



University of Pennsylvania
ScholarlyCommons

Publicly Accessible Penn Dissertations

2017

Individual Differences In Value-Based Decision-Making: Learning And Time Preference

Marieta Pehlivanova

University of Pennsylvania, marieta.p@gmail.com

Follow this and additional works at: <https://repository.upenn.edu/edissertations>

 Part of the [Cognitive Psychology Commons](#)

Recommended Citation

Pehlivanova, Marieta, "Individual Differences In Value-Based Decision-Making: Learning And Time Preference" (2017). *Publicly Accessible Penn Dissertations*. 2520.
<https://repository.upenn.edu/edissertations/2520>

This paper is posted at ScholarlyCommons. <https://repository.upenn.edu/edissertations/2520>
For more information, please contact repository@pobox.upenn.edu.

Individual Differences In Value-Based Decision-Making: Learning And Time Preference

Abstract

Human decisions are strongly influenced by past experience or by the subjective values attributed to available choice options. Although decision processes show some common trends across individuals, they also vary considerably between individuals. The research presented in this dissertation focuses on two domains of decision-making, related to learning and time preference, and examines factors that explain decision-making differences between individuals. First, we focus on a form of reinforcement learning in a dynamic environment. Across three experiments, we investigated whether individual differences in learning were associated with differences in cognitive abilities, personality, and age. Participants made sequential predictions about an on-screen location in a video game. Consistent with previous work, participants showed high variability in their ability to implement normative strategies related to surprise and uncertainty. We found that higher cognitive ability, but not personality, was associated with stronger reliance on the normative factors that should govern learning. Furthermore, learning in older adults (age 60+) was less influenced by uncertainty, but also less influenced by reward, a non-normative factor that has substantial effects on learning across the lifespan. Second, we focus on delay discounting, the tendency to prefer smaller rewards delivered soon over larger rewards delivered after a delay. Delay discounting has been used as a behavioral measure of impulsivity and is associated with many undesirable real-life outcomes. Specifically, we examined how neuroanatomy is associated with individual differences in delay discounting in a large adolescent sample. Using a novel multivariate method, we identified networks where cortical thickness varied consistently across individuals and brain regions. Cortical thickness in several of these networks, including regions such as ventromedial prefrontal cortex, orbitofrontal cortex, and temporal pole, was negatively associated with delay discounting. Furthermore, this brain data predicted differences beyond those typically accounted for by other cognitive variables related to delay discounting. These results suggest that cortical thickness may be a useful brain phenotype of delay discounting and carry unique information about impulsivity. Collectively, this research furthers our understanding of how cognitive abilities, brain structure and healthy aging relate to individual differences in value-based decision-making.

Degree Type

Dissertation

Degree Name

Doctor of Philosophy (PhD)

Graduate Group

Psychology

First Advisor

Joseph W. Kable

Subject Categories

Cognitive Psychology

INDIVIDUAL DIFFERENCES IN VALUE-BASED DECISION-MAKING:
LEARNING AND TIME PREFERENCE

Marieta Pehlivanova

A DISSERTATION

in

Psychology

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2017

Supervisor of Dissertation

Joseph W. Kable, Ph.D., Baird Term Associate Professor, Psychology

Graduate Group Chairperson

Sara R. Jaffee, Ph.D., Professor, Psychology

Dissertation Committee

Theodore D. Satterthwaite, M.D., M.A., Assistant Professor, Psychiatry

Russell A. Epstein, Ph.D., Professor, Psychology

Danielle S. Bassett, Ph.D., Associate Professor, Bioengineering

INDIVIDUAL DIFFERENCES IN VALUE-BASED DECISION-MAKING:

LEARNING AND TIME PREFERENCE

COPYRIGHT

2017

Marieta Pehlivanova

This work is licensed under the
Creative Commons Attribution-
NonCommercial-ShareAlike 3.0
License

To view a copy of this license, visit

<https://creativecommons.org/licenses/by-nc-sa/3.0/us/>

For Maika, Tatko, and Ina

ABSTRACT

INDIVIDUAL DIFFERENCES IN VALUE-BASED DECISION-MAKING:

LEARNING AND TIME PREFERENCE

Marieta Pehlivanova

Joseph W. Kable

Human decisions are strongly influenced by past experience or by the subjective values attributed to available choice options. Although decision processes show some common trends across individuals, they also vary considerably between individuals. The research presented in this dissertation focuses on two domains of decision-making, related to learning and time preference, and examines factors that explain decision-making differences between individuals. First, we focus on a form of reinforcement learning in a dynamic environment. Across three experiments, we investigated whether individual differences in learning were associated with differences in cognitive abilities, personality, and age. Participants made sequential predictions about an on-screen location in a video game. Consistent with previous work, participants showed high variability in their ability to implement normative strategies related to surprise and uncertainty. We found that higher cognitive ability, but not personality, was associated with stronger reliance on the normative factors that should govern learning. Furthermore, learning in older adults (age 60+) was less influenced by uncertainty, but also less influenced by reward, a non-normative factor that has substantial effects on learning across the lifespan. Second, we focus on delay discounting, the tendency to prefer smaller rewards delivered soon over

larger rewards delivered after a delay. Delay discounting has been used as a behavioral measure of impulsivity and is associated with many undesirable real-life outcomes. Specifically, we examined how neuroanatomy is associated with individual differences in delay discounting in a large adolescent sample. Using a novel multivariate method, we identified networks where cortical thickness varied consistently across individuals and brain regions. Cortical thickness in several of these networks, including regions such as ventromedial prefrontal cortex, orbitofrontal cortex, and temporal pole, was negatively associated with delay discounting. Furthermore, this brain data predicted differences beyond those typically accounted for by other cognitive variables related to delay discounting. These results suggest that cortical thickness may be a useful brain phenotype of delay discounting and carry unique information about impulsivity. Collectively, this research furthers our understanding of how cognitive abilities, brain structure and healthy aging relate to individual differences in value-based decision-making.

TABLE OF CONTENTS

| | |
|--|-------------|
| ABSTRACT..... | IV |
| LIST OF TABLES | VIII |
| LIST OF ILLUSTRATIONS..... | IX |
| CHAPTER 1 – INTRODUCTION..... | 1 |
| CHAPTER 2 — AGE AND COGNITIVE ABILITIES PREDICT LEARNING IN A DYNAMIC ENVIRONMENT | 18 |
| <i>Abstract.....</i> | <i>18</i> |
| <i>Introduction.....</i> | <i>19</i> |
| <i>Experiment 1: Method.....</i> | <i>22</i> |
| <i>Experiment 1: Results</i> | <i>33</i> |
| <i>Experiment 1: Discussion</i> | <i>36</i> |
| <i>Experiment 2: Method.....</i> | <i>38</i> |
| <i>Experiment 2: Results and Discussion.....</i> | <i>41</i> |
| <i>Experiment 3: Method.....</i> | <i>45</i> |
| <i>Experiment 3: Results and Discussion.....</i> | <i>47</i> |
| <i>Group differences between young and elderly participants</i> | <i>49</i> |
| <i>General Discussion.....</i> | <i>52</i> |
| CHAPTER 3 — DIMINISHED CORTICAL THICKNESS IS ASSOCIATED WITH IMPULSIVE CHOICE IN ADOLESCENCE..... | 63 |
| <i>Abstract.....</i> | <i>63</i> |
| <i>Significance.....</i> | <i>64</i> |
| <i>Introduction.....</i> | <i>64</i> |
| <i>Methods.....</i> | <i>67</i> |

| | |
|---|------------|
| <i>Results</i> | 75 |
| <i>Discussion</i> | 79 |
| CHAPTER 4 — GENERAL DISCUSSION | 92 |
| BIBLIOGRAPHY | 102 |

LIST OF TABLES

| | |
|--|----|
| 2.1 Summary of regression results for three models predicting NLF in Experiment 1 ... | 55 |
| 2.2 Summary of regression results for three models predicting NLF in Experiment 2 ... | 56 |
| 2.3 Summary of regression results for three models predicting NLF in Experiment 3 ... | 57 |
| 3.1 Association between delay discounting and NMF-derived structural covariance networks | 85 |

LIST OF ILLUSTRATIONS

| | |
|--|----|
| 2.1 Overview of helicopter task and normative learning factors | 58 |
| 2.2 Experiment 1 results: normative learning behavior and its associations with IQ | 59 |
| 2.3 Experiment 2 results: normative learning behavior and its associations with memory | 60 |
| 2.4 Experiment 3 results: normative learning behavior and its associations with memory | 61 |
| 2.5 Group differences between young and older participants in updating based on normative factors and the incidental reward factor | 62 |
| 3.1 Schematic of non-negative matrix factorization and example data for each matrix ... | 86 |
| 3.2 NMF reconstruction error plot | 87 |
| 3.3 Structural covariance networks delineated by NMF | 88 |
| 3.4 Significant associations between delay discounting and cortical thickness | 89 |
| 3.5 Age effects on cortical thickness | 90 |
| 3.6 Multivariate prediction of delay discounting | 91 |

CHAPTER 1 – Introduction

Value-based decision-making: examples and main effects vs. individual differences

In life, people make a multitude of decisions in the span of a single day, week, month, or even across many years. Some of these decisions involve choices rooted in *preferences*, such as choosing between spending available money now or saving for retirement or future consumption (time preference); or choosing between a risky but possibly profitable career move versus a less risky but also less beneficial alternative (risk preference); or choosing between an option of narrow-self interest and an option that may benefit another person or group (social preference). Alternatively, other decisions are more strongly rooted in *learning* from past experience, through trial and error, such as choosing a driving route based on previously experienced traffic and delays on specific roads. A large effort in the research field of decision-making has focused on studying the main effects of choices under different conditions, i.e. elucidating the behavioral principles and neural mechanisms of how people make decisions. To this end, relevant measures such as response rates, reaction times, or neural activation in response to task stimuli are typically averaged between different experimental conditions, while inter-individual variability is treated as idiosyncratic noise (Kanai & Rees, 2011). For example, people tend to view losses more unfavorably than gains of the same magnitude, a phenomenon called loss aversion (Tversky & Kahneman, 1992).

Yet, one of the most consistent and salient observations in both everyday life and scientific inquiry is that people differ from each other in a variety of ways. In everyday

life, some differences are easily and immediately observable (for example, height), while others take time to uncover (e.g. personality or cognitive ability), and others may yet be inconspicuous to the naked eye (e.g. brain size). Within the scientific study of decision-making, research has shown, for example, that people vary widely in the level of risk they accept in life (Weber, Blais, & Betz, 2002) or in the degree to which they can delay gratification (Mischel, Shoda, & Rodriguez, 1989). The research presented in this dissertation aims at furthering our understanding of cognitive and neuroanatomical features that relate to individual differences in decision-making, specifically in the domains of learning and time preference.

Sources of individual differences in differential psychology

Differences between individuals that relate to psychological and behavioral characteristics have been studied within the purview of differential psychology. Here we consider the following major domains of individual differences between humans: cognitive abilities, personality and age (Kanai & Rees, 2011; Lubinski, 2000). One of the most consequential dimensions of human variability is general intelligence, a factor that subsumes variability in reasoning, mental processing speed, executive function and memory, and is measured by the intelligence quotient (IQ). Higher IQ is associated with higher educational achievement, successful job performance and functioning in modern life (Gottfredson, 1997) and decreased mortality (Batty, Deary, & Gottfredson, 2007), among other beneficial outcomes. Personality is another dimension of variability, with one prominent model (“Big Five”) identifying five traits along which humans vary: neuroticism, conscientiousness, agreeableness, extraversion, and openness (Costa &

McCrae, 1992). As measured by this model, personality is associated with differences in job performance (Barrick & Mount, 1991), health behaviors (McAdams & Donnellan, 2009; Rhodes & Smith, 2006), and even political views (Gerber, Huber, Doherty, & Dowling, 2011) and predicts mortality, divorce, and occupational attainment at least as well as cognitive ability does (Roberts, Kuncel, Shiner, Caspi, & Goldberg, 2007). Finally, both cognition and personality change across the lifespan (Craig & Bialystok, 2006; Specht, Egloff, & Schmukle, 2011), and healthy aging is generally associated with changes in decision-making, resulting from changes in neurophysiology, cognitive abilities and affective motivation (Samanez-Larkin & Knutson, 2015).

Over the past couple of decades, differential psychology has increasingly begun to examine the brain anatomy and function underlying psychological differences among individuals. This process has been facilitated by advances in neuroimaging and further reinforced by discoveries that the individual difference factors described above may have a neurobiological basis and a genetic component, thus giving rise to the field of personality neuroscience (DeYoung & Gray, 2009). Indeed, general intelligence is linked to structural brain differences in frontal, temporal, and parietal lobes—where more gray matter or thicker cortex is associated with higher IQ— and is highly heritable (Deary, Penke, & Johnson, 2010; Haier, Jung, Yeo, Head, & Alkire, 2004). The study of personality has also yielded structural correlates of distinct Big Five traits (DeYoung et al., 2010) and evidence of heritability (Riemann, Angleitner, & Strelau, 1997). Among the above factors, perhaps the most significant source of neural variability is associated with age, as the human brain undergoes dramatic changes through childhood and

adolescence (Giedd et al., 1999; Lenroot et al., 2007) to healthy old age (Buckner, 2004; Raz et al., 2005).

The following section will discuss cognition, personality, and age as sources of individual differences in specific decision-making tasks related to learning and time preference.

Reinforcement learning and individual differences

Brief background on reinforcement learning

Much of human decision-making is based on learning from experience, through trial and error, so as to avoid unpleasant outcomes and achieve desirable ones. This phenomenon is known as *reinforcement learning* and has been robustly observed across species beyond humans. Early explorations of such learning came from the field of animal behavior. In a classical conditioning paradigm, an animal was repeatedly presented with a pairing of an initially neutral and unconditioned stimulus and a reward, such as food, which produces a natural response from the animal (Daw & Tobler, 2013; Pavlov, 1927). After repeated exposure to this pairing, the animal develops an implicit response to the (now conditioned) stimulus, even when the stimulus is presented without the reward. It was proposed that the animal learns the conditioned response based on comparisons between the observed reward and the expected/predicted reward based on prior experience (Bush & Mosteller, 1951). The difference between the two is termed the *prediction error*, and learning is greatest when that error is large. This observation gave rise to the Rescorla-Wagner model of classical conditioning (Rescorla and Wagner,

1972). According to this model, the animal updates its predictions in the direction of the prediction error and the extent of updating is governed by a parameter called the *learning rate* (constrained between 0 and 1, where larger values indicate updating towards more recent outcomes). Most of the work on the neural basis of reinforcement learning computations has focused on prediction errors. Prediction error-like signals have been associated with phasic activity of the midbrain dopamine neurons (Glimcher, 2011; Schultz, Dayan, & Montague, 1997); over the course of learning stimulus-reward associations, dopamine neurons gradually become responsive to the reward-predicting stimulus, rather than the reward itself (Schultz et al., 1997).

Adaptive learning

In this dissertation, we will focus on a specific form of reinforcement learning that takes into account abrupt changes in the environment. Recent work has begun to characterize the principles of learning specifically in a changing environment (Behrens et al., 2007; Nassar et al., 2010; McGuire et al., 2014). We refer to such learning as “adaptive” because it allows and requires *adapting* one’s rate of learning to the type of environment. In a stable but noisy environment, long-run experience is the best predictor of the future, and it is reasonable to maintain a low learning rate favoring the averaged experienced outcomes (Behrens et al., 2007; Nassar et al., 2010). Conversely, in a volatile environment, characterized by unexpected changes, recent experience is the best predictor of the future, and it is reasonable to use a high learning rate such that the most recently experienced outcome is weighted more heavily. In such learning paradigms, people are generally sensitive to changes in the environment— they can detect volatility

and respond appropriately by using higher learning rates— but these adaptive tendencies vary widely across individuals (Behrens et al., 2007; McGuire et al., 2014; Nassar et al., 2010).

Individual differences in reinforcement learning

In this section, I will review literature on individual difference factors, as outlined above, that have been found to explain variability in reinforcement learning.

Cognitive factors. Some work has been done in a framework that distinguishes between two learning algorithms: “model-based” versus “model-free” learning. In this framework, model-free learning relies on simply repeating actions that have previously been rewarded without explicitly learning the structure of the reward environment (Daw, Niv, & Dayan, 2005). This is the type of learning mechanism that is encoded in dopaminergic neurons’ responses. In contrast, model-based learning relies on building a sophisticated mental model of the reward environment, evaluating choices in the context of that model, and flexibly updating one’s learning rate depending on the circumstances. Otto and colleagues (2013, 2014) have recently shown that better working memory and cognitive control are associated with increased reliance on model-based rather than model-free learning. In addition, high working memory capacity protects individuals against the deleterious effect of stress on model-based learning (Otto et al., 2013).

