. *NeuroImage*, 14, 1136-1149.
- Clarke, D. D., & Sokoloff, L. (1998). Circulation and energy metabolism of the brain. In G. Siegel, B. Agranoff, R. Albers, S. Fisher & M. Uhler (Eds.), *Basic neurochemistry: Molecular, cellular and medical aspects*. (6th ed., pp. 637-669). Philadelphia, PA: Lippincott Raven.
- Cobb, C. W., & Douglas, P. H. (1928). A theory of production. *The American Economic Review*, 18, 139-165.
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*, 97, 332-361.

- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362, 933-942.
- Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J., & Smith, E. E. (1997). Temporal dynamics of brain activation during a working memory task. *Nature*, 386, 604-608.
- Cohen, J. D., & Servan-Schreiber, D. (1992). Context, cortex, and dopamine: A connectionist approach to behavior and biology in schizophrenia. *Psychological Review*, 99, 45-77.
- Colzato, L. S., Waszak, F., Nieuwenhuis, S., Posthuma, D., & Hommel, B. (2010). The flexible mind is associated with the catechol-O-methyltransferase (COMT) Val158Met polymorphism: Evidence for a role of dopamine in the control of task-switching. *Neuropsychologia*, 48, 2764-2768.
- Conover, K. L., & Shizgal, P. (2005). Employing labor-supply theory to measure the reward value of electrical brain stimulation. *Games and Economic Behavior*, 52, 283-304.
- Cools, R., & D'Esposito, M. (2011). Inverted-U-Shaped Dopamine Actions on Human Working Memory and Cognitive Control. *Biological Psychiatry*, 69, e113-e125.
- Cools, R., Gibbs, S., Miyakawa, A., Jagust, W., & D'Esposito, M. (2008). Working memory capacity predicts dopamine synthesis capacity in the human striatum. *Journal of Neuroscience*, 28, 1208-1212.
- Cools, R., & Robbins, T. (2004). Chemistry of the adaptive mind. *Philosophical Transactions of the Royal Society A: Mathematical, Physical, and Engineering Sciences*, 362, 2871-2888.
- Cools, R., Sheridan, M., Jacobs, E., & D'Esposito, M. (2007). Impulsive personality predicts dopamine-dependent changes in frontostriatal activity during component processes of working memory. *Journal of Neuroscience*, 27, 5506-5514.

- Cousins, M. S., Atherton, A., & Turner, L. (1996). Nucleus accumbens dopamine depletions alter relative response allocation in a T-maze cost/benefit task. *Behavioural Brain Research*, 74, 189-197.
- Cox, R. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, 29, 162-173.
- Croxson, P. L., Walton, M. E., O'Reilly, J. X., Behrens, T. E. J., & Rushworth, M. F. S. (2009). Effort-based cost-benefit valuation and the human brain. *Journal of Neuroscience*, 29, 4531-4541.
- Cuvo, A. J., Lerch, L. J., Leurquin, D. A., Gaffaney, T., & Poppen, R. L. (1998). Response allocation to concurrent fixed-ratio reinforcement schedules with work requirements by adults with mental retardation and typical preschool children. *Journal of Applied Behavior Analysis*, 31, 43-63.
- De Jong, R. (2000). An intention-activation account of residual switch costs. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance*. (pp. 357-376). Cambridge, MA: MIT Press.
- de Wit, S., Standing, H. R., Devito, E. E., Robinson, O. J., Ridderinkhof, K. R., Robbins, T. W., & Sahakian, B. J. (2012). Reliance on habits at the expense of goal-directed control following dopamine precursor depletion. *Psychopharmacology*, 219, 621-631.
- Del Arco, A., & Mora, F. (2008). Prefrontal cortex–nucleus accumbens interaction: In vivo modulation by dopamine and glutamate in the prefrontal cortex. *Pharmacology Biochemistry and Behavior*, 90, 226-235.
- Denk, F., Walton, M. E., Jennings, K. A., Sharp, T., Rushworth, M. F. S., & Bannerman, D. M. (2004). Differential involvement of serotonin and dopamine systems in cost-benefit decisions about delay or effort. *Psychopharmacology*, 179, 587-596.