Neural factors. Some work has also identified a neural basis of individual differences in reinforcement learning, specifically linked to the striatum, which is one of the main target areas of dopamine neurons coding for prediction errors. In a reward

learning task, individuals who were able to learn the reward contingencies of the task showed pronounced striatal activation in response to prediction errors, relative to those who were unable to learn, and the magnitude of activation was positively associated with performance across both groups (Schönberg, Daw, Joel, & O'Doherty, 2007). Another study has directly linked individual differences in learning to differences in baseline striatal dopamine synthesis capacity, such that individuals with high and low capacity show distinct patterns of responses to learning from rewards versus punishments in a reversal learning task (Cools et al., 2009).

In adaptive learning specifically, prior work has reported psychophysiological and neural correlates of individual variability in learning. Pupil diameter tracks normative learning factors derived from a computational model, and the degree of pupil metrics' sensitivity to environmental statistics reflects the extent to which a person's behavior was influenced by these normative factors (Nassar et al., 2012). In an fMRI study, activity in the anterior insula and dorsomedial frontal cortex, regions linked to arousal and salience (Seeley et al., 2007), is modulated by these normative factors (McGuire et al., 2014). Furthermore, the extent to which activity in anterior insula and dorsomedial frontal cortex was modulated by these normative factors was associated with the extent to which a person's behavior was influenced by these same factors (McGuire et al., 2014).

Personality. The literature on personality factors that explain individual variability in learning behavior is sparse. In particular, extraversion— a personality trait that has been hypothesized to relate to dopamine function— is associated with individual differences in learning, as measured by distinct EEG response patterns (Cooper, Duke,

Pickering, & Smillie, 2014; Smillie, Cooper, & Pickering, 2011). In an associative learning task, more extraverted individuals show a greater difference in sensitivity between unpredicted non-rewards and predicted rewards, compared to introverted individuals (Cooper et al., 2014; Smillie et al., 2011).

Age. Reinforcement learning processes show age-related differences, with children and older adults generally able to learn reward contingencies, but learning more slowly and with more difficulty than adolescents and younger adults (Eppinger, Kray, Mock, & Mecklinger, 2008; Eppinger, Mock, & Kray, 2009; Hämmerer, Li, Müller, & Lindenberger, 2011). In old adults, some of these differences have been hypothesized to result from age-related changes in the dopaminergic and serotonergic systems (Marschner et al., 2005). Healthy aging is specifically associated with decreased sensitivity to reward in a probabilistic learning task, which is neurobiologically mediated by reduced white matter integrity in select prefrontal pathways (Samanez-Larkin et al., 2012).

Important open questions in individual differences in adaptive learning

As outlined above, most of the work on individual differences in reinforcement learning has been done in paradigms that do not explicitly model fundamental changes in the environment. Differentiating between variability due to changes as apposed to noise is an important skill, and people vary in their degree of adaptability in dynamic environments— some people are quick learners in the face of change, while others require more time and experience to adapt. Given that most realistic environments humans experience, such as relationships, the work place or financial markets, are rarely stable over time, understanding factors that explain variability in individuals' decision-

making in dynamic environments is important. Previous individual difference work in the adaptive learning framework has identified psychophysiological (Nassar et al., 2012) and neural correlates (McGuire et al., 2014) of variability. However, differences have not been examined in the context of other factors that are generally known to relate to inter-individual variability, such as cognitive and personality factors. In addition, healthy aging could be an important source of individual differences, as it is associated with changes in personality (Roberts & Mroczek, 2008) and cognition (McArdle, Ferrer-Caja, Hamagami, & Woodcock, 2002), possibly mediated by neurological age-dependent changes (Marschner et al., 2005). Alternatively, age may play a role in adaptive learning independently of changes in the above factors. The systematic study of these factors together could be fruitful, as it will allow us to study the adaptive learning process in its complexity. Accordingly, here we extend the study of individual differences in adaptive learning by investigating the effect of cognitive abilities, personality, and age on learning strategies.

Time preference: delay discounting and individual differences

In adaptive learning, the decision-maker is influenced by the past, and recent outcomes compete with more distant ones to shape current expectations about the environment. Next, we transition to another domain of decision-making— time preference— where decisions are in turn based on the influence of the recent or immediate versus distant future.

Introduction to delay discounting

In the domain of preference, this dissertation specifically focuses on time preference which entails a choice between receiving something desirable soon and receiving something desirable after some time. This type of decision is common in real life and consists of a comparative evaluation of costs and benefits of options occurring at different points in time. For example, a student has a choice between playing video games now and using that time to study for the SAT, which might lead to a larger return on investment in the future; an employee has a choice between using their discretionary income for hobbies or leisure now or investing that income for larger future monetary returns. Such intertemporal choices have been studied in the laboratory using a paradigm called “delay discounting” (DD), which measures the degree to which people “discount” the value of rewards received after a delay (Kirby, Petry, & Bickel, 1999). The DD task consists of a series of questions of the type “Would you rather receive \$10 now or \$18 in 30 days?” By varying the amounts and delays, one can estimate the subjective discount rate at which each individual devalues future outcomes. The discount rate (DR) estimates the steepness of the reduction of present value with increases in delays. Individuals with higher DRs are considered more impulsive, while individuals with low DRs are considered patient. DD is particularly fruitful for the study of individual differences for the following reasons: First, there exists large variability between individuals in the degree of discounting. Second, behavioral variability on this lab task predicts variability in many real-life behaviors, including addiction. Third, DD has been used extensively as a measure of impulsivity.

Delay Discounting, real-life outcomes, and relation to impulsivity

A large body of literature has shown that discount rates measured in the lab correlate with consequential real-life behaviors and outcomes. Higher discounting is associated with relationship infidelity (Reimers, Maylor, Stewart, & Chater, 2009), lower creditworthiness (Meier & Sprenger, 2012), and poor health habits and obesity (Amlung, Petker, Jackson, Balodis, & MacKillop, 2016; Chabris, Laibson, Morris, Schuldt, & Taubinsky, 2008). Conversely, lower discounting is associated with greater life satisfaction (Becker, Deckers, Dohmen, Falk, & Kosse, 2012) greater social competence in adolescents (Mischel et al., 1989), and higher GPA in college students (Kirby, Winston, & Santiesteban, 2005). The largest literature on the real-life correlates of delay discounting comes from the field of addiction research (see MacKillop et al., 2011 for a meta-analysis and Reynolds, 2006 for a qualitative review). Higher discounting has been reliably associated with smoking (Epstein et al., 2003; Reynolds, 2004; Reynolds, Richards, Horn, & Karraker, 2004); with higher and/or problematic alcohol use (Courtney et al., 2012; Petry, 2001a; Vuchinich & Simpson, 1998); with heroin and cocaine addiction (Coffey, Gudleski, Saladin, & Brady, 2003; Kirby et al., 1999; Kirby & Petry, 2004; Madden, Petry, Badger, & Bickel, 1997); and with pathological gambling (Petry, 2001b). In the aggregate, greater discounting of future rewards appears to be robustly related to less desirable outcomes, while low discounting appears to be associated with more desirable outcomes.

Impulsivity is defined as a tendency to act without deliberation or adequate regard for consequences (Evenden, 1999). DD has been proposed as a model of impulsivity, where steady preference for immediate rewards is considered impulsive whereas steady

preference for delayed rewards is interpreted as a display of self-control (Ainslie, 1975). DD has also specifically been proposed as a framework for understanding substance abuse as an impulsive choice of an immediate reward over the more patient and beneficial choice of abstinence (Bickel et al., 2007).

Individual differences in delay discounting

In this section, I will review individual difference factors that have been found to explain variability in discount rates among people.

Cognitive factors. In relation to cognitive factors, lower discounting is robustly associated with higher intelligence (Shamosh and Gray, 2008), and also with better working memory (Shamosh et al., 2008; Bobova et al., 2009). As a potential neural mechanism of these associations, Shamosh and colleagues (2008) found that working memory-related activity in anterior prefrontal cortex partially mediates the relationship between DD and IQ.

Personality. In contrast, relationships between personality traits and DD appear to be less stable. High extraversion (characterized by an orientation towards people and external events), high agreeableness (characterized by a tendency to act in a cooperative manner), and low conscientiousness (characterized by a tendency to be less organized and responsible) have been reported to predict steeper discounting (Anderson, Burks, DeYoung, & Rustichini, 2011; Becker et al., 2012; Daly, Harmon, & Delaney, 2009; Ostaszewski, 1996), though replication of these effects has been inconsistent (Becker et al., 2012; Dohmen et al., 2010). Interestingly, some research has also examined the

interaction of cognitive abilities and personality in predicting discount rates (Hirsh, Morisano, & Peterson, 2008). At the low end of the cognitive ability distribution, high extraversion and low neuroticism, defined as emotional stability, predict higher discounting, while at the high end extraversion has no effect and high neuroticism is associated with steeper discounting (Hirsh et al., 2008). Traits related to time perception also appear to play a role in temporal discounting (Kim & Zauberman, 2009): feeling weak connectedness to one's future self (Ersner-Hersfield et al., 2009a; 2009b) is associated with steeper discounting. Both differences in time perception and future connectedness have been associated with differences in brain activity (Cooper, Kable, Kim, & Zauberman, 2013; Ersner-Hersfield et al., 2009b).

Age. Multiple studies have examined age as a source of individual differences in DD. One of the earliest such investigations reported that young adults were lower discounters than children but higher discounters than older adults, suggesting a linear decrease in discounting across age groups (Green, Fry, & Myerson, 1994). Despite these differences, though, the mechanism of choice in DD appears to be qualitatively similar across age groups (Green et al, 1994). However, reports about discounting differences between young and old adults are inconsistent. In addition to several studies reporting that older adults are lower discounters compared to young adults (Eppinger, Nystrom, & Cohen, 2012; Green et al., 1994; Reimers et al., 2009), there is evidence of old adults being higher discounters (Read & Read, 2004), as well as reports of no differences between these age groups (Chao, Szrek, Pereira, & Pauly, 2009; Samanez-Larkin et al., 2011). Notably, Read and Read (2004) reported a curvilinear relationship between adult

age and DD, such that elderly adults (average age of 75) discounted more steeply than young adults (average age of 25), with middle-aged adults (average age of 44) exhibiting the lowest discount rates. In youth, age differences have been reported, such that younger adolescents are higher discounters than older adolescents (Olson, Hooper, Collins, & Luciana, 2007; Scheres et al., 2006; Steinberg et al., 2009).

Brain structure. A number of studies with adults have investigated individual differences in DD in relation to neuroanatomy. Negative associations between gray matter volume and discounting have been found in lateral prefrontal cortex (Bjork et al., 2009), superior frontal gyrus (Schwartz et al., 2010), putamen (Dombrovski et al., 2012; Cho et al., 2013); and positive associations have been found in posterior cingulate cortex (PCC) and ventral striatum (Schwartz et al., 2010), medial prefrontal regions and anterior cingulate cortex (ACC, Cho et al., 2013), middle frontal gyrus and frontal pole (Wang et al., 2016). Decreased CT in medial prefrontal cortex and ACC is associated with higher discounting (Bernhardt et al., 2014). Generally, these results have been inconsistent in terms of directionality and regional specificity of effects, and might have been limited by small or clinical samples, and region-of-interest analyses (for a review see Kable & Levy, 2015).

Important open questions in individual differences in delay discounting

The neurofunctional correlates of DD have been characterized extensively (Bartra, McGuire, & Kable, 2013; Clithero & Rangel, 2013; Kable & Glimcher, 2007; Peters & Büchel, 2011), and some research has begun to investigate its neurostructural correlates.

However, most of the previous work relating brain structure to DD has been conducted with adult samples, and it is unclear whether documented age differences in DD might also contribute to differences in neurostructural correlates. Such differences are especially important to consider in relation to adolescence because this developmental period is characterized by dramatic structural brain changes which may contribute to increased impulsivity (Giedd et al., 1999; Sowell et al., 2004; Van Leijenhorst et al., 2010), as well as changes in cognitive abilities, such as intelligence and executive function (Blakemore & Choudhury, 2006; McArdle, Ferrer-Caja, Hamagami, & Woodcock, 2002), which are associated with DD. In addition, delay discounting, as a behavioral measure of impulsivity, has been associated with many risky behaviors, such as reckless driving and alcohol abuse, which are an acute source of morbidity in adolescence (Eaton et al., 2011). Here we extend the study of individual differences in DD by investigating how differences in neuroanatomy, specifically cortical thickness, relate to variability in discounting in adolescence.

Research overview

This dissertation specifically focuses on investigating individual differences in two types of value-based decision-making processes: adaptive learning and delay discounting. Both of these processes involve processing time-dependent information: in learning, the decision-maker incorporates information from the past to accurately estimate the current expectations about the environment, while in delay discounting the decision-maker incorporates information about delays into the future. Conversely, the processes differ in that adaptive learning is based on experienced outcomes, whereas

delay discounting is based on preference rather than experience. Understanding the factors that contribute to individual differences in these processes is important because successful functioning in the world depends on the abilities to detect changes in one's environment, delay gratification and accumulate resources for future consumption.

In Chapter 2, in a series of three behavioral experiments, we investigate whether individual differences in adaptive learning are associated with differences in cognitive abilities, personality, and age. In a simple video game task, participants make sequential predictions about an on-screen location. This task provides trial-by-trial learning rate estimates and has been successfully used in eliciting individual differences in learning behavior (Nassar et al., 2010, 2012, 2016; McGuire et al., 2014). In addition, the task has an underlying computational model, which has characterized two distinct normative factors that should drive learning: *change-point probability*, which is related to how surprising and unexpected a new outcome is, and *relative uncertainty*, which tracks the reliability of the current beliefs about the state of the environment (Wilson et al., 2010; Nassar et al., 2010, 2012). Consistent with previous work, we find that participants exhibit common strategies in adaptive learning but also show large individual differences in implementing these strategies. We find that increased cognitive ability is associated with stronger reliance on the normative factors that should govern learning in this task. In contrast, we find no reliable evidence that personality traits, including trait anxiety, influence learning strategies in this dynamic environment. Furthermore, we find that adaptive learning in older adults (age 60+) is less influenced by relative uncertainty, but

also less influenced by reward, a non-normative factor that has substantial effects on adaptive learning across the lifespan.

In Chapter 3, we examine neuronatomy, specifically cortical thickness, as a source of individual differences in delay discounting in large sample of adolescents. Using a novel multivariate method, we describe networks where cortical thickness varies consistently across individuals and brain regions. Cortical thickness in several of these networks of regions shows a negative relationship with impulsivity, such that diminished cortical thickness is associated with greater discounting. The strongest effects were found in regions typically implicated in delay discounting, such as ventromedial prefrontal cortex and orbitofrontal cortex. Brain data predicted differences in discounting above and beyond cognitive variables typically found to correlate with delay discounting. Combined, these results suggest that cortical thickness may be a useful brain phenotype of delay discounting and carry unique information about impulsivity.

CHAPTER 2 — Age and cognitive abilities predict learning in a dynamic environment

Abstract

People are adaptive learners in a changing environment and are able to update their beliefs using relevant cues about the state of the world. Despite common trends in behavior, there is great variability in participants' learning strategies. Across three experiments, we investigate cognitive abilities, personality, and age as potential sources of these individual differences. In a simple video game task, participants made sequential predictions about an on-screen location. This task provides estimates of trial-by-trial learning rates and has been successfully used in eliciting individual differences in learning behavior. In addition, the task has an underlying computational model, which has characterized two distinct normative factors that should drive learning: *change-point probability*, which is related to how surprising a new outcome is, and *relative uncertainty*, which tracks the reliability of the current beliefs about the state of the environment. Consistent with previous work, we found that participants exhibit common strategies in adaptive learning but also show large individual differences in implementing these strategies. We found that increased cognitive ability, and specifically better memory performance, is associated with stronger reliance on the normative factors that should govern learning in this task. In contrast, we found no reliable evidence that personality traits influence learning strategies in a dynamic environment. Furthermore, adaptive learning in older adults (age 60+) was less influenced by relative uncertainty,

consistent with previous reports, but also less influenced by reward, a non-normative factor that has substantial effects on adaptive

Introduction

Imagine dining at the same restaurant on repeated occasions and having delicious meals, making it your favorite place to dine. Unexpectedly, on one particular occasion they serve you a bad meal. Is this negative experience just an unfortunate exception in an otherwise stable but noisy world, or does it indicate that something fundamental about the restaurant has changed, such as a new chef? How much you should update your beliefs in response to new experiences depends on the stability of the environment. In an environment that is stable but noisy, long-run experience is the best predictor of the future, and beliefs should be updated minimally in the face of new evidence (Behrens, Woolrich, Walton, & Rushworth, 2007; Nassar, Wilson, Heasly, & Gold, 2010). In this case, if a restaurant has a reliable and stable history with the same chef, it is sensible to disregard the occasional negative experience if you have had a steady streak of good meals there. Conversely, in a volatile environment, characterized by unexpected changes, recent experience is the best predictor of the future, and beliefs should be updated more rapidly in the face of new data (Behrens et al., 2007; Nassar et al., 2010). In this case, if a restaurant has gone through many chef changes recently, a surprisingly negative experience is more likely to mean that another change has occurred. Recent work has shown that, while on average people follow these strategies, individuals differ substantially in the extent to which they adhere to these principles of “adaptive learning”

(Behrens et al., 2007; McGuire, Nassar, Gold, & Kable, 2014; Nassar et al., 2010).

However, the different factors that might influence these individual differences in adaptive learning have not been systematically studied. Here we investigate cognitive abilities, personality, and age as possible sources of these individual differences.

A learning model derived from Bayesian theory has characterized two distinct normative factors that should drive learning in volatile environments, and individual differences in adaptive learning can be described by variability in responding to these normative factors (McGuire et al., 2014; Nassar, Bruckner, Gold, Li, Heekeren, & Eppinger, 2016; Nassar, Rumsey, Wilson, Parikh, Heasly, & Gold, 2012; Nassar et al., 2010; Wilson, Nassar, & Gold, 2010). The first factor, which we call *change-point probability*, tracks the likelihood of a fundamental shift in the environment and is related to how surprising and unexpected a new outcome is (Nassar et al., 2012). In the restaurant example, a meal of unexpected quality, relative to one's experience, might suggest that something important about the restaurant has changed. The second factor, which we call *relative uncertainty*, tracks the uncertainty, and thus reliability, of the current beliefs about the state of the environment, and is related to the number of observed outcomes consistent with the current state (Nassar et al., 2012). In the restaurant example, the longer the streak of delicious meals, the more confident you are of the quality of the meals under the current chef. Though on average individuals update their beliefs in a dynamic environment according to both of these factors, there is large inter-individual variability in doing so (Nassar et al., 2012; McGuire et al., 2014).

What accounts for these individual differences in sensitivity to relevant environmental cues? To date, studies have identified associations between individual differences in adaptive learning and psychophysiological and neural activity linked to arousal systems. Pupil diameter reflects both normative factors described above, with evoked changes in pupil diameter tracking change-point probability and average baseline pupil diameter tracking relative uncertainty (Nassar et al., 2012). Furthermore, the degree of pupil metrics' sensitivity to environmental statistics predicted the extent to which a person's behavior was influenced by the normative factors (Nassar et al., 2012). In an fMRI study, activity in the anterior insula and dorsomedial frontal cortex, regions linked to arousal and salience (Seeley et al., 2007), was modulated by both change-point probability and relative uncertainty (McGuire et al., 2014). Furthermore, the extent to which activity in anterior insula and dorsomedial frontal cortex was modulated by these normative factors was associated with the extent to which a person's behavior was influenced by these same factors (McGuire et al., 2014).

Here we extend these findings in a series of three behavioral experiments investigating whether individual differences in adaptive learning are associated with differences in cognitive abilities, personality, and age. There are strong reasons to hypothesize that each of these classic individual difference variables might influence adaptive learning. Cognitive abilities are associated with individual differences in many decision-making and learning tasks (Burks, Carpenter, Goette, & Rustichini, 2009; Otto, Skatova, Madlon-Kay, & Daw, 2014; Shamosh & Gray, 2008) but have not been explicitly studied in adaptive learning. Given that sensitivity to relevant cues from the environment is associated with arousal responses (Nassar et al., 2012), personality and

trait-based affect may influence adaptive learning. Indeed, a recent report links high trait anxiety with a reduced ability to adjust learning between stable and volatile environments (Browning, Behrens, Jocham, O'Reilly, & Bishop, 2015). Healthy aging is also associated with changes in decision-making, resulting from changes in neurophysiology, cognitive abilities and affective motivation (Samanez-Larkin & Knutson, 2015). In particular, Nassar et al. (2016) have shown that learning from uncertainty in a dynamic environment is reduced in older, relative to younger, adults. Here we systematically explore the interplay of cognitive abilities, personality, and aging on adaptive learning.

Consistent with previous reports (Nassar et al., 2010, 2012; McGuire et al., 2014), we found that people exhibited common strategies in adaptive learning, but also showed large individual differences in implementing these strategies. We found that increased cognitive ability, and specifically better memory performance, is associated with stronger reliance on the normative factors that should govern learning in this task. In contrast, we found no reliable evidence that personality traits, including trait anxiety, influence learning strategies in a dynamic environment. Finally, we found that adaptive learning in older adults (age 60+) is less influenced by relative uncertainty, consistent with previous reports (Nassar et al., 2016), but also less influenced by reward, a non-normative factor that has substantial effects on adaptive learning across the lifespan.

Experiment 1: Method

Participants. The experimental protocol was approved by the University of Pennsylvania Internal Review Board. Informed consent was obtained from all

participants. We recruited 49 participants for this experiment (age 18–34, $M = 25.5$, $SD = 4.7$; 33% males) from the University of Pennsylvania and the surrounding local community. Inclusion criteria included comfort with using a right-handed trackball and fluency in English. Exclusion criteria included a history of alcohol or drug abuse, major psychiatric disorders not in remission for > 6 months, and current use of psychotropic medications (including antidepressants, anxiolytics, and antipsychotics).

Procedure. Participants first completed a variant of a predictive inference task that we have previously used to study adaptive learning, the “helicopter task,” which is described in more detail below. Subsequently, participants’ IQ was assessed with the similarities and matrix reasoning subtests of the Wechsler Abbreviated Intelligence Scale (WASI; Wechsler, 1999). Finally, participants completed the Big Five Inventory-44 (John, Donahue, & Kentle, 1991), Brief Sensation Seeking Scale (Hoyle, Stephenson, Palmgreen, Lorch, & Donohew, 2002), and the State-Trait Anxiety Inventory (STAI) trait subscale (Spielberger, Gorsuch, Lushene, Vagg, & Jacobs, 1983) via Qualtrics (Qualtrics Labs Inc., Provo, UT). We also administered an intertemporal choice task (Senecal, Wang, Thompson, & Kable, 2012) for exploratory purposes, and do not describe the results of that task here. All tests were administered by the same experimenter (MP), in the same manner, and in the same testing room. One notable exception to the experimental protocol is that we started administering the STAI trait subscale mid-data collection upon becoming aware of Browning et al.’s (2015) reported anxiety effects on adaptive learning. We retroactively collected the anxiety data for the first half of our sample via an online survey.

Helicopter Task. Participants performed four blocks (150 trials each) of a predictive-inference task in which learning rates can be tracked trial-by-trial (Nassar et al., 2010, 2012, 2016; McGuire et al., 2014). The task was programmed in Matlab (The MathWorks, Natick, MA) using MGL (<http://justingardner.net/mgl>) and snowDots (<https://code.google.com/archive/p/snow-dots/>) extensions. On every trial, participants made a prediction about the mean of an underlying generative distribution based on observed random draws from that distribution. The task was presented as a computer game with the following description: a helicopter in the sky drops a bag of coins on each trial; once the bag reaches the ground, it turns into a coin explosion, and the participant's job is to position a bucket on the ground prior to the bag fall to collect as many coins as possible (Figure 1A). Importantly, the helicopter is hidden in clouds, so participants must use previously observed bag drop locations to infer the helicopter's position. Participants made trial-by-trial predictions about the location of the helicopter by positioning the rectangular bucket using a trackball. At the beginning of every trial, the bucket was repositioned to its default position in the middle of the screen. Once a participant submitted their prediction, they could not move the bucket until the bag-drop outcome was realized. Participants were instructed that the best strategy to maximize earnings is to position the bucket directly underneath the helicopter, even when the helicopter is obscured by clouds, thus implicitly discouraging guessing of individual bag drop locations. Participants received feedback on every trial in the form of a red bar spanning the difference between the bag drop location and the selected bucket position, indicating the *prediction error* (PE). This feedback was added to minimize the working memory burden

in this inference problem. The helicopter location changed occasionally, without warning, giving rise to unsignaled *change points* (CPs).

Possible bag drop and bucket locations were mapped onto arbitrary screen units valued from 0 to 300. Bag drop locations were drawn at random from a Normal distribution around the center of the helicopter. The mean of the distribution, and thus the helicopter position, was held fixed on most trials (with 100% probability for the first 3 trials in a new location). Occasionally there were unsignaled change points: with a .125 probability after the third stable trial, the mean of the generative distribution was re-drawn from a uniform distribution on (0, 300). The standard deviation of the distribution of bag drops varied between blocks and was set to either 10 or 25 screen units, creating different levels of noise in the environment. Participants were instructed to think of this noise as different strengths of “wind” which occasionally blew the bags away from the center of the helicopter. The width of the bucket scaled with the strength of the wind, and was set to 2.5 and 2.2 times the *SD* of the bag drop distribution, in the low and high wind conditions, respectively. The width of the bucket was chosen to balance earnings between different noise blocks, and the set of other parameter values were chosen based on simulations seeking to maximize the difference in payoff between an approximately Bayesian learning model (described below) and an observer that uses the simple strategy of placing the bucket at the location of the last bag drop. The main part of the task consisted of four blocks (150 trials each) alternating in wind strength, with the type of the first block (low vs. high wind/noise) counterbalanced across participants.

The task also included a manipulation of reward with each bag containing one of two colors of coins, determined independently at random with 50% probability on each

trial. For half the participants, yellow coins were rewarded and gray coins were unrewarded. For the remaining participants, the colors were flipped with gray coins being rewarded and yellow coins being unrewarded. The rewarded color was also counterbalanced across participants and participants were assigned to one of the four conditions obtained by crossing the factors initial noise level and rewarded color. Participants were told which color would be rewarded, and were specifically instructed that (1) they will not see the coin color until after they have made a bucket prediction and observed the bag fall; (2) the color of the coin should not influence their strategy in the game. All rewarded coins gathered were redeemed for money. Participants received feedback after each block in the form of a percentage score indicating the participant's earnings relative to earnings obtained by the approximate Bayesian learning model described in the following section.

Prior to the four experimental blocks, participants completed a training session that was more extensive than in earlier versions of the task (Nassar et al., 2012; McGuire et al., 2014). The training built understanding of the environment by gradually introducing more complexity, such as wind, clouds and changes in the helicopter's position. In the first training block, the helicopter was visible, fixed and the bag always dropped right underneath the helicopter (i.e. no wind); participants were simply required to repeatedly position the bucket underneath the helicopter to ensure they can manipulate the trackball. This first block was performed to criterion: 15 consecutive trials where the center of bucket was positioned within 12 screen units of the helicopter position. In the second training block of 25 trials, a distribution of bag drops with $SD = 10$ units around the helicopter center was introduced and explained to participants as "wind." The third

block was 35 trials of a “windier” condition with $SD = 25$ units and a larger bucket, conditions that were maintained throughout the remainder of the training session. In the fourth block of 40 trials we introduced clouds on about half the trials (randomly determined on each trial); the clouds obscured the stationary helicopter from view. In the following two blocks of 50 trials each it was cloudy at all times, so participants never saw the helicopter and had to infer its position based on observed bag drops. In the next block of 50 trials the clouds were fully removed to make visible occasional shifts of the helicopter (i.e., change points were added). In the final training block of 60 trials, participants experienced high wind, clouds at all times, and occasional change points, and the only difference from the experimental blocks was that there was only one type of coin. The unrewarded coins of different color were introduced in the main experimental blocks. Participants did not earn money from the training session.

Normative Model. We used a well-described approximately Bayesian belief updating model to simulate (nearly) optimal performance in the helicopter task (Nassar et al., 2010, 2012, 2016; McGuire et al., 2014). The model computes trial-by-trial estimates of the location of the helicopter in the form of a delta-rule (Nassar et al., 2010, 2012, 2016):

$$B_{t+1} = B_t + \alpha_t \times \delta_t \tag{1}$$

$$\delta_t = X_t - B_t \tag{2}$$

where B_t and B_{t+1} are beliefs about the current and next locations of the helicopter, respectively, X_t is the current observed bag drop, δ_t is the prediction error for the current trial, and a_t is the learning rate.

The ideal observer infers the position of the helicopter, *i.e.*, the mean of the distribution generating bag drops, from the positions of previously observed bag drops according to:

$$p(\mu_t | X_{1:t}) = \frac{p(X_{1:t} | \mu_t) p(\mu_t)}{p(X_{1:t})} \quad (3)$$

where μ_t is the position of the helicopter on trial t , and $X_{1:t}$ denotes the locations of bag drops observed on trials 1 through t .

The model posits that trial-by-trial learning rates are driven by two factors computed on every trial from the sequence of experienced bag drop locations (Figure 1B; Nassar et al., 2012, 2016). The first factor is *change point probability* (CPP, denoted as Ω), which measures the probability of the helicopter having switched locations since the last trial, given the observed prediction error. The second factor is *relative uncertainty* (RU, denoted as τ), which estimates the uncertainty in the exact location of the helicopter (often caused by a recent change point) as a function of the total uncertainty about the next bag drop.

Trial-by-trial learning rate is computed according to:

$$\alpha_t = \Omega_t + \tau_t \times (1 - \Omega_t) \quad (4)$$

Learning accelerates and favors recent outcomes when the estimated likelihood of a change point is high or when the uncertainty in the precise location of the helicopter is high. The effect of RU in determining the learning rate is highest when CPP is zero and, conversely, the effect of CPP is highest when uncertainty about the mean of the generative distribution is low.

In a recent instantiation of this normative model (Nassar et al., 2016), these two factors are computed recursively by first estimating RU as follows:

$$\tau_{t+1} = \frac{\Omega_t \sigma_N^2 + (1 - \Omega_t) \tau_t \sigma_N^2 + \Omega_t (1 - \Omega_t) (\delta_t (1 - \tau_t))^2}{\Omega_t \sigma_N^2 + (1 - \Omega_t) \tau_t \sigma_N^2 + \Omega_t (1 - \Omega_t) (\delta_t (1 - \tau_t))^2 + \sigma_N^2} \quad (5)$$

where Ω_t , τ_t , and δ_t are, respectively, the change point probability, relative uncertainty and prediction error on the previous trial; σ_N is the standard deviation of the bag drop distribution (previously defined as “noise”). The numerator of Equation 5 denotes the variance of the predictive distribution over possible helicopter locations, and includes terms for the variability in the bag drop distribution (σ_N) weighted under conditions of a recent change point or a stable environment, respectively, as well as an adjustment term for the variance arising from the difference in means of these distributions (McGuire et al., 2014). The denominator in Equation 5 includes all the terms in the numerator, plus a term to account for the remaining total uncertainty, namely uncertainty arising strictly from variability in the distribution of bag drops (σ_N). Total uncertainty in bag locations in this environment is attributable to both uncertainty about the center of the helicopter

(numerator) and uncertainty about the standard deviation of the bag drop distribution around the helicopter center. Accordingly, RU is calculated as the proportion of total uncertainty about the next bag location that is due to having an imprecise estimate of the location of the helicopter (McGuire et al., 2014; Nassar et al., 2016).

To calculate CPP, the model considers whether the generative mean has been resampled. The overall probability of resampling is called the Hazard rate, denoted by H . CPP on every trial is then computed from H , relative uncertainty (τ_t), SD of the bag drop distribution (σ_N) and the current prediction error (δ_t) as follows (Nassar et al., 2012, 2016):

$$\Omega_t = \frac{U(X_t|0, 300)H}{U(X_t|0, 300)H + N(\delta_t|0, \sigma_N^2/1 - \tau_t)(1 - H)} \quad (6)$$

where $U(X_t|0, 300)$ represents the probability of the current observation being randomly drawn from the possible helicopter locations and $N(\delta_t|0, \sigma_N^2/1 - \tau_t)$ is the likelihood of observing the current prediction error under the current predictive distribution. In principle, CPP is higher as the base-rate of change points increases; however, H in our experiment was held fixed throughout. Thus, CPP only depended on how unlikely an observed prediction error was under the current beliefs.