- Dickinson, D. L. (1999). An experimental examination of labor supply and work intensities. *Journal of Labor Economics*, 17, 638-670.
- Doll, B., Hutchison, K., & Frank, M. (2011). Dopaminergic genes predict individual differences in susceptibility to confirmation bias. *The Journal of Neuroscience*, 31, 6188-6198.
- Droll, J. A., & Hayhoe, M. M. (2007). Trade-offs between gaze and working memory use. *Journal of Experimental Psychology: Human Perception and Performance*, 33, 1352-1365.
- Duckworth, A. L., & Seligman, M. E. P. (2005). Self-discipline outdoes IQ in predicting academic performance of adolescents. *Psychological Science*, 16, 939-944.
- Duncan, J. (2010). The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends in Cognitive Sciences*, 14, 172-179.
- Duncan, J., & Owen, A. M. (2000). Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends in Neurosciences*, 23, 475-483.
- Egeland, J., Nordby Johansen, S., & Ueland, T. (2010). Do low-effort learning strategies mediate impaired memory in ADHD? *Journal of Learning Disabilities*, 43, 430-440.
- Eichele, T., Debener, S., Calhoun, V., Specht, K., Engel, A., Hugdahl, K., von Cramon, D., & Ullsperger, M. (2008). Prediction of human errors by maladaptive changes in event-related brain networks. *Proceedings of the National Academy of Sciences USA*, 105, 6173-6178.
- Eisenberger, N. I., Lieberman, M. D., & Williams, K. D. (2003). Does rejection hurt? An fMRI study of social exclusion. *Science*, 302, 290.
- Engelmann, J. B., Damaraju, E., Padmala, S., & Pessoa, L. (2009). Combined effects of attention and motivation on visual task performance: transient and sustained motivational effects. *Frontiers in Human Neuroscience*, 3, 4.
- Engelmann, J. B., & Pessoa, L. (2007). Motivation sharpens exogenous spatial attention. *Emotion*, 7, 668-674.

- Engle, R., & Kane, M. (2003). Executive attention, working memory capacity, and a two-factor theory of cognitive control. *Psychology of Learning and Motivation*, 44, 145-199.
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, 16, 143-149.
- Farber, H. S. (2005). Is tomorrow another day? The labor supply of New York City cabdrivers. *Journal of Political Economy*, 113, 46-82.
- Fehr, E., & Goette, L. (2007). Do workers work more if wages are high? Evidence from a randomized field experiment. *American Economic Review*, 97, 298-317.
- Fehr, E., Goette, L., & Zehnder, C. (2009). A behavioral account of the labor market: The role of fairness concerns. *Annual Review of Economics*, 1, 355-384.
- Feng, S. F., Schwemmer, M., Gershman, S. J., & Cohen, J. D. (2014). Multitasking versus multiplexing: Toward a normative account of limitations in the simultaneous execution of control-demanding behaviors. *Cognitive, Affective, & Behavioral Neuroscience*, 14, 129-146.
- Figner, B., Knoch, D., Johnson, E. J., Krosch, A. R., Lisanby, S. H., Fehr, E., & Weber, E. U. (2010). Lateral prefrontal cortex and self-control in intertemporal choice. *Nature*, 13, 538–539.
- Fiske, S. T., & Taylor, S. E. (1991). *Social Cognition*. New York: McGraw-Hill.
- Floresco, S., Tse, M., & Ghods-Sharifi, S. (2008). Dopaminergic and glutamatergic regulation of effort- and delay-based decision making. *Neuropsychopharmacology*, 33, 1966-1979.
- Fox, M., Snyder, A., Vincent, J., Corbetta, M., Van Essen, D., & Raichle, M. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences USA*, 102, 9673-9678.
- Frank, M., & Hutchison, K. (2009). Genetic contributions to avoidance-based decisions: striatal D2 receptor polymorphisms. *Neuroscience*, 164, 131-140.