We obtained trial-by-trial estimates of CPP and RU by fitting the model to the sequences of experienced prediction errors using the true values of the hazard rate ($H = 0.125$) and noise in the generative distribution ($SD = 10$ or 25).

Behavioral Analysis. We used linear regression to test the effect of normative model factors and the non-normative reward factor on participants' trial-by-trial bucket updates, an analysis strategy successfully used in recent work with this task (McGuire et al., 2014; Nassar et al., 2016). Specifically, our analysis of task behavior was closely modeled after that used by Nassar and colleagues (2016). In other work with this task (Nassar et al., 2010, 2012, 2016), the normative model was alternatively, or additionally, fit to subjects' bucket predictions to infer latent variables, such as the baseline rate of change points. In contrast, here we simulated normative influences on behavior (trial-by-trial CPP and RU; by fitting the normative model to the sequence of observed bag drops), and examined how those relate to actual participant behavior.

We fit the regression models explaining behavior separately for each participant, combining data from all 4 experimental blocks. Each model included the following independent variables calculated for each trial: prediction error and the interaction between prediction error and CPP, RU, an indicator variable for noise (the *SD* of the bag distribution), reward value (1 for rewarded trial and 0 for non-rewarded trial), and the quadratic weighted distance from screen center (to adjust for bias towards the default bucket position in the center of the screen). All predictors were mean-centered. As discussed in Nassar et al. (2016), residuals from this regression scale with prediction errors, such that they are larger with larger absolute errors. To account for this violation of regression assumptions, we used weighted linear regression, as described in Nassar et al. (2016). In an initial regression, we estimated the variance of residuals in sliding windows of prediction error. The variance estimates were used to weight errors; in

addition, a ridge penalty was applied to shrink estimated coefficients. We used a sliding window method, where data were binned based on relative prediction error (absolute prediction error divided by the *SD* of the bag drop distribution), which indexed surprise. We applied the binning to account for a non-monotonic relationship between updating and surprise. Each bin contained 10% of the total data and bins were incremented by one percentile of the data. For further details on analysis strategy, refer to Nassar et al. (2016).

The regression coefficients for each independent variable in the above-described model were taken as the effects of each factor on the participant's updates. Group effects across participants for each factor were tested against zero using Wilcoxon signed-rank tests. We then tested the CPP, RU and reward effects for association with personality and cognitive variables of interest. In our and previous studies (McGuire et al., 2014), the CPP and RU effects were strongly positively correlated; we therefore averaged them into a "Normative Learning Factor" (NLF) to use in a first step analysis and minimize the number of tests conducted. We used separate regression models with NLF as the dependent variable and either the cognitive or personality factors as independent variables to test the separate associations with personality and cognition. We used F-tests to formally compare a full model of personality and cognitive predictors to a nested model with just the cognitive variable, IQ, thus testing if personality predicts normative learning above and beyond cognition.

Although the CPP and RU effects are strongly conceptually and quantitatively related, prior work has also shown some distinctions between the two (McGuire et al., 2014; Nassar et al., 2016). In cases where we found significant associations with

normative learning, we further unpacked these effects by looking at bivariate correlations between the CPP and RU effects and the significant individual difference variables. We used Spearman's rho for all reported correlations, to account for slight non-normality, especially in the cognitive and personality variables.

Experiment 1: Results

In this experiment, we investigated whether cognitive or personality factors explain individual differences in adaptive learning in young participants.

Model-free analyses showed that participants' learning was sensitive to normative model factors. Specifically, we examined if trial-by-trial learning was responsive to: (1) different magnitudes of spatial prediction error, and (2) sudden changes in the environment.

Both of these factors have previously been shown to influence belief updating in a changing environment (Nassar et al., 2010, 2012). Consistent with previous findings, participants' learning rates increased monotonically as the prediction error increased, in both low- and high-noise blocks (Figure 2A). Learning rates were also higher in the low-noise blocks than in the high-noise blocks for an equivalent PE magnitude (Figure 2A; median p across prediction error bins $< .0001$). Combined, these results show that participants are sensitive to CPP, as CPP increases with prediction error and is greater for low noise than high noise for a fixed PE (Nassar et al., 2012). Also consistent with previous findings, participants used the highest learning rates on trials right after a change point, and learning rates decreased steeply thereafter, stabilizing around the third trial

post change point (Figure 2B). This result shows that participants are also sensitive to RU, as RU peaks after a change point and declines thereafter (Nassar et al., 2012).

We used model-based analyses to further quantify how behavior depended on both normative and other incidental factors (Nassar et al., 2016). Trial-by-trial bucket update was the dependent variable, and the regression model included trial-wise PE, and the interaction between PE and CPP, RU, and reward as independent variables. The regression coefficients are a measure of the influence of that factor on each participant's learning behavior, controlling for the influence of the other factors in the regression model. Consistent with previous reports (McGuire et al., 2014; Nassar et al., 2016), participants, as a group, made larger updates based on both CPP ($Mdn = 0.37$, $Q1 = 0.22$, $Q3 = 0.50$) and RU ($Mdn = 0.19$, $Q1 = 0.07$, $Q3 = 0.32$), with both group effects being significantly different from zero ($Z = -6.02$ and $Z = -5.03$, respectively, $ps < .0001$; Figure 2C). Consistent with their normative roles, both CPP- and RU-based updating were positively correlated with total number of coins earned ($\rho = .70$, $p < .0001$ and $\rho = .44$, $p = .002$, respectively).

Participants' learning was not fully accounted for by the normative model, however. The normative model posits that learning is fully driven by CPP and RU (Nassar et al., 2012; 2016). In contrast, participants reliably updated based on the observed spatial prediction error alone, with higher PE engendering larger bucket updates ($Mdn = 0.61$, $Q1 = 0.50$, $Q3 = 0.71$, $Z = -6.09$, $p < .0001$). The regression coefficient on PE alone captures a tendency to rely on a fixed learning rate instead of learning adaptively according to the changing environment. Indeed, the tendency to update based

on PE alone is strongly negatively correlated with the tendency to update based on CPP or RU ($\rho = -.91$ and $\rho = -.58$, respectively, $ps < .0001$).

Participants' belief updating was also influenced by reward: participants updated more after trials with rewarded versus unrewarded coins ($Mdn = 0.06$, $Q1 = 0.01$, $Q3 = 0.10$, $Z = -5.00$, $p < .0001$). Reward-based updating is *non-normative* in our task because the value of coins in each bag does not carry information about the distribution of bag locations. Accordingly, the reward effect was negatively correlated with number of coins earned ($\rho = -.29$, $p = .046$), with those who were more influenced by bags with rewarded coins performing worse on the task.

As is clear in Figure 2C, however, there were individual differences in the degree of normative updating. These individual differences were associated with cognitive but not personality measures. We used regression models to test the association between cognitive and personality factors and normative learning. The two normative behavioral effects, CPP and RU, were strongly positively correlated across participants ($\rho = .58$, $p < .0001$; McGuire et al., 2014). Thus, to reduce the number of tests conducted, we used the average of the two effects (the “Normative Learning Factor”, NLF) as a dependent variable in the regression. A model with just IQ significantly predicted NLF (Table 1, $F(1,47) = 6.82$, $p = .012$), while a model only using personality measures did not significantly predict NLF ($F(7,41) = 0.78$, $p = .61$). A formal comparison of the full model (including cognitive—IQ—plus all personality measures) to a nested model with just IQ indicated that the nested model fits our data better than the full model ($F(7,40) = 0.63$, $p = .73$). Higher IQ was associated with more adaptive learning (higher NLF, Table 1, $t(47) = 2.61$, $p = .012$) and higher IQ was directly associated with better performance

(*i.e.*, more total coins caught) in the game ($\rho = .50, p = .0003$). To further unpack the effects of IQ, we examined the correlation between IQ and the normative model-based factors separately (CPP and RU). IQ was moderately positively correlated with both CPP- and RU- based updating ($\rho = .38, p = .006$ and $\rho = .32, p = .027$, respectively; Figures 2D and 2E), with higher IQ participants being more sensitive to these model-based factors.

However, IQ did not predict reward-based updating ($F(1,47) = 0.49, p = .49$; Figure 2F). Reward-based updating was also not significantly associated with personality measures (including “Big Five” dimensions of personality, sensation seeking, and trait anxiety; $F(7,41) = 0.69, p = .68$).

Experiment 1: Discussion

Experiment 1 demonstrated that learning behavior in a dynamic environment is influenced by normative factors from an approximate Bayesian model, and individual differences in normative learning correlate with cognitive abilities but not personality traits. Specifically, trial-by-trial surprise (CPP) and belief uncertainty (RU) increase learning from the most recent outcome (as shown previously in McGuire et al., 2014), and higher IQ participants show more sensitivity to these normative factors. Participants’ learning is also improperly influenced by incidental reward (McGuire et al., 2014); however, this tendency shows no association with cognitive abilities or personality.

Cognitive abilities are an important source of individual differences and have been associated with decision-making across a variety of tasks. Cognitive abilities are

robustly correlated with temporal discounting—the tendency to devalue future rewards—with higher IQ individuals exhibiting less discounting of delayed rewards (Shamosh & Gray, 2008). IQ is also associated with consistency in risky decisions, and in general with higher willingness to take calculated risks, perhaps due to an increased capacity to evaluate different options (Burks et al., 2009). Successful performance in our task requires inferring the underlying statistical structure of the environment from noisy evidence. Accurate inference could rely on cognitive abilities in a variety of different ways: it may reflect the ability to integrate abstract information from different sources (wind and helicopter movement; Gottfredson, 1997), or to actively maintain a representation of the environment, focusing attention on relevant information while ignoring the interference of noise (Kane & Engle, 2002). That is, though we observed a relationship between adaptive learning and IQ, measures of IQ do not isolate a single, specific cognitive process. Therefore, the relationship we observed might be traceable to the effects of specific cognitive processes, such as memory, which are correlated with IQ (Ackerman, Beier, & Boyle, 2005). We therefore examined associations with measures of more specific cognitive processes, in addition to the more general measure of IQ, in the following two experiments.

We investigated several dimensions of personality (“Big Five” dimensions of personality, sensation seeking, and trait anxiety), and we did not find any association between personality traits and adaptive learning. We had hypothesized that traits such as neuroticism or anxiety might increase participants’ perceived rate of change points, and thus lead to overall increased learning rates (Nassar et al., 2012). High trait anxiety has been associated with reduced adjustment of learning rates between stable and volatile

periods in an aversive learning task (Browning et al., 2015). A notable difference between their study and ours—that might explain why we do not see an association with anxiety—is that they used anxiety-provoking electrical shocks as stimuli.

In the next experiment, we set out to investigate if individual differences in adaptive learning are attributable to similar cognitive (versus personality) factors in a sample of elderly adults. This would replicate the findings of Experiment 1 and demonstrate their generalizability across the lifespan. Healthy aging itself might also influence performance in our task. Advanced age is associated with a decline in the ability to learn changing reward contingencies (Eppinger, Hämmerer, & Li, 2011), and a recent study with a different version of our predictive inference task showed that older participants exhibited a decreased sensitivity to uncertainty compared to young adults (Nassar et al., 2016).

Experiment 2: Method

In Experiment 2, we extended our investigation of cognitive and personality factors that influence adaptive learning in a sample of cognitively healthy older adults. Below (after Experiment 3), we will also use these data to consider healthy aging as a source of individual differences in adaptive learning by comparing performance between older and younger adults.

Participants. We recruited 41 cognitively healthy older participants (age 60–84, $M = 71.2$, $SD = 6.3$; 32% males). Exclusion criteria were identical to those used in Experiment 1; inclusion criteria were further restricted to older adults who had been

screened as cognitively normal in the past 2 years. To this end, participants were recruited from a pool of cognitively normal controls followed as part of a larger longitudinal study at the University of Pennsylvania's Alzheimer's Disease Core Center (ADCC). All controls undergo longitudinal medical, neurologic and psychiatric assessments, as well as a standard battery of psychometric measures, including those described by the National Alzheimer's Coordinating Center's (NACC) Uniform Data Set (UDS; Morris et al., 2006). Additional psychometric measures included the Mini Mental State Examination (Folstein, Folstein, & McHugh, 1975), Boston Naming Test (Williams, Mack, & Henderson, 1989), Trail Making Test (Crowe, 1998), Wechsler Digit Span Test (Wechsler, 1997), and the Consortium to Establish a Registry for Alzheimer's Disease (CERAD) Word List Memory Test (Morris et al., 1989). Designation of cognitively normal is determined by a consensus group of neurologists, geriatricians, neuropsychologists, and psychiatrists at the ADCC. Data from three participants were excluded because the participants were unable ($N = 2$) or refused ($N = 1$) to finish the task, so the effective sample size for analysis was 38 participants.

Procedure. The experimenter, order of administration, and training session were the same as in Experiment 1, but there were several slight differences in procedure. Starting with the sixth participant, to reduce task fatigue, we cut the number of trials in each experimental block from 150 to 100 (a 33% reduction). We collected the same personality measures as in Experiment 1, except for the State-Trait Anxiety Inventory (STAI) trait subscale (Spielberger et al., 1983). We offered participants the option to complete the questionnaires on paper in order to minimize computer use; five participants chose this option.

Although we did not collect any additional data from older adults at the time of testing, we were able to acquire measures of cognitive performance collected as part of their participation in the ADCC. Psychometric testing was performed within a median of 119 days (4 months) of testing in our experiment. Here, we specifically focus on the following measures: (a) total score from the Mini Mental State Examination, a questionnaire of cognitive impairment including questions about orientation, concentration, memory, and language; (b) total score from the Boston Naming Test, a test of identification of 30 objects presented as line drawings; (c) two timing measures from the Trail Making Test, a test of visual attention, processing speed, and mental flexibility. Participants are asked to sequentially connect 25 dots. In Part A, all dots contain numbers from 1 to 25, whereas Part B consists of alternating numbers and letters (1, A, 2, B, etc.). The timing measures for each part indicate time to successful completion of the task; (d) scores from the Wechsler Digit Span Test, a measure of short term memory. Participants are verbally presented with lists of digits and asked to repeat the digits in the presented order (forward test) or in reverse order (backward test). The number of digits increases by one until a participant fails two trials of the same length, such that scores indicate highest successfully repeated span lengths in each direction. (e) two memory measures derived from the Word List Memory task, a working memory and learning task. Participants are presented with a list of 10 high-frequency words which are read to them at a constant rate of 1 word every 2 seconds. The word list is presented 3 consecutive times, in randomized order. After every presentation, participants are asked to recall the words and responses are recorded. The first measure of memory performance is the total number of words immediately recalled across all three presentations (maximum is 30). The second measure

of memory performance is the total number of words recalled out of 10 after a delay. For one participant, Trail Making B score was not available at the time of collecting the other psychometric measures; we instead used a score obtained a year earlier. We note that reported results did not change when using a Trail Making B score imputed from other available cognitive variables.

Experiment 2: Results and Discussion

Model-free analyses in Experiment 2 revealed similar influences on learning rates as in Experiment 1 and previous reports (Nassar et al., 2010, 2012). Participants used higher learning rates with higher spatial prediction error and lower noise (Figure 3A; median p across prediction error bins = .0082), suggesting that learning was sensitive to the normative CPP factor (Nassar et al., 2012). Average learning rates peaked right after a change point and stabilized within three trials (Figure 3B), consistent with the normative influence of RU (Nassar et al., 2012).

Model-based analyses further corroborated that older participants' learning was sensitive to normative factors (Figure 3C; McGuire et al., 2014). Older participants, as a group, made larger updates based on the two normative factors, CPP ($Mdn = 0.33$, $Q1 = 0.22$, $Q3 = 0.47$, $Z = -5.36$, $p < .0001$) and RU ($Mdn = 0.18$, $Q1 = -0.02$, $Q3 = 0.31$, $Z = -2.99$, $p = .003$). As with the younger participants in Experiment 1, the behavior of older participants in Experiment 2 also systematically departed from the normative model. Older participants' updates showed a residual effect of observed spatial PE ($Mdn = 0.65$, $Q1 = 0.47$, $Q3 = 0.71$, $Z = -5.37$, $p < .0001$), suggesting that participants had some

tendency towards fixed learning rates rather than adapting learning rates based solely on CPP and RU. Reward also influenced belief updating: older participants updated more after trials with rewarded versus unrewarded coins ($Mdn = 0.02$, $Q1 = -0.01$, $Q3 = 0.04$, $Z = -2.53$, $p = .011$).

Similar to our results in young adults, individual differences in adaptive learning in older adults were more strongly associated with cognitive than personality measures. We again examined the separate effects of cognitive (IQ and cognitive performance) and personality variables (“Big Five” personality and sensation seeking) on the Normative Learning Factor (average of CPP and RU). A model with cognitive measures alone significantly predicted NLF (Table 2, $F(9,28) = 3.17$, $p = .009$), while a model with only personality measures did not significantly predict NLF ($F(6,31) = 0.47$, $p = .82$). The nested model with just cognitive variables was preferred over the full model with both cognitive and personality variables ($F(6,22) = 1.25$, $p = .32$). In this regression model, unlike in the sample of young adults in Experiment 1, IQ was not a significant predictor of adaptive learning (Table 2, $t(28) = -0.767$, $p = .45$). However, several other cognitive variables were significantly associated with adaptive learning (Table 2). Namely, two measures of memory performance— Digit Span forward score ($t(28) = 2.54$, $p = .017$) and Total Recall from the Word List Memory task ($t(28) = 2.05$, $p = .049$) — as well as Trail Making B score ($t(28) = 2.05$, $p = .0498$).

Unpacking these cognitive effects on adaptive learning, we found that Total Recall was positively correlated with both CPP- and RU-based updating ($\rho = .42$, $p = .009$; Figure 3D and $\rho = .45$, $p = .004$; Figure 3E, respectively), while the other two variables— Digit Span forward score ($\rho = .25$, $p = .14$ and $\rho = .28$, $p = .09$,

respectively) and Trail Making B score ($\rho = .0004$, $p = .99$ and $\rho = .17$, $p = .30$, respectively)— did not reach significance when examined separately for CPP- and RU-based updating and independently of other cognitive variables. Participants with better memory (Total Recall) made updates that were more sensitive to the model-derived normative factors that should govern learning. We note, however, that memory was not directly correlated with better performance (*i.e.*, total number of coins caught) in the game ($\rho = .13$, $p = .45$). In contrast to normative learning, cognitive abilities (including memory performance) in older adults were not associated with non-normative reward-based updating ($F(9,28) = 1.54$, $p = .18$; Figure 3F). As in Experiment 1, none of the personality measures were significantly associated with any of the factors influencing belief updating ($F(6,31) = 0.58$, $p = 0.75$).

Similar to what we observed in younger adults in Experiment 1, in Experiment 2, older adults' learning in a dynamic environment was influenced by both normative factors, CPP and RU (Nassar et al., 2012), as well as a non-normative factor, outcome reward (McGuire et al., 2014). Also as in Experiment 1, the effect of normative factors on adaptive learning in older adults was associated with differences in cognitive abilities and not with differences in personality. Combined, these two experiments suggest that cognitive factors have a stronger influence than personality on individual differences in adaptive learning.

However, the specific cognitive factors that were associated with normative learning differed across the two experiments. While in younger adults in Experiment 1 adaptive learning was linked to IQ (the only cognitive measure collected), in older adults in Experiment 2, individual differences in adaptive learning were correlated with memory

but not with IQ. One possible explanation is that the different findings in the two experiments are due to differences in IQ between the two samples. Even though median IQ in the young sample was quite high ($Mdn = 117$, 87th percentile of the IQ distribution), median IQ in the older sample was significantly higher ($Mdn = 124.5$, 95th percentile, $Z = 2.61$, $p = .009$). The extremely high level of cognitive ability in our sample of older adults is not too surprising, given that highly functioning older adults may be more likely to volunteer to participate in research, and that our recruitment specifically excluded individuals with cognitive impairments (see Method section), which might be common and go undiagnosed in a more representative sample of that age. It is possible that the association between IQ and adaptive learning might be obscured by restricted range of IQ in our sample of older adults.

Another possibility is that individual differences in adaptive learning might be driven by a specific aspect of cognitive function, which is generally correlated with the broader statistical concept of IQ (which would explain the association we observed between IQ and adaptive learning in young adults), but is dissociated from IQ in our sample of older adults. In fact, memory performance is typically positively correlated with IQ (Ackerman et al., 2005), but in our sample of older adults IQ was not significantly correlated with recall ($\rho = .11$, $p = .53$). In addition, memory ability is highly relevant to learning in this task. Remembering previous outcomes is essential to building and maintaining a mental representation of the task environment and statistics (e.g., the current helicopter location, one's confidence in that estimate, the noise in the bag drop distribution), which is necessary for adaptive modulation of learning rates according to normative factors. Updating with a fixed learning rate does not require such

memory demands. Therefore, it is highly plausible that memory capacity constrains one's ability to perform adaptive inference. It is also possible, however, that this is only true in older individuals, where memory abilities start to decline.

We cannot distinguish between these different possibilities, however, since we did not collect memory measures in the young adults in Experiment 1. Therefore, we conducted a third experiment, with young adults, in an attempt to replicate and disentangle the different cognitive influences on adaptive learning.

Experiment 3: Method

While our first two experiments both point to cognitive factors as explaining individual differences in adaptive learning, the specific associations we observed in the two samples were with different cognitive measures—IQ in younger adults and memory (but not IQ) in older adults. In Experiment 3, we aimed to distinguish potential explanations for this discrepancy by measuring both IQ and memory in the same sample of young adults who performed our adaptive learning task.

Participants. For Experiment 3 we recruited 40 participants from the University of Pennsylvania community (age 18–31, $M = 23.2$, $SD = 3.4$; 38% males). Inclusion and exclusion criteria were the same as in Experiment 1. Data from one participant were excluded from analysis because the participant fell asleep. The effective sample for analysis was therefore $N = 39$.

Procedure. Participants completed a memory task (described below), the similarities and matrix reasoning subtests of the WASI (Wechsler, 1999), and the

helicopter task, in that order. Experimenter MP obtained informed consent from the participants, and administered the IQ test and the helicopter task. Experimenter YX administered the memory task. We used the shorter version of the helicopter task (100 trials per experimental block), but the training was identical to that used in Experiments 1 and 2.

Memory Task. Participants performed a free recall task in a single session consisting of eleven word lists in total. The first list was for practice and the ten subsequent lists were scored. The procedures were modeled based on previous experimental studies of free recall (Polyn, Erlikhman, & Kahana, 2011; Zaromb, et al., 2006). Each list consisted of 16 words and each word was presented one at a time on a computer screen using E-Prime software (Version 2.0, Psychology Software Tools, Pittsburgh, PA). All text was presented in white with a black background. At the end of each list, participants were asked to name, in any order, as many words as they could recall from the just-presented list. The words and their order within lists were identical for all participants. All words were drawn randomly without replacement from the Toronto Noun Pool (retrieved from <http://memory.psych.upenn.edu/files/wordpools/nouns.txt>). The word pool consists of 480 words with estimates of word frequency and concreteness. We divided the pool into 16 groups (obtained by crossing quartiles of frequency and concreteness), and constructed each list by drawing a word from each of the 16 groups. At the beginning of each list, there was a 1500-ms delay before the first word was shown on the screen. Each word then appeared on screen for 3000-ms followed by a jittered interstimulus interval (uniformly drawn between 800 and 1200-ms). After the last item in each list, there was a

jittered delay period (uniform on 1200-1400-ms) before a 1000-ms tone sounded. The tone signaled to the participant the beginning of the recall period. From the tone, participants had 75 seconds to attempt to recall any words from the list that was just shown. Participants were signaled by a 2000-ms tone at the end of the 75-second period to stop recalling. The experimenter recorded participants' recalled words in order, on a pre-prepared answer sheet.

Statistical Analysis. Data were analyzed as in Experiments 1 and 2. In addition, we used the Aroian version of the Sobel's test (Aroian, 1947; Sobel, 1982) to evaluate the mediating effect of normative learning on the relationship between age group and task performance.

Experiment 3: Results and Discussion

In this third experiment, we administered the predictive-inference task to a second sample of younger adults, and collected measures of both IQ and free recall to investigate the differential effects of these two cognitive variables on adaptive learning.

Similar to Experiments 1 and 2, participants used higher learning rates with higher spatial prediction error and lower noise (Figure 4A; median p across prediction error bins $< .0001$) and higher learning rates right after a change point (Figure 4B). Also like the previous experiments, model-based analysis revealed a significant positive influence of both normative factors, CPP ($Mdn = 0.41$, $Q1 = 0.24$, $Q3 = 0.57$; $Z = -5.44$, $p < .0001$) and RU ($Mdn = 0.31$, $Q1 = 0.13$, $Q3 = 0.51$; $Z = -5.44$, $p < .0001$), and non-normative factors, PE ($Mdn = 0.56$, $Q1 = 0.42$, $Q3 = 0.71$; $Z = -5.36$, $p < .0001$) and

reward ($Mdn = 0.05$, $Q1 = 0.03$, $Q3 = 0.09$; $Z = -4.48$, $p < .0001$), on trial-by-trial updating (Figure 4C).

The results of Experiment 3 replicated the pattern observed in Experiment 2, in which memory (free recall), but not IQ, predicted normative-based belief updating. A model with IQ as the sole independent variable did not significantly predict NLF (Table 3, $F(1,37) = 0.13$, $p = .72$), while a model with only free recall did significantly predict NLF ($F(1,37) = 5.32$, $p = .027$). A formal comparison showed that adding IQ in addition to free recall did not improve prediction of NLF ($F(1,36) = 0.42$, $p = .52$). There was a small but insignificant direct association between free recall and helicopter task performance ($\rho = .19$, $p = .25$).

In looking at the effects of memory performance on CPP- and RU-based updating separately, we found that both were positively correlated with percent recall ($\rho = .41$, $p = .01$, Figure 4D; and $\rho = .27$, $p = .09$, Figure 4E, respectively). (Note that while we report two-tailed hypothesis tests to be consistent throughout the manuscript, the RU effect would reach significance in a one-tailed test, which would be merited given our stated attempt to replicate Experiment 2.) Finally, reward-based updating was not significantly associated with either IQ ($F(1,37) = 0.14$, $p = .71$) or percent recall ($F(1,37) = 2.32$, $p = .14$; Figure 4F).

In this second experiment with young participants, we replicated the memory effect seen in older participants— free recall performance predicted the degree to which participants update beliefs based on normative learning factors. Because we used previously collected data from the older sample in Experiment 2, the exact memory measures in the two experiments differed, though both are measures of immediate free

recall. When we z-scored the immediate memory measures within each sample and examined the association across the combined sample (total $N = 77$), memory performance was robustly associated with both CPP- and RU-based belief updating ($\rho = .42, p = .0001$ and $\rho = .35, p = .0019$, respectively).

However, we did not replicate the significant association between IQ and belief updating based on normative factors that we observed in the first sample of young adults in Experiment 1. If we combine data across the two young samples (total $N = 88$), the association between IQ and CPP-based updating remains statistically significant ($\rho = .22, p = 0.04$), while that between IQ and RU-based updating does not ($\rho = .17, p = .11$). Furthermore, if we combine across all three samples of (total $N = 126$), IQ is no longer significantly associated with either CPP- ($\rho = .15, p = .09$) or RU-based belief updating ($\rho = -.058, p = .52$). However, we hesitate to conclude that the correlation between IQ and adaptive learning (particularly CPP-based learning) is zero. To have 80% power to detect a true correlation of $\rho = .15$, we would need a sample size of 346, and to detect a true correlation of $\rho = .20$, we would need a sample size of 194, both of which are larger than the samples we collected in any individual experiment or our combined sample across all three experiments (calculated with “pwr” R package).

Group differences between young and elderly participants

We hypothesized that age would also be a source of individual differences in adaptive learning. To test this, we pooled data from the two samples of younger participants from Experiments 1 and 3, as their performance was comparable, and we had no *a priori*

reason to expect differences between the two samples. We then compared this combined young sample ($N = 88$) to the sample of older participants from Experiment 2 ($N = 38$, Figure 5). We looked at the two normative learning factors (CPP and RU effects) separately because prior research has shown differential effects of age on these two factors (Nassar et al., 2016). Overall, the younger group performed slightly better at the task, gathering more coins ($M = 29\%$ of total, $SD = 2.3$) than the older group ($M = 27.6\%$, $SD = 2.6$, $t(124) = 3.04$, $p = .003$). There was no significant difference between the groups in CPP-based updating ($t(124) = 1.07$, $p = .29$). However, younger participants were more influenced by the other model-based factor, RU ($M = 0.26$, $SD = .25$), than older participants ($M = 0.14$, $SD = .27$, $t(124) = 2.52$, $p = .013$). Younger participants were also more influenced by the non-normative reward factor ($M = 0.07$, $SD = .09$) than older participants ($M = 0.03$, $SD = .07$, $t(124) = 2.42$, $p = .017$).

Given that both task performance and normative learning differ between young and elderly participants, we asked whether differences in learning behavior account for differences in coins collected between age groups. Age group was a significant predictor of the RU effect ($b = -0.126$, $t(124) = -2.52$, $p = .013$), and the RU effect significantly predicted coins collected, while controlling for age group ($b = 0.032$, $t(123) = 4.05$, $p < .0001$). The effect of age group in predicting coins collected, while significant, decreased once the mediating variable (RU effect) was controlled for ($b = -0.010$, $t(123) = -2.25$, $p = .026$), compared to a model without the RU effect ($b = -0.014$, $t(124) = -3.04$, $p = .003$). The indirect (mediating) effect was significant (Sobel's $Z = -2.09$, $p = .036$), suggesting that RU-based updating partially mediates the relationship between age group and task

performance. However, reward-based updating, which also differed between age groups, was not a significant mediator of that relationship (Sobel's $Z = 1.43$, $p = .15$).

In this predictive-inference task, the updating behavior of older adults was less influenced by uncertainty. This replicates previous findings with a different version of the helicopter task (Nassar et al., 2016). Nassar and colleagues used computational modeling to show that older adults are specifically impaired in representing and using uncertainty to drive learning. As the task used by Nassar and colleagues (2016) included dramatic manipulations of uncertainty (the helicopter's location was occasionally directly revealed), our results show that this deficit in uncertainty-based learning extends to situations with more subtle fluctuations in uncertainty. Nassar and colleagues (2016) further showed that differences between age groups cannot be accounted for by differences in fluid intelligence or working memory. Our data are consistent with this interpretation, inasmuch as IQ scores were higher in the older adults. Furthermore, in a regression model predicting the RU effect ($F(3,122) = 3.76$, $p = .013$), age group is close to significance ($b = 1.08$, $t(122) = 1.95$, $p = .054$), after controlling for IQ ($b = 0.005$, $t(122) = 1.45$, $p = .15$) and its interaction with age group ($b = -0.01$, $t(122) = -2.18$, $p = .031$). Unfortunately, we cannot perform a similar test looking at the interaction between age and memory because we have different memory measures for the two age groups. Nassar et al. (2016) also reported that older adults are more sensitive to CPP than young adults, but we did not find such an effect.

Older participants were also less influenced than young participants by non-informative rewards. This finding is consistent with previous evidence that healthy aging is associated with decreased sensitivity to reward in a probabilistic learning task, which is

neurobiologically mediated by reduced white matter integrity in select prefrontal pathways (Samanez-Larkin, Levens, Perry, Dougherty, & Knutson, 2012). Interestingly, older adults' diminished sensitivity to reward is present both in contexts where reward is informative and advantageous (Samanez-Larkin et al., 2012) and where reward is non-informative and reward sensitivity might weaken performance.

General Discussion

Across three experiments, we investigated how personality, cognitive abilities and age affect individual differences in adaptive learning in volatile environments. We used a modified version of a predictive-inference task that allows tracking of trial-by-trial learning rates (McGuire et al., 2012; Nassar et al., 2010, 2012, 2016). Consistent with several previous reports using this task (McGuire et al., 2014; Nassar et al., 2010, 2012), we found that learning rates were influenced by two normative factors related to surprise (CPP) and uncertainty (RU), as well as by incidental rewards. As a unique contribution of this paper, we found that the degree of reliance on normative learning factors is positively associated with cognitive abilities — specifically memory abilities measured with free recall — but not with personality. In addition, we found that advanced age was associated with a reduced influence of the normative RU factor and a reduced influence of the non-normative reward factor. Thus, age and cognitive abilities had distinct patterns of overall influence on adaptive learning.