- Frank, R. H. (1999). *Microeconomics and behavior*. New York: McGraw-Hill.
- Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives*, 19, 25-42.
- Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40, 351-401.
- Friedrich, A. M., & Zentall, T. R. (2004). Pigeons shift their preference toward locations of food that take more effort to obtain. *Behavioural Processes*, 67, 405-415.
- Fry, P. S. (1975). Affect and resistance to temptation. *Developmental Psychology*, 11, 466-472.
- Fudenberg, D., & Levine, D. K. (2006). A dual-self model of impulse control. *American Economic Review*, 96, 1449-1476.
- Fürst, A. J., & Hitch, G. J. (2000). Separate roles for executive and phonological components of working memory in mental arithmetic. *Memory & Cognition*, 28, 774-782.
- Gailliot, M. T., & Baumeister, R. F. (2007). The physiology of willpower: Linking blood glucose to self-control. *Personality and Social Psychology Review*, 11, 303-327.
- Gailliot, M. T., Baumeister, R. F., DeWall, C. N., Maner, J. K., Plant, E. A., Tice, D. M., Brewer, L. E., & 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, 325-336.
- Gaultney, J. F., Kipp, K., Weinstein, J., & McNeill, J. (1999). Inhibition and mental effort in attention deficit hyperactivity disorder. *Journal of Developmental and Physical Disabilities*, 11, 105-114.
- Geyer, M. A., Puerto, A., Dawsey, W. J., Knapp, S., & Bullard, W. P. (1976). Histologic and enzymatic studies of the mesolimbic and mesostriatal serotonergic pathways. *Brain Research*, 106, 241-256.

- Gibson, E. L. (2007). Carbohydrates and mental function: feeding or impeding the brain? *Nutrition Bulletin*, 32, 71-83.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review*, 103, 650-669.
- Gläscher, J., Daw, N., Dayan, P., & O'Doherty, J. (2010). States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, 66, 585-595.
- Glass, D. C., Siger, J. E., & Friedman, L. N. (1969). Psychic cost of adaptation to an environmental stressor. *Journal of Personality and Social Psychology*, 12, 200-210.
- Gold, M., Blum, K., Oscar-Berman, M., & Braverman, E. (2014). Low Dopamine Function in Attention Deficit/Hyperactivity Disorder: Should Genotyping Signify Early Diagnosis in Children? *Postgraduate Medicine*, 126, 153-177.
- Gratton, G., Coles, M. G., & Donchin, E. (1992). Optimizing the use of information: strategic control of activation of responses. *Journal of Experimental Psychology: General*, 121, 480-506.
- Gray, W. D. (2000). The nature and processing of errors in interactive behavior. *Cognitive Science*, 24, 205-248.
- Green, L., & Rachlin, H. (1991). Economic substitutability of electrical brain stimulation, food, and water. *Journal of the Experimental Analysis of Behavior*, 55, 133-143.
- Greene, J., Nystrom, L., Engell, A., Darley, J., & Cohen, J. D. (2004). The neural bases of cognitive conflict and control in moral judgment. *Neuron*, 44, 389-400.
- Gul, F., & Pesendorfer, W. (2001). Temptation and self-control. *Econometrica*, 69, 1403-1435.
- 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, 495-525.

- Hare, T. A., Camerer, C. F., & Rangel, A. (2009). Self-control in decision-making involves modulation of the vmPFC valuation system. *Science*, *324*, 646-648.
- Harrison, A. G., Flaro, L., & Armstrong, I. (2014). Rates of effort test failure in children with ADHD: An exploratory study. *Applied Neuropsychology: Child*, Advance online publication. doi:10.1080/21622965.2013.21850581.
- Heatherton, T. F., Striepe, M., & Wittenberg, L. (1998). Emotional distress and disinhibited eating: The role of self. *Personality and Social Psychology Bulletin*, *24*, 301-313.
- Heatherton, T. F., & Wagner, D. D. (2011). Cognitive neuroscience of self-regulation failure. *Trends in Cognitive Sciences*, *15*, 132-139.
- Heinz, A., Goldman, D., Jones, D. W., Palmour, R., Hommer, D., Gorey, J. G., Lee, K. S., Linnoila, M., & Weinberger, D. R. (2000). Genotype influences in vivo dopamine transporter availability in human striatum. *Neuropsychopharmacology*, *22*, 133-139.
- Hirvonen, M., Laakso, A., Nägren, K., Rinne, J. O., Pohjalainen, T., & Hietala, J. (2004). C957T polymorphism of the dopamine D2 receptor (DRD2) gene affects striatal DRD2 availability in vivo. *Molecular Psychiatry*, *9*, 1060-1061.
- Hitch, G. J. (1978). The role of short-term working memory in mental arithmetic. *Cognitive Psychology*, *16*, 393-408.
- Hommel, B. (2004). Event files: Feature binding in and across perception and action. *Trends in Cognitive Sciences*, *8*, 494-500.
- Hull, C. L. (1943). *Principles of behavior*. New York: Appleton Century
- Inzlicht, M., & Schmeichel, B. (2012). What is ego depletion? Toward a mechanistic revision of the resource model of self-control. *Perspectives on Psychological Science*, *7*, 450-463.