Our finding that better memory abilities are associated with the influence of normative learning factors is broadly consistent with a body of work showing that

reliance on “model-based” rather than “model-free” reinforcement learning is dependent on cognitive resources (Otto, Raio, Chiang, Phelps, & Daw, 2013, Otto et al., 2014). In this framework, model-based learning relies on building a sophisticated mental model of the environment and evaluating choices in the context of that model, in the same way that adaptive inference in the helicopter task requires building an internal model of the structure of the environment (incorporating components such as wind, clouds and a moving helicopter) and flexibly updating one’s learning rate depending on the circumstances. In contrast, model-free learning relies on the simpler approach of caching the value of different actions (Daw, Niv, & Dayan, 2005). Recent work using this reinforcement learning framework has shown that individual differences in cognitive abilities, specifically better working memory and cognitive control, predict increased model-based vs. model-free contributions to learning (Otto et al., 2013, 2014). In addition, high working memory capacity protects individuals against the deleterious effect of stress on model-based learning (Otto et al., 2013).

With our results, two independent studies have now shown that learning in older participants is less driven by uncertainty than in young participants. Our previous fMRI study has shown that activity in anterior prefrontal regions increases with RU, while activity in the medial temporal lobe decreases with RU (McGuire et al., 2014). Both of these brain regions show structural and functional impairment in healthy aging (Buckner, 2004; Fjell et al., 2014; Raz et al., 2005;), which might contribute to a diminished ability to compute or use an RU signal. Future work could investigate whether reduced uncertainty-driven learning in older adults is associated with changes in prefrontal or hippocampal activity, or both.

Prior research has hinted at a potential role for affect and arousal in adaptive learning. Pupil diameter, which is a measure of arousal, tracks normative factors that should influence learning rates, and the extent to which pupil diameter tracks these factors across participants predicts learning rates (Nassar et al., 2012). BOLD activity in dorsomedial frontal cortex and anterior insula, two regions linked with arousal (Seeley et al., 2007), is modulated by factors that influence learning rate, and the extent to which activity in these regions is modulated by normative factors predicts the degree of behavioral sensitivity to these same factors (McGuire et al., 2014). However, we did not find that adaptive learning was associated with any measures of personality, including some that should capture differences in affective arousal. This result differs from another recent report, which found that anxiety was associated with a reduced effect of volatility on learning rates (Browning et al., 2015). As discussed above, this study may have found an association as the stimuli being predicted were aversive (and potentially anxiety-inducing) electric shocks. This difference in findings suggests that the influence of personality on adaptive learning may be context-sensitive. More broadly, our results suggest that adaptive learning is a result of the interplay between both cognitive and affective factors, with cognitive abilities playing a crucial role in constructing mental models and expectations, the deviations from which drive arousal and other affective responses. Future work is needed to more completely characterize this complex interplay.

Overall, our findings further extend the aspects of learning and decision-making that are affected by age and cognitive abilities. In addition to known and specific benefits of better memory, our results highlight how this core psychological process enables people to behave more adaptively in a changing world.

Tables

Table 1. Summary of regression results for three models predicting NLF in Experiment 1. Model 1 includes just one cognitive measure, IQ. Model 2 includes only personality measures. Model 3 includes both cognitive and all personality measures of interest.

| Variable | Model 1 | | | Model 2 | | | Model 3 | | |
|----------|----------------|-------|--------|----------------|-------|--------|----------------|-------|--------------------|
| | B | SE B | t(47) | B | SE B | t(41) | B | SE B | t(40) |
| I | -0.675 | 0.347 | -1.804 | 0.515 | 0.520 | 0.991 | -0.526 | 0.684 | -0.768 |
| IQ | 0.008 | 0.003 | 2.611* | | | | 0.008 | 0.003 | 2.208 ^a |
| A | | | | 0.026 | 0.061 | 0.435 | 0.025 | 0.058 | 0.436 |
| E | | | | -0.070 | 0.043 | -1.633 | -0.066 | 0.041 | -1.609 |
| C | | | | -0.032 | 0.054 | -0.586 | -0.020 | 0.052 | -0.383 |
| N | | | | 0.060 | 0.053 | 1.130 | 0.055 | 0.051 | 1.089 |
| O | | | | -0.013 | 0.056 | -0.240 | -0.012 | 0.054 | -0.221 |
| STAI | | | | -0.007 | 0.005 | -1.526 | -0.005 | 0.005 | -1.032 |
| SS | | | | 0.065 | 0.059 | 1.097 | 0.066 | 0.057 | 1.171 |
| Model F | F(1,47)= 6.82* | | | F(7,41) = 0.78 | | | F(8,40) = 1.36 | | |

I = Intercept; A = Agreeableness; E = Extraversion; C = Conscientiousness; N = Neuroticism; O = Openness; SS = Sensation Seeking; B = unstandardized regression coefficient; SE B = standard error of B; t(df) = t-test statistic and degrees of freedom

* $p < 0.013$

^a $p = 0.033$

Table 2. Summary of regression results for three models predicting NLF in Experiment 2. Model 1 includes only cognitive measures. Model 2 includes only personality measures. Model 3 includes all cognitive and personality measures of interest.

| Variable | Model 1 | | | Model 2 | | | Model 3 | | |
|----------|----------------|-------|--------------------|----------------|-------|--------|------------------------------|-------|--------------------|
| | B | SE B | t(28) | B | SE B | t(31) | B | SE B | t(22) |
| I | -1.747 | 1.168 | -1.496 | 0.367 | 0.402 | 0.912 | -2.433 | 1.272 | -1.913 |
| IQ | -0.002 | 0.003 | -0.767 | | | | -0.000002 | 0.004 | -0.001 |
| MMSE | 0.008 | 0.032 | 0.256 | | | | -0.008 | 0.039 | -0.201 |
| DSpanF | 0.053 | 0.021 | 2.537* | | | | 0.082 | 0.027 | 3.000 ^b |
| DSpanB | -0.004 | 0.017 | -0.213 | | | | -0.024 | 0.022 | -1.084 |
| Trails A | -0.004 | 0.005 | -0.724 | | | | -0.003 | 0.006 | -0.469 |
| Trails B | 0.003 | 0.001 | 2.050 ^a | | | | 0.004 | 0.002 | 2.167 ^c |
| Boston N | 0.034 | 0.023 | 1.445 | | | | 0.045 | 0.026 | 1.704 |
| T Recall | 0.020 | 0.010 | 2.053 ^a | | | | 0.026 | 0.011 | 2.415 ^d |
| D Recall | 0.005 | 0.027 | 0.191 | | | | -0.014 | 0.034 | -0.402 |
| A | | | | 0.014 | 0.079 | 0.180 | 0.006 | 0.076 | 0.080 |
| E | | | | -0.067 | 0.053 | -1.254 | 0.005 | 0.048 | 0.110 |
| C | | | | -0.030 | 0.067 | -0.455 | -0.068 | 0.062 | -1.106 |
| N | | | | -0.045 | 0.051 | -0.935 | -0.013 | 0.044 | -0.298 |
| O | | | | 0.052 | 0.062 | 0.829 | 0.120 | 0.054 | 2.214 ^c |
| SS | | | | 0.035 | 0.061 | 0.582 | 0.075 | 0.058 | 1.306 |
| Model F | F(9,28)= 3.17* | | | F(6,31) = 0.47 | | | F(15,22) = 2.50 ^d | | |

I = Intercept; MMSE = Mini Mental State Examination score; DSpanF = Digit Span Forward score; DSpanB = Digit Span Backward score; Boston N = Boston Naming score; T Recall = Word List Memory Total Recall score; D Recall = Word List Memory Delayed Recall score; A = Agreeableness; E = Extraversion; C = Conscientiousness; N = Neuroticism; O = Openness; SS = Sensation Seeking; B = unstandardized regression coefficient; SE B = standard error of B; t(df) = t-test statistic and degrees of freedom; * $p < 0.017$; ^a $p < 0.0498$; ^b $p < 0.007$; ^c $p < 0.042$; ^d $p < 0.025$

Table 3. Summary of regression results for models predicting NLF in Experiment 3. Model 1 only includes Percent Recall. Model 2 only includes IQ. Model 3 includes both Percent Recall and IQ.

| | Model 1 | | | Model 2 | | | Model 3 | | |
|-----------|-----------------|-------|--------|---------------|-------|--------|----------------|-------|--------|
| Variable | B | SE B | t(37) | B | SE B | t(37) | B | SE B | t(36) |
| Intercept | -0.017 | 0.173 | -0.096 | 0.600 | 0.626 | 0.957 | 0.354 | 0.601 | 0.589 |
| RP | 0.633 | 0.274 | 2.307* | | | | 0.653 | 0.278 | 2.346* |
| IQ | | | | -0.002 | 0.005 | -0.361 | -0.003 | 0.005 | -0.645 |
| Model F | F(1,37) = 5.32* | | | F(1,37)= 0.13 | | | F(2,36) = 2.83 | | |

PR = percent recall; B = unstandardized regression coefficient; SE B = standard error of B; t(dfs) = t-test statistic and degrees of freedom

* $p < 0.027$

Figures

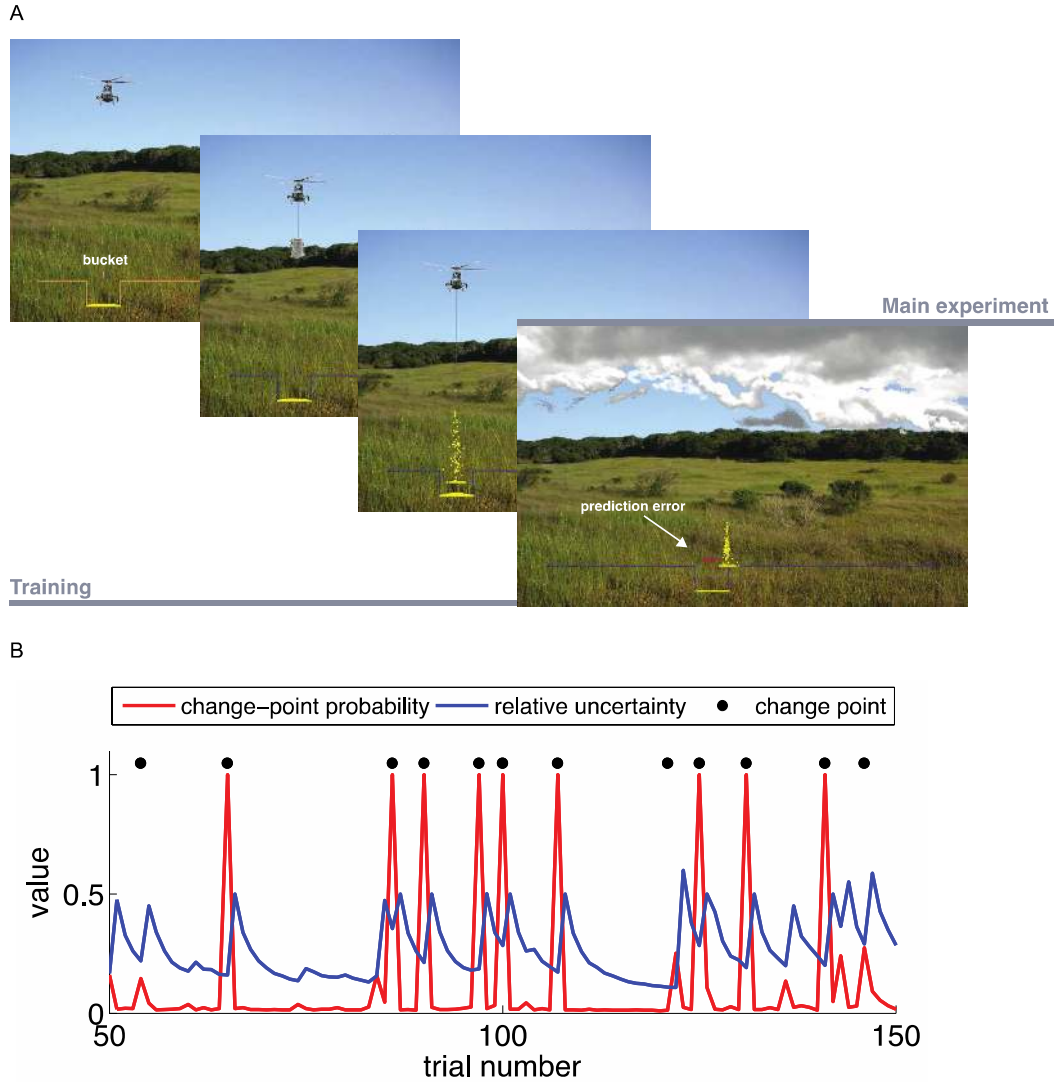


Figure 1. Overview of helicopter task and normative learning factors. **A.** Screenshots of helicopter task taken from training and experimental blocks. Participants were asked to position a bucket (orange rectangle) right under the helicopter to collect coins. On every trial, the helicopter drops bags which explode into coins upon reaching the ground. Once participants position the bucket and commit the selection, the bucket turns blue and can no longer be moved until the next trial. Bucket shows the gradual accumulation of coins over the course of a block. On every trial, participants receive visual feedback about the distance between the bag drop location and the last selected bucket position (prediction error, in red). All blocks in the main experiment included clouds, which obscured the helicopter from view. **B.** Fluctuation of normative learning factors in a sample experimental block. Change points (black dots) typically result in large prediction errors, and change-point probability (in red) is highest on such trials, while relative uncertainty (in blue) increases right after a change point and gradually decreases thereafter, as learner acquires additional evidence about the new helicopter position.

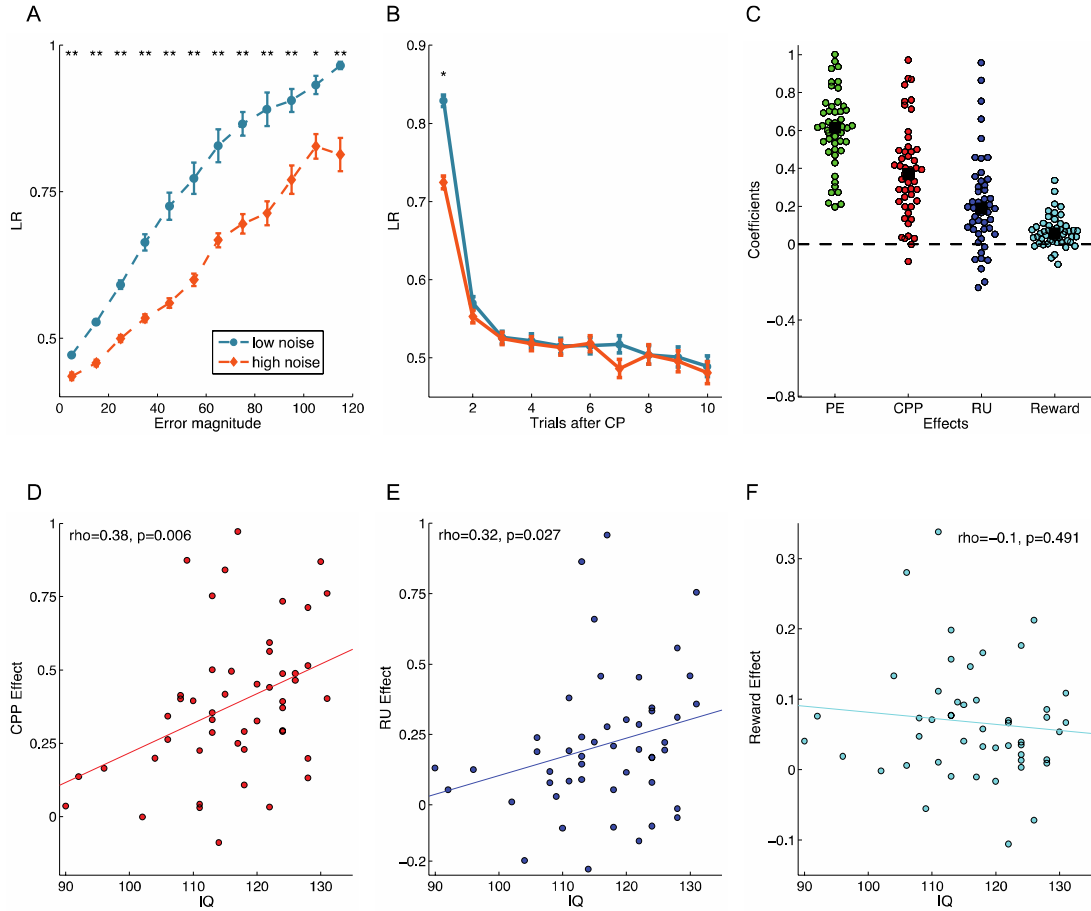


Figure 2. Experiment 1 results: normative learning behavior and its associations with IQ. **A.** Relationship between error magnitude and learning rates, shown separately for low (blue) versus high noise/wind (orange) blocks. Learning rates scaled with the magnitude of prediction errors, and were higher in the low wind blocks. Average learning rates (circles) were calculated in bins of error magnitude in intervals of 10 units centered around the midpoint of each bin, combining data from all participants. In both **A** and **B**: error bars indicate standard error of the mean; significant group differences between low and high noise blocks are indicated by * ($p < 0.05$) or ** ($p < 0.0001$). **B.** Relationship between number of trials after a change point and learning rates, shown separately by type of block. Learning rates were highest on trials right after a change point and decreased steeply thereafter. Each circle represents average learning rate grouped by number of trials after a change point, combining data from all participants. **C.** Individual differences in effects of normative and non-normative factors on updating behavior. Each effect is calculated as regression coefficients for respective factor in predicting bucket updates. Each circle represents the respective coefficient for one participant; black square represents median coefficient for each factor. Circles are jittered for better visibility. Reward effect is the effect of updating based on rewarded versus non-rewarded coins. **D.** Scatterplot of relationship between IQ and CPP effect (as represented in **C**), including best-fit line. Higher IQ was associated with higher CPP-based updating. **E.** Scatterplot of relationship between IQ and RU effect (as represented in **C**), including best-fit line. Higher IQ was associated with higher RU-based updating. **F.** Scatterplot of relationship between IQ and Reward effect (as represented in **C**), including best-fit line. All correlations are Spearman's rho.

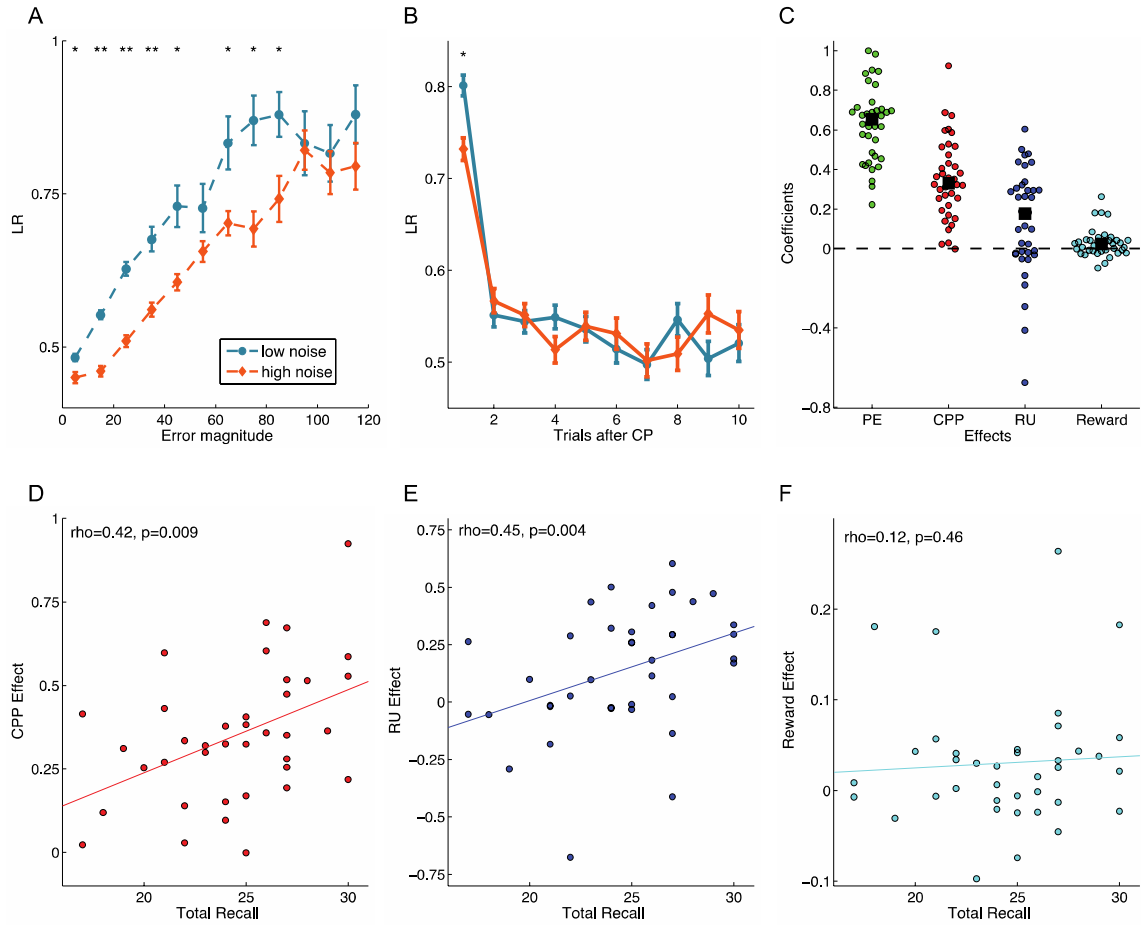


Figure 3. Experiment 2 results: normative learning behavior and its associations with memory. **A.** Relationship between error magnitude and learning rates, shown separately for low versus high noise/wind blocks. Learning rates scaled with the magnitude of prediction errors, and were higher in the low wind blocks. Average learning rates (circles) were calculated in bins of error magnitude in intervals of 10 units centered around the midpoint of each bin, combining data from all participants. In both **A** and **B**: error bars indicate standard error of the mean; significant group differences between low and high noise blocks are indicated by * ($p < 0.05$) or ** ($p < 0.0001$). **B.** Relationship between number of trials after a change point and learning rates, shown separately by type of block. Learning rates were highest on trials right after a change point and decreased steeply thereafter. Each circle represents average learning rate grouped by number of trials after a change point, combining data from all participants. **C.** Individual differences in effects of normative and non-normative factors on updating behavior. Each effect is calculated as regression coefficients for respective factor in predicting bucket updates. Each circle represents the respective coefficient for one participant; black square represents median coefficient for each factor. Circles are jittered for better visibility. Reward effect is the effect of updating based on rewarded versus non-rewarded coins. **D, E,** and **F.** Scatterplots of relationship between memory and CPP-, RU-, and Reward effects, respectively, including best-fit lines. All correlations are Spearman's rho.

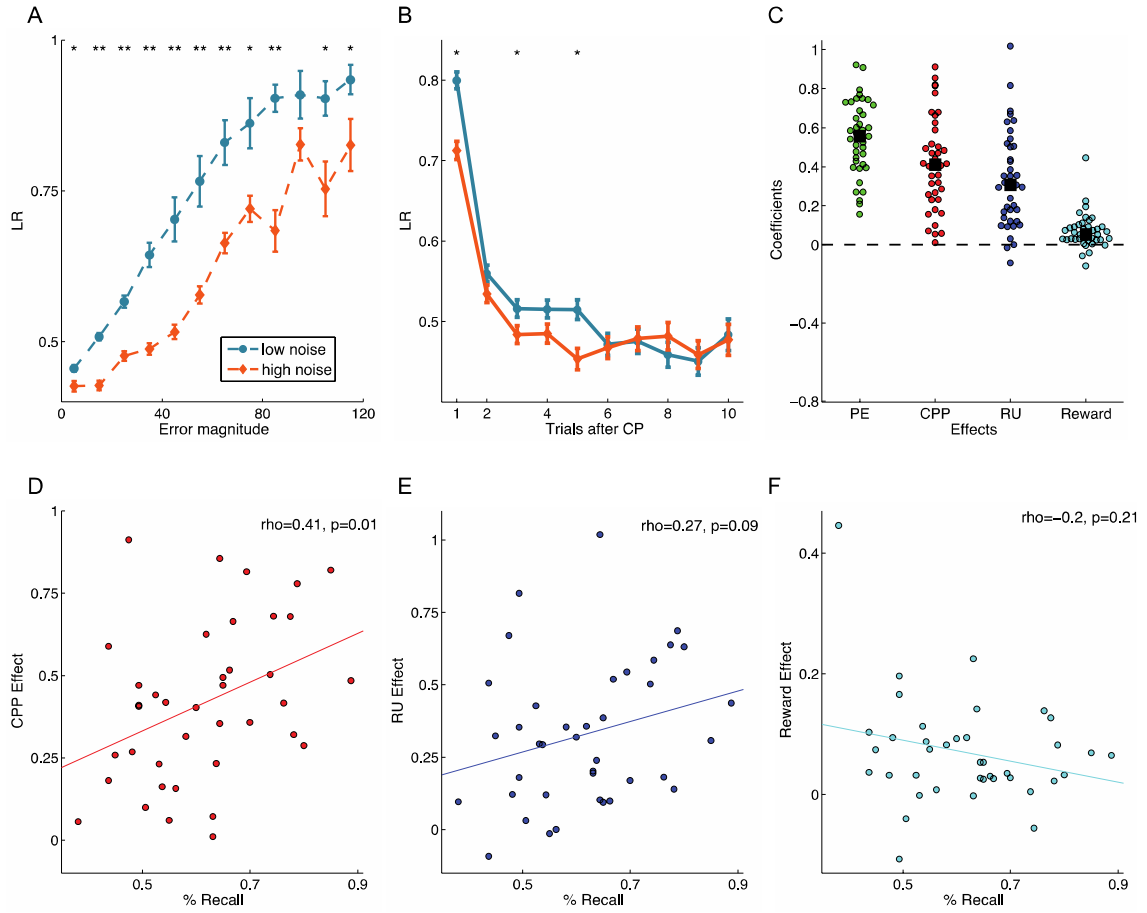


Figure 4. Experiment 3 results: normative learning behavior and its associations with memory. **A.** Relationship between error magnitude and learning rates, shown separately for low (blue) versus high noise/wind (orange) blocks. Learning rates scaled with the magnitude of prediction errors, and were higher in the low wind blocks. Average learning rates (circles) were calculated in bins of error magnitude in intervals of 10 units centered around the midpoint of each bin, combining data from all participants. In both **A** and **B**: error bars indicate standard error of the mean; significant group differences between low and high noise blocks are indicated by * ($p < 0.05$) or ** ($p < 0.0001$). **B.** Relationship between number of trials after a change point and learning rates, shown separately by type of block. Learning rates were highest on trials right after a change point and decreased steeply thereafter. Each circle represents average learning rate grouped by number of trials after a change point, combining data from all participants. **C.** Individual differences in effects of normative and non-normative factors on updating behavior. Each effect is calculated as regression coefficients for respective factor in predicting bucket updates. Each circle represents the respective coefficient for one participant; black square represents median coefficient for each factor. Circles are jittered for better visibility. Reward effect is the effect of updating based on rewarded versus non-rewarded coins. **D, E, and F.** Scatterplots of relationship between memory and CPP-, RU-, and Reward effects, respectively, including best-fit lines. All correlations are Spearman's rho.




Figure 5. Group differences between young and older participants in updating based on normative factors and the incidental reward factor. Older participants update less based on RU and coin value. Barplots show average effects for young (dark gray) and old (light gray) and error bars show standard error of the mean. Asterisk represents a significant group difference with $p < 0.02$.

CHAPTER 3 — Diminished cortical thickness is associated with impulsive choice in adolescence

Marieta Pehlivanova, Daniel H. Wolf, Aristeidis Sotiras, Antonia Kaczurkin, Tyler M. Moore, Rastko Ciric, Philip A. Cook, Angel Garcia de La Garza, Adon Rosen, Kosha Ruparel, Anup Sharma, Russell T. Shinohara, David R. Roalf, Ruben C. Gur, Christos Davatzikos, Raquel E. Gur, Joseph W. Kable, & Theodore D. Satterthwaite

Abstract

Adolescence is characterized by both maturation of brain structure and increased risk of negative outcomes from behaviors associated with impulsive decision-making, such as substance abuse and automobile accidents. One important index of impulsive choice is delay discounting (DD), which measures the tendency to prefer smaller rewards available soon to larger rewards delivered after a delay. However, it remains largely unknown how individual differences in structural brain development may be associated with impulsive choice during adolescence. Leveraging a unique large sample of 427 youths (208 males and 219 females) studied as part of the Philadelphia Neurodevelopmental Cohort, we examined associations between delay discounting and cortical thickness within structural covariance networks. These structural networks were derived using non-negative matrix factorization, an advanced multivariate analysis technique for dimensionality reduction, and analyzed using generalized additive models with penalized splines to capture both linear and nonlinear developmental effects. We found that impulsive choice was most strongly associated with diminished cortical thickness in structural brain networks that

encompassed the ventromedial prefrontal cortex, orbitofrontal cortex, temporal pole, and temporal-parietal junction. Furthermore, structural brain networks predicted DD above and beyond cognitive performance. Taken together, these results suggest that reduced cortical thickness in regions known to be involved in value-based decision-making is a marker of impulsive choice during the critical period of adolescence.

Significance

Risky behaviors during adolescence, such as initiation of drug use or reckless driving, are a major source of morbidity and mortality. In this study, we present evidence from a large sample of youth that diminished cortical thickness in specific structural brain networks is associated with impulsive choice. Notably, the strongest association between impulsive choice and brain structure was seen in regions implicated in value-based decision-making, namely the ventromedial prefrontal and orbitofrontal cortex. Moving forward, such neuroanatomical markers of impulsivity may aid in the development of personalized interventions targeted to reduce risk of negative outcomes during the critical period of adolescence.

Introduction

Adolescence is marked by an increased vulnerability to risky behaviors, such as tobacco, alcohol, and drug use, reckless driving, and unprotected sex, which can lead to increased morbidity and mortality (Eaton et al., 2011). During this vulnerable period, the

brain undergoes dramatic structural changes (Giedd et al., 1999; Sowell et al., 2004). Some evidence suggests that risk during adolescence is associated with differential maturation of brain regions related to reward processing (such as the orbitofrontal cortex and ventral striatum) and those necessary for cognitive control (such as the dorsolateral prefrontal cortex, dlPFC; Casey, Jones, & Hare, 2008; Van Leijenhorst et al., 2010). One of the most commonly used indices of impulsive choice is delay discounting (DD)— a behavioral measure of impulsivity where one chooses between a smaller reward delivered sooner, and a larger reward with a longer delay (Kable, 2013; Kirby & Maraković, 1995; Peters & Büchel, 2011). Delay discounting engages regions known to mature at different rates in adolescence, including dlPFC (Peters & Büchel, 2011), orbitofrontal cortex and ventral striatum (Bartra, McGuire, & Kable, 2013; Kable & Glimcher, 2007). Increased DD has been proposed as a framework for understanding substance abuse and other risky decisions as reflecting impulsive choices of immediate reward (Bickel et al., 2007). Indeed, studies of adolescents show that higher impulsivity, as indexed by higher discounting, is associated with increased smoking frequency (Reynolds, 2004), greater alcohol consumption (Field, Christiansen, Cole, & Goudie, 2007), and predicts longitudinal increase in both smoking (Audrain-McGovern et al., 2009) and alcohol use (Fennie et al., 2013).

At present, it remains relatively unknown how individual differences in structural brain development may relate to DD in adolescents. Neuroanatomical studies in adults are more numerous, but have yielded inconsistent results, perhaps due to small samples and focused region-of-interest analyses (for a review see Kable & Levy, 2015). For

example, it has been reported that greater DD (more impulsive choice) is associated with reduced gray matter volume in lateral prefrontal cortex (Bjork, Momenan, & Hommer, 2009), superior frontal gyrus (Schwartz et al., 2010), and putamen (Cho et al., 2013; Dombrovski et al., 2012). Furthermore, greater DD has been associated with larger volume of the ventral striatum and posterior cingulate cortex (PCC, Schwartz et al., 2010), medial prefrontal regions and anterior cingulate cortex (ACC, Cho et al., 2013), and prefrontal cortex (Wang et al., 2016). One study of cortical thickness (CT) in adults revealed an association between higher DD and decreased CT in both medial prefrontal cortex and the ACC (Bernhardt et al., 2014). To our knowledge, there have been no neuroanatomical studies in adolescents to specifically examine the relationship between DD and cortical thickness. Notably, findings from adults may not necessarily extend to adolescents, given the dynamic re-modeling of brain structure that occurs during this critical period (Sowell et al., 2004).