- Jimura, K., Locke, H. S., & Braver, T. S. (2010). Prefrontal cortex mediation of cognitive enhancement in rewarding motivational contexts. *Proceedings of the National Academy of Sciences USA*, *107*, 8871-8876.
- Job, V., Dweck, C., & Walton, G. (2010). Ego depletion—Is it all in your head? *Psychological Science*, *21*, 1686–1693.
- Kable, J. W., & Glimcher, P. (2010). An “as soon as possible” effect in human intertemporal decision making: behavioral evidence and neural mechanisms. *Journal of Neurophysiology*, *103*, 2513-2531.
- Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, *10*, 1625-1633.
- Kagel, J. H., Battalio, R. C., & Green, L. (1995). *Economic choice theory: An experimental analysis of animal behavior*. Cambridge: Cambridge University Press.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263-292.
- Kelly, A. M., Uddin, L. Q., Biswal, B. B., Castellanos, F. X., & Milham, M. P. (2008). Competition between functional brain networks mediates behavioral variability. *NeuroImage*, *39*, 527-537.
- Kerns, J. G., Cohen, J. D., suffix, A. W. M. s. I., Cho, R. Y., Stenger, V. A., & Carter, C. S. (2004). Anterior cingulate conflict monitoring and adjustments in control. *Science*, *303*, 1023-1026.
- Knapp, A., & Clark, M. S. (1991). Some detrimental effects of negative mood on individuals' ability to solve resource dilemmas. *Personality and Social Psychology Bulletin*, *17*, 678-688.
- Kolling, N., Behrens, T. E. J., Mars, R. B., & Rushworth, M. F. S. (2012). Neural mechanisms of foraging. *Science*, *336*, 95-98.
- Kool, W., & Botvinick, M. M. (2013). The intrinsic of cognitive control. *Behavioral and Brain Sciences*, *36*, 697-698.

- Kool, W., & Botvinick, M. M. (2014). A labor/leisure trade-off in cognitive control. *Journal of Experimental Psychology: General*, 143, 131-141.
- 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, 665-682.
- 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, e72626.
- Kouneiher, F., Charron, S., & Koechlin, E. (2009). Motivation and cognitive control in the human prefrontal cortex. *Nature Neuroscience*, 12, 939-945.
- Kurzban, R. (2010). Does the brain consume additional glucose during self-control tasks? *Evolutionary Psychology*, 8, 244-259.
- Kurzban, R., Duckworth, A. L., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences*, 36, 661-726.
- Landau, S. M., Lal, R., O'Neil, J. P., Baker, S., & Jagust, W. J. (2009). Striatal dopamine and working memory. *Cerebral Cortex*, 19, 445-454.
- Locke, H. S., & Braver, T. S. (2008). Motivational influences on cognitive control: Behavior, brain activation, and individual differences. *Cognitive Affective, & Behavioral Neuroscience*, 8, 99-112.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279-280.
- MacLeod, C. M., Hunt, E. B., & Mathews, N. N. (1978). Individual differences in the verification of sentence—picture relationships. *Journal of Verbal Learning and Verbal Behavior*, 17, 493-507.
- Madsen, P. L., Hasselbalch, S. G., Hagemann, L. P., Olsen, K. S., Bülow, J., Holm, S., Wildschjødztz, G., Paulson, O. B., & Lassen, N. A. (1995). Persistent resetting of the cerebral oxygen/glucose uptake ratio by brain activation: evidence obtained with the Kety-Schmidt technique. *Journal of Cerebral Blood Flow and Metabolism*, 15, 485-491.