Accordingly, here we investigated how individual differences in DD may be associated with differences in brain structure during adolescence. To do this, we capitalized upon a large sample of 427 youths imaged as part of the Philadelphia Neurodevelopmental Cohort (Satterthwaite et al., 2014a; 2016). We delineated covariance networks of cortical thickness using a recently-developed application of non-negative matrix factorization for the multivariate analysis of high-dimensional neuroimaging data (Sotiras, Resnick, & Davatzikos, 2015). Subsequently, we evaluated the association between DD and CT in each network while specifically modeling both linear and nonlinear developmental effects using penalized splines. We hypothesized that

we would find associations between DD and CT in brain regions encoding reward value such as the ventromedial prefrontal cortex (vmPFC; Bartra et al., 2013; Kable & Glimcher, 2007), as well as regions subserving cognitive control (e.g., dlPFC). As described below, diminished CT in these as well as other networks was associated with impulsive choice, and predicted individual variation in DD above and beyond that explained by cognitive performance.

Methods

Participants and sample construction

Participants were a subsample of 1,601 youths recruited as part of the Philadelphia Neurodevelopmental Cohort (PNC) who underwent neurocognitive assessment (Gur et al., 2010; 2012) as well as multimodal neuroimaging (Satterthwaite et al., 2014a; 2016). A sub-sample of PNC participants ($n = 453$) completed the delay discounting (DD) task. Of those, $n = 2$ did not pass the quality control criteria for the task (described below). Additionally, $n = 24$ participants were excluded for the following reasons: health conditions that could impact brain structure ($n = 19$), scanning performed more than 12 months from DD testing ($n = 1$), or inadequate structural image quality ($n = 4$). The remaining $n = 427$ participants constituted our final sample for analysis (mean age at scanning: 17.0 ± 3.2 years, age range: 9.3–24.3 years; 48.7%, $n = 208$ males).

Delay discounting task

The DD task consisted of 34 self-paced questions where the participant chose between a smaller amount of money available immediately or a larger amount available after a delay. This task was modeled after the work of Senecal et al. (2012). The smaller, immediate rewards ranged between \$10 and \$34 and were always displayed at the top of the computer screen. The larger, delayed rewards were fixed at \$25, \$30, or \$35, with the delays ranging between 1 and 171 days. Larger, delayed rewards were always displayed on the bottom of the screen. All rewards were hypothetical but participants were instructed to make decisions as if the choices were real. Discount rates based on hypothetical choices have shown no systematic differences from discount rates based on real rewards, in the same subjects (Johnson & Bickel, 2002). The set of choices was identical in content and order for all participants. The DD task was administered as part of an hour-long web-based battery of neurocognitive tests (Computerized Neurocognitive Battery, described below), on a separate day from the imaging session. The mean interval between the DD task and imaging was 0.44 months with a SD of 1 month (range 0–8 months).

Discount rates from the delay discounting task were calculated assuming a hyperbolic discounting model of the form: $SV = A/(1+kD)$, where SV is the subjective value of the delayed reward, A is amount of the delayed reward, D is the delay in days, and k is the subject-specific discount rate (Mazur, 1987). We used the *fmincon* optimization algorithm in MATLAB (Mathworks, Natick, MA) to estimate the best-fitting k from each participant's choice data. A higher k value indicates steeper

discounting of delayed rewards and thus more impulsive choices. As the distribution of discount rates is highly right-skewed, we used log-transformed k ($\log k$) in all analyses.

We performed quality control to ensure that participants were not responding randomly, and verified that their responses were a function of task variables which should be relevant to the choice. Although a hyperbolic discounting model has been shown to fit discounting data better than an exponential model (Kirby & Maraković, 1995), quality control was performed independently of assumptions about the shape of the discount function. Specifically, each participant's responses were fit with a logistic regression model, with predictors including the immediate amount, delayed amount, delay, their respective squared terms, and two-way interaction terms. We assessed goodness of fit of this model using the coefficient of discrimination (Tjur, 2009), and discarded DD data from any participant who had a value of less than 0.20.

Neurocognitive battery

Cognition was assessed using the University of Pennsylvania Computerized Neurocognitive Battery (Penn CNB, Gur et al., 2010; 2012) during the same session that delay discounting was evaluated. Briefly, this hour-long battery consisted of 14 tests administered in a fixed order, evaluating various aspects of cognition, including executive control, episodic memory, complex reasoning, social cognition, and sensorimotor and motor speed. Except for the motor tests that only measure speed, each test provides measures of both accuracy and speed. Performance on the tests for each domain is summarized as cognitive factors obtained with exploratory factor analysis with an oblique rotation (Moore, Reise, Gur, Hakonarson, & Gur, 2015). Prior work has

demonstrated that accuracy on this battery can be parsimoniously summarized as either one overall cognitive performance factor or three domain-specific factors, including executive function and complex reasoning combined, social cognition, and episodic memory (Moore et al., 2015). Associations between DD and factor scores for each of these dimensions were analyzed, as described below.

Image acquisition and quality assurance

Image acquisition and processing are reported in detail elsewhere (Satterthwaite et al., 2014a; 2016). Briefly, all data were acquired on a single scanner (Siemens TIM Trio 3 Tesla, Erlangen, Germany; 32-channel head coil) using the same imaging sequences. Structural brain scanning was completed using a magnetization-prepared, rapid acquisition gradient-echo (MPRAGE) T1-weighted image with the following parameters: TR 1810 ms; TE 3.51 ms; FOV 180x240 mm; matrix 192x256; 160 slices; slice thickness/gap 1/0 mm; TI 1100 ms; flip angle 9 degrees; effective voxel resolution of 0.93 x 0.93 x 1.00 mm; total acquisition time 3:28 min. T1 image quality was independently assessed by three expert image analysts, who were trained to >85% concordance with faculty consensus rating on an independent dataset; images with substantial artifact were excluded from analysis.

Image processing and cortical thickness estimation

Structural image processing for estimating cortical thickness (CT) used tools included in Advanced Normalization Tools (ANTs, Tustison et al., 2014). In order to avoid registration bias and maximize sensitivity to detect regional effects that can be impacted by registration error, a custom adolescent template and tissue priors were

created. Structural images were then processed and registered to this template using the ANTs CT pipeline (Tustison et al., 2014). This procedure includes brain extraction, N4 bias field correction (Tustison et al., 2010), Atropos probabilistic tissue segmentation (Avants et al., 2011b), the top-performing SyN diffeomorphic registration method (Klein et al., 2010; Avants et al., 2011a), and direct estimation of cortical thickness in volumetric space (Das, Avants, Grossman, & Gee, 2009). Large-scale evaluation studies have shown that this highly accurate procedure for estimating CT is more sensitive to individual differences over the lifespan than comparable techniques (Tustison et al., 2014). CT images were down-sampled to 2 mm voxels before applying non-negative matrix factorization, but no additional smoothing was performed.

Non-negative matrix factorization

Cortical thickness was estimated as described above over the entire cortical surface. We sought to reduce CT in our sample into fewer dimensions, for two reasons. First, an efficient summary of CT data would allow us to evaluate only a small number of associations, rather than conduct voxel-wise inference that may be vulnerable to substantial Type I error (Eklund, Nichols, & Knutsson, 2016). Second, and importantly, prior work has shown that there are inherent patterns of covariance in CT (Alexander-Bloch, Giedd, & Bullmore, 2013; Sotiras et al., 2015, 2017; Zielinski, Gennatas, Zhou, & Seeley, 2010), and analyzing the data according to this covariance structure may enhance interpretability.

Accordingly, we achieved both goals by using non-negative matrix factorization (NMF) to identify structural networks where cortical thickness co-varies consistently

across individuals and brain regions (Sotiras et al., 2015). NMF has previously been shown to yield more interpretable and reproducible components than other decomposition techniques such as Principal Component Analysis or Independent Component Analysis (Sotiras et al., 2015, 2017). In contrast to the other techniques, NMF only yields compact networks with positive weights, which facilitates interpretation of effects.

The NMF algorithm takes as input a matrix X containing voxel-wise CT estimates (dimensions: 128,155 voxels x 427 participants), and approximates that matrix as a product of two matrices with non-negative elements: $X \approx BC$ (Figure 1). The first matrix, B , is of size $V \times K$ and contains the estimated non-negative networks and their respective loadings on each of the V voxels, where K is the user-specified number of networks. The B matrix (“CT loadings”) is composed of coefficients that denote the relative contribution of each voxel to a given network. These non-negative coefficients of the decomposition by necessity represent the entirety of the brain as a subject-specific addition of various parts. The second matrix, C , is of size $K \times N$ and contains subject-specific scores for each network. These subject-specific scores (“CT network scores”) indicate the contribution of each network in reconstructing the original CT map for each individual, and were evaluated for associations with DD as described below. We examined multiple NMF solutions requesting 2 to 30 networks (in steps of 2) and calculated reconstruction error for each solution as the Frobenius norm between the CT data matrix and its NMF approximation (Sotiras et al., 2015, 2017). The optimal number of components was chosen based on the elbow of the gradient of the reconstruction error, such that the solution is adequate to model the structure of the data without modeling random noise

(Sotiras et al., 2017). Network loadings were visualized on the inflated Population-Average, Landmark-, and Surface-based (PALS) cortical surfaces (Van Essen, 2005) using Caret software (Van Essen et al., 2001).

Experimental design and statistical analyses

To examine associations between DD and brain structure, we used a cross-sectional sample of youths recruited as part of a large neurodevelopmental study. As described above, our analysis sample consisted of 427 young participants who had usable data from both the DD task and structural neuroimaging.

Brain development is frequently a nonlinear process (Giedd et al., 1999; Lenroot et al., 2007; Satterthwaite et al., 2014b). In order to capture both linear and nonlinear age effects, we modeled age with a penalized spline within Generalized Additive Models (GAMs; Wood, 2004; 2011; Vandekar et al., 2015). In this type of model, a penalty is assessed on nonlinearity using restricted maximum likelihood in order to avoid overfitting. GAMs were implemented in the R package ‘mgcv’ (<https://cran.r-project.org/web/packages/mgcv/index.html>).

GAMs were first used to test for associations between DD and demographic variables such as age and sex. Next, we evaluated the association between DD and cognitive performance (as summarized by the overall cognitive performance factor and three domain-specific factor scores described above), while co-varying for sex and age. In both sets of analyses, DD was used as the dependent variable. Finally, univariate associations between DD and NMF-derived structural covariance networks were evaluated, with CT scores as the dependent variables and controlling for sex and age.

Interactions between DD and age were evaluated but were not found to be significant, and were thus not included in the univariate models. To control multiple testing across either cognitive factors or structural covariance networks, we used the False Discovery Rate (FDR, $Q < 0.05$; Benjamini & Hochberg, 1995).

In order to ensure that our results were not driven by potentially confounding factors, we conducted several sensitivity analyses. First, to ensure that our results were not driven by socio-economic status (SES) or non-specific neurostructural effects, we repeated these analyses while including maternal education and total brain volume as model covariates in separate models. Second, we repeated our analyses while excluding participants who were taking a psychotropic medication at the time of scan ($n = 52$) or for whom medication data was not available ($n = 3$) to ensure that these participants did not bias the observed results.

Multivariate analyses

The analyses described above examined univariate associations between each structural covariance network and DD. As a final step, we also investigated the multivariate predictive power of all cortical networks considered simultaneously, over and above that of a reduced model that included only demographics and cognitive data. The full model predicted DD using all 19 NMF networks, as well as age, sex, and the cognitive factors that were significantly associated with DD. This full model was compared to the reduced model (without the CT networks) using an F-test.

Results

Impulsive choice is associated with reduced cognitive performance

Mean discount rate in our sample was 0.073 ± 0.088 . Delay discounting was not related to demographic variables including age ($p = .387$). There was a non-significant trend toward more impulsive discounting in males ($p = .07$), and this trend was most prominent at younger ages (age by sex interaction: $p = .09$). In contrast, delay discounting was significantly associated with cognitive performance: youth who had higher discount rates also tended to have lower overall cognitive performance (partial $r = -.26$, $p < .0001$). Follow-up analyses with a three-factor model describing specific cognitive domains revealed that this effect was driven primarily by an association with a combined executive functioning and complex reasoning factor (partial $r = -.29$, $p < .0001$). Greater discounting was also associated with diminished memory accuracy (partial $r = -.20$, $p < .0001$), whereas there was no significant relationship between DD and social cognition (partial $r = -.08$, $p = .10$).

Non-negative matrix factorization identifies structural covariance networks

Next, we sought to identify structural covariance networks in CT using NMF. NMF provides a data-driven way to identify structural covariance networks, where cortical thickness varies in a consistent way across individuals. As NMF identifies structural networks at a resolution set by the user, we examined solutions ranging from 2 to 30 networks (in steps of 2). As expected, reconstruction error consistently decreased as the number of networks increased. Similar to previous applications of this method

(Sotiras et al., 2015), reconstruction error stabilized at 20 networks (Figure 2).

Accordingly, the 20-network solution was used for all subsequent analyses (Figure 3).

As in prior work using NMF (Sotiras et al., 2017), the structural covariance networks identified were highly symmetric bilaterally. Networks included specific cortical regions that are relevant to decision-making, such as ventromedial prefrontal cortex (vmPFC) and orbitofrontal cortex (OFC). Notably, when combined, several of the networks corresponded to aspects of functional brain networks. For example, networks 1 and 3 loaded on ACC and anterior insula, respectively, similar to the “salience network” (Seeley et al., 2007). Furthermore, specific networks defined lower-order systems, including motor (network 11) and visual (network 12) cortex. The 20-network solution also included a noise component (network 17), which was subsequently excluded from all analyses, resulting in 19 networks evaluated in total.

Greater delay discounting is associated with diminished cortical thickness

Having identified 19 interpretable structural covariance networks using NMF, we next examined associations with delay discounting while controlling for sex as well as linear and nonlinear age effects using penalized splines. Univariate analyses revealed that there was a significant association (after FDR correction) in eleven networks (Table 1). In each of these networks, impulsive choice, indicated by high discount rates, was associated with diminished cortical thickness. Notably, the strongest effects were found in two networks including the ventromedial prefrontal cortex and orbitofrontal cortex, both regions known to be critical for reward-related decision-making. These two networks also included parts of the temporal pole and temporoparietal junction, TPJ

(networks 14 and 15; Figure 4). Other networks where DD was associated with reduced CT included the temporal poles (network 9), lateral (network 8) and posterior temporal lobe (network 20), dorsolateral prefrontal cortex (network 18), insula (network 3), fusiform gyrus (network 7), fronto-parietal cortex (network 11), and visual cortex (network 12).

Association between cortical thickness and delay discounting is independent of age-related changes in cortical thickness

Having established that individual differences in DD are associated with CT, we next examined whether this effect was moderated by age. Notably, there was no significant age by DD interaction on any network (median $p = .77$, range: .09—.94). Thus, age-related changes in CT were similar in both high and low discounters, but those with a higher discount rate had thinner cortex across the age range examined (Figure 5).

Sensitivity analyses provide convergent results

We conducted sensitivity analysis to evaluate potentially confounding variables including maternal education, total brain volume, and psychotropic medications. First, we examined if results could be explained by differences in maternal education, a proxy of socioeconomic status. Discount rate was significantly associated with maternal education (partial $r = -.164$, $p = .0007$), but including it in the model did not have a great impact on results. Specifically, 7 of 11 networks found to be related to DD remained FDR-significant, including the vmPFC and OFC networks; the other 4 networks trended towards significance ($p_{fdr} < .067$). Second, we examined the effect of total brain volume on our findings. After adding total brain volume as a covariate, 10 of 11 networks

remained FDR-significant for association with DD, with the remaining network showing a trend towards FDR-significance ($p_{fdr} = .0762$). Finally, we repeated this analysis after excluding 52 participants who were taking psychotropic medication at the time of scan and 3 participants for whom medication data were missing. Despite the reduced power of this smaller sample, 10 of 11 networks remained FDR-significant, with the final network showing a trend trend towards significance ($p_{fdr} = .0503$).

Covariance networks provide improved prediction of DD over demographic and cognitive data

The univariate analyses described above demonstrated that reduced CT in several structural covariance networks is associated with impulsive choice. Next, we tested whether a multivariate model including all structural networks could accurately predict DD on an individual basis. Delay discounting predicted from a model of CT scores in all 19 networks, as well demographic data (age and sex), was significantly correlated with actual delay discounting behavior ($r = .33, p < .0001$; Figure 6). Adding CT scores to a reduced model with demographics alone improved model fit ($F_{(405,424)} = 2.37, p = .001$); DD predicted from this reduced model with demographics