- Mathews, N. N., Hunt, E. B., & MacLeod, C. M. (1980). Strategy choice and strategy training in sentence—picture verification. *Journal of Verbal Learning and Verbal Behavior*, 19, 531-548.
- Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. K. Commons, J. E. Mazur, J. A. Nevin & H. Rachlin (Eds.), *Quantitative analysis of Behavior: The effects of delay and intervening events on reinforcement value*. (pp. 55-73). Hillsdale, NJ: Lawrence Erlbaum Associates.
- McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2007). Time discounting for primary rewards. *Journal of Neuroscience*, 27, 5796-5804.
- McClure, S. M., Laibson, D., Loewenstein, G., & Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science*, 306, 503-507.
- McGuire, J. T., & Botvinick, M. M. (2010). Prefrontal cortex, cognitive control, and the registration of decision costs. *Proceedings of the National Academy of Sciences USA*, 107, 7922–7926.
- McGuire, W. J. (1969). The nature of attitudes and attitude change. In G. Lindzey & E. Aronson (Eds.), *Handbook of Social Psychology*. (2nd ed., Vol. 3, pp. 136-314). Reading, MA: Addison-Wesley.
- Menon, V., & Uddin, L. Q. (2010). Saliency, switching, attention and control: a network model of insula function. *Brain Structure and Function*, 214, 655-667.
- Messier, C. (2004). Glucose improvement of memory: a review. *European Journal of Pharmacology*, 490, 33-57.
- Meyer-Lindenberg, A., Kohn, P. D., Kolachana, B., Kippenhan, S., McInerney-Leo, A., Nussbaum, R., Weinberger, D. R., & Berman, K. F. (2005). Midbrain dopamine and prefrontal function in humans: interaction and modulation by COMT genotype. *Nature Neuroscience*, 8, 594-596.
- Meyer-Lindenberg, A., Straub, R. E., Lipska, B. K., Verchinski, B. A., Goldberg, T., Callicott, J. H., Egan, M. F., Huffaker, S. S., Mattay, V. S., Kolachana, B., Kleinman, J. E., & Weinberger, D.

- R. (2007). Genetic evidence implicating DARPP-32 in human frontostriatal structure, function, and cognition. *Journal of Clinical Investigation*, 117, 672-682.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Reviews in Neuroscience*, 24, 167-202.
- Miner, M., & Park, D. C. (2004). A lifespan database of adult facial stimuli. *Behavior Research Methods, Instruments, & Computers*, 36, 630-633.
- Mischel, W. (2011). Self-control theory. In A. W. Kruglanski, E. T. Higgins & P. A. M. van Lange (Eds.), *Handbook of Theories of Social Psychology* (Vol. 2, pp. 1-22). Thousand Oaks, CA: Sage.
- Mischel, W., Shoda, Y., & Rodriguez, M. L. (1989). Delay of gratification in children. *Science*, 244, 933-938.
- Mitaki, S., Isomura, M., Maniwa, K., Yamasaki, M., Nagai, A., Nabika, T., & Yamaguchi, S. (2012). Impact of five SNPs in dopamine-related genes on executive function. *Acta Neurologica Scandinavica*, 127, 70-76.
- Moffitt, T. E., Arseneault, L., Belsky, D., Dickson, N., Hancox, R. J., Harrington, H., Houts, R., Poulton, R., Roberts, B. W., Ross, S., Sears, M. R., Thomson, W. M., & Caspi, A. (2011). A gradient of childhood self-control predicts health, wealth, and public safety. *Proceedings of the National Academy of Sciences USA*, 108, 2693-2698.
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, 7, 134-140.
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *The Journal of Neuroscience*, 16, 1936-1947.
- Muraven, M. (2011). Ego-depletion: Theory and evidence. In R. M. Ryan (Ed.), *The Oxford handbook of human motivation*. Oxford: Oxford University Press.
- Muraven, M., & Baumeister, R. F. (2000). Self-regulation and depletion of limited resources: Does self-control resemble a muscle? *Psychological Bulletin*, 126, 247-259.

- Muraven, M., & Slessareva, E. (2003). Mechanisms of self-control failure: Motivation and limited resources. *Personality and Social Psychology Bulletin*, 29, 894-906.
- Muraven, M., Tice, D. M., & Baumeister, R. F. (1998). Self-control as limited resource: regulatory depletion patterns. *Journal of Personality and Social Psychology*, 74, 774-789.
- Murphy, F. C., Rubinsztein, J. S., Michael, A., Rogers, R. D., Robbins, T. W., Paykel, E. S., & Sahakian, B. J. (2001). Decision-making cognition in mania and depression. *Psychological Medicine*, 31, 679-693.
- Nicholson, W., & Snyder, C. M. (2008). *Microeconomic theory: Basic principles and extensions*. Boston, MA: Cengage Learning.
- Nieuwenhuis, S., Aston-Jones, G., & Cohen, J. (2005). Decision making, the P3, and the locus coeruleus-norepinephrine system. *Psychological Bulletin*, 131, 510-532.
- Okuyama, Y., Ishiguro, H., Toru, M., & Arinami, T. (1999). A genetic polymorphism in the promoter region of DRD4 associated with expression and schizophrenia. *Biochemical and Biophysical Research*, 258, 292-295.
- Otto, A. R., Raio, C. M., Chiang, A., Phelps, E., & Daw, N. (2013). Working-memory capacity protects model-based learning from stress. *Proceedings of the National Academy of Sciences USA*, 110, 20941-20946.
- Pasquereau, B., & Turner, R. S. (2013). Limited Encoding of Effort by Dopamine Neurons in a Cost-Benefit Trade-off Task. *Journal of Neuroscience*, 33, 8288-8300.
- Paxton, J. M., Ungar, L., & Greene, J. D. (2011). Reflection and reasoning in moral judgment. *Cognitive Science*, 36, 163-177.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge, England: Cambridge University Press.

- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437-442.
- Phillips, P., Walton, M., & Jhou, T. (2007). Calculating utility: preclinical evidence for cost–benefit analysis by mesolimbic dopamine. *Psychopharmacology*, 191, 483-495.
- Posner, M. I., & DiGirolamo, G. J. (1998). Executive attention: Conflict, target detection, and cognitive control. In R. Parasuraman (Ed.), *The attentive brain*. (pp. 401-423). Cambridge, MA: MIT Press.
- Pyka, M., Beckmann, C., Schöning, S., Hauke, S., Heider, D., Kugel, H., Arolt, V., & Konrad, C. (2009). Impact of working memory load on fMRI resting state pattern in subsequent resting phases. *PLoS ONE*, 4, e7198.
- Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences USA*, 98, 676-682.
- Reichle, E. D., Carpenter, P. A., & Just, M. A. (2000). The Neural Bases of Strategy and Skill in Sentence–Picture Verification. *Cognitive Psychology*, 40, 261-295.
- Ridderinkhof, K. R., Ullsperger, M., Crone, E. A., & Nieuwenhuis, S. (2004). The role of the medial frontal cortex in cognitive control. *Science*, 306, 443-447.
- Rodriguez-Jimenez, R., Hoenicka, J., Jimenez-Arriero, M. A., Ponce, G., Bagney, A., Aragues, M., & Palomo, T. (2006). Performance in the Wisconsin Card Sorting test and the C957T polymorphism of the DRD2 gene in healthy volunteers. *Neuropsychobiology*, 54, 166-170.
- Rossi, A., Pessoa, L., Desimone, R., & Ungerleider, L. G. (2009). The prefrontal cortex and the executive control of attention. *Experimental Brain Research*, 192, 489-497.
- Rudebeck, P. H., Walton, M. E., Smyth, A. N., Bannerman, D. M., & Rushworth, M. F. S. (2006). Separate neural pathways process different decision costs. *Nature Neuroscience*, 9, 1161-1168.

- Rushworth, M. F. S., & Behrens, T. E. J. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. *Nature Neuroscience*, 11, 389-397.
- Ryan, R. M., & Deci, E. L. (2000a). Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary Educational Psychology*, 25, 54-67.
- Ryan, R. M., & Deci, E. L. (2000c). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, 55, 68-78.
- Salamone, J. D., & Correa, M. (2012). The mysterious motivational functions of mesolimbic dopamine. *Neuron*, 76, 470-485.
- Salamone, J. D., Correa, M., Farrar, A. M., & Mingote, S. M. (2007). Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. *Psychopharmacology*, 191, 461-482.
- Salamone, J. D., Correa, M., Farrar, A. M., Nunes, E. J., & Pardo, M. (2009). Dopamine, behavioral economics, and effort. *Frontiers in Behavioral Neuroscience*, 3, 2-12.
- Salamone, J. D., Correa, M., Mingote, S. M., Weber, S. M., & Farrar, A. M. (2006). Nucleus accumbens dopamine and the forebrain circuitry involved in behavioral activation and effort-related decision making: Implications for understanding anergia and psychomotor slowing in depression. *Current Psychiatry Reviews*, 2, 267-280.
- Salamone, J. D., Correa, M., & Weber, S. M. (2003). Nucleus accumbens dopamine and the regulation of effort in food-seeking behavior: implications for studies of natural motivation, psychiatry, and drug abuse. *Journal of Pharmacology and Experimental Therapeutics*, 305, 1-8.
- Salamone, J. D., Cousins, M., & Bucher, S. (1994). Anhedonia or anergia? Effects of haloperidol and nucleus accumbens dopamine depletion on instrumental response selection in a T-maze cost/benefit procedure. *Behavioural Brain Research*, 65, 221-229.

- Salamone, J. D., Wisniecki, A., Carlson, B. B., & Correa, M. (2001). Nucleus accumbens dopamine depletions make animals highly sensitive to high fixed ratio requirements but do not impair primary food reinforcement. *Neuroscience*, *105*, 863-870.
- Samanez-Larkin, G. R., Buckholtz, J. W., Cowan, R. L., Woodward, N. D., Li, R., Ansari, M. S., Arrington, C. M., Baldwin, R. M., Smith, C. E., Treadway, M. T., Kessler, R. M., & Zald, D. H. (2013). A Thalamocortico-striatal Dopamine Network for Psychostimulant-Enhanced Human Cognitive Flexibility. *BPS*, *74*, 99-105.
- Sarter, M., Gehring, W. J., & Kozak, R. (2006). More attention must be paid: the neurobiology of attentional effort. *Brain Research Reviews*, *51*, 145-160.
- Schoupe, N., Ridderinkhof, K. R., Verguts, T., & Notebaert, W. (2014). Context-specific control and context selection in conflict tasks. *Acta Psychologica*, *146*, 63-66.
- Schultz, W. (1999). The reward signal of midbrain dopamine neurons. *News in Physiological Sciences*, *14*, 249-255.
- Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron*, *36*, 241-263.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*, 1593-1599.
- Schweimer, J., & Hauber, W. (2006). Dopamine D1 receptors in the anterior cingulate cortex regulate effort-based decision making. *Learning & Memory*, *13*, 777.
- Shah, A., & Oppenheimer, D. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, *134*, 207.
- Shamosh, N. A., Deyoung, C. G., Green, A. E., Reis, D. L., Johnson, M. R., Conway, A. R. A., Engle, R. W., Braver, T. S., & Gray, J. R. (2008). Individual differences in delay discounting: Relation to intelligence, working memory and anterior prefrontal cortex. *Psychological Science*, *19*, 904-911.

- Shenhav, A., Botvinick, Matthew M., & Cohen, Jonathan D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron*, 79, 217-240.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84, 127-190.
- Shugan, S. M. (1980). The cost of thinking. *Journal of Consumer Research*, 7, 99-111.
- Simon, D. A., & Daw, N. D. (2011). Neural correlates of forward planning in a spatial decision task in humans. *Journal of Neuroscience*, 31, 5526-5539.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69, 99-118.
- Simon, H. A. (1956). A comparison of game theory and learning theory. *Psychometrika*, 21, 267-272.
- Smith, V. L., & Walker, J. M. (1993). Monetary rewards and decision cost in experimental economics. *Economic Inquiry*, 31, 245-261.
- Sohn, M. H., Ursu, S., Anderson, J. R., Stenger, V. A., & Carter, C. S. (2000). The role of prefrontal cortex and posterior parietal cortex in task switching. *Proceedings of the National Academy of Sciences USA*, 97, 13448-13453.
- Solomon, R. L. (1948). The influence of work on behavior. *Psychological Bulletin*, 45, 1-40.
- Spanagel, R., & Weiss, F. (1999). The dopamine hypothesis of reward: past and current status. *Trends in Neurosciences*, 22, 521-527.
- Sridharan, D., Levitin, D., & Menon, V. (2008). A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. *Proceedings of the National Academy of Sciences USA*, 105, 12569-12574.
- Stephens, D. W., & Krebs, J. R. (1986). *Foraging Theory*. Princeton, NJ: Princeton University Press.

- Stewart, C. C., Wright, R. A., Hui, S.-K. A., & Simmons, A. (2009). Outcome expectancy as a moderator of mental fatigue influence on cardiovascular response. *Psychophysiology*, *46*, 1141-1149.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *18*, 643-662.
- Taatgen, N. A. (2007). The minimal control principle. In W. Gray (Ed.), *Integrated models of cognitive systems*. (pp. 368-379). Oxford, England: Oxford University Press.
- Tangney, J. P., Baumeister, R. F., & Boone, A. L. (2004). High self-control predicts good adjustment, less pathology, better grades, and interpersonal success. *Journal of Personality*, *72*, 271-324.
- Taylor, S. E. (1981). The interface of cognitive and social psychology. In J. H. Harvey (Ed.), *Cognition, social behavior, and the environment*. (pp. 189-211). Hillsdale, NJ: Erlbaum.
- Thaler, R., & Shefrin, H. (1981). An economic theory of self-control. *Journal of Political Economy*, *89*, 392-406.
- Thompson, M. E. (1944). Learning as a function of the absolute and relative amounts of work. *Journal of Experimental Psychology*, *34*, 506-515.
- Tice, D. M., Baumeister, R. F., Shmueli, D., & Muraven, M. (2007). Restoring the self: Positive affect helps improve self-regulation following ego depletion. *Journal of Experimental Social Psychology*, *43*, 379-384.
- Tierney, J. (2011, August 17). Do you suffer from decision fatigue? *New York Times*.
- Treadway, M. T., Buckholtz, J. W., Cowan, R. L., Woodward, N. D., Li, R., Ansari, M. S., Baldwin, R. M., Schwartzman, A. N., Kessler, R. M., & Zald, D. H. (2012). Dopaminergic mechanisms of individual differences in human effort-based decision-making. *Journal of Neuroscience*, *32*, 6170-6176.

- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124.
- Tzschentke, T. M. (2001). Pharmacology and behavioral pharmacology of the mesocortical dopamine system. *Progress in Neurobiology*, 63, 241-320.
- Ullsperger, M., Bylsma, L. M., & Botvinick, M. M. (2005). The conflict adaptation effect: It's not just priming. *Cognitive, Affective, & Behavioral Neuroscience*, 5, 467-472.
- van der Plas, E. A. A., Crone, E. A., van den Wildenberg, W. P. M., Tranel, D., & Bechara, A. (2008). Executive control deficits in substance-dependent individuals: A comparison of alcohol, cocaine, and methamphetamine and of men and women. *Journal of Clinical and Experimental Neuropsychology*, 31, 706-719.
- Vijayraghavan, S., Wang, M., Birnbaum, S., Williams, G., & Arnsten, A. (2007). Inverted-U dopamine D1 receptor actions on prefrontal neurons engaged in working memory. *Nature Neuroscience*, 10, 376-384.
- Wagner, D. D., & Heatherton, T. F. (2012). Self-regulatory depletion increases emotional reactivity in the amygdala. *Social Cognitive and Affective Neuroscience*, 8, 410-417.
- Walton, M. E., Kennerley, S. W., Bannerman, D. M., Phillips, P. E. M., & Rushworth, M. F. S. (2006). Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Networks*, 19, 1302-1314.
- Wardle, M. C., Treadway, M. T., Mayo, L. M., Zald, D. H., & de Wit, H. (2011). Amping up effort: Effects of d-amphetamine on human effort-based decision-making. *Journal of Neuroscience*, 31, 16597-16602.
- Weiner, B. (1985). An attributional theory of achievement motivation and emotion. *Psychological Review*, 92, 548-573.

- Weissman, D. H., Roberts, K. C., Visscher, K. M., & Woldorff, M. G. (2006). The neural bases of momentary lapses in attention. *Nature Neuroscience*, 9, 971-978.
- Wenzlaff, R. M., & Bates, D. E. (1998). Unmasking a cognitive vulnerability to depression: how lapses in mental control reveal depressive thinking. *Journal of Personality and Social Psychology*, 75, 1559-1571.
- 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*, 22, e68210.
- Wilkinson, L. S. (1997). The nature of interactions involving prefrontal and striatal dopamine systems. *Journal of Psychopharmacology*, 11, 143-150.
- Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, 143, 2074-2081.
- Wunderlich, K., Smittenaar, P., & Dolan, R. (2012). Dopamine enhances model-based over model-free choice behavior. *Neuron*, 75, 418–424.
- Yeung, N., Botvinick, M., & Cohen, J. (2004). The neural basis of error detection: Conflict monitoring and the error-related negativity. *Psychological Review*, 111, 931-959.
- Yeung, N., & Monsell, S. (2003). Switching between tasks of unequal familiarity: the role of stimulus-attribute and response-set selection. *Journal of Experimental Psychology: Human Perception and Performance*, 29, 455-469.
- Zhang, Y., Bertolino, A., Fazio, L., Blasi, G., Rampino, A., Romano, R., Lee, M.-L., Xiao, T., Papp, A., & Wang, D. (2007). Polymorphisms in human dopamine D2 receptor gene affect gene expression, splicing, and neuronal activity during working memory. *Proceedings of the National Academy of Sciences USA*, 104, 20552-20557.

Zysset, S., Müller, K., Lohmann, G., & von Cramon, D. Y. (2001). Color-word matching stroop task: Separating interference and response conflict. *NeuroImage*, 13, 29-36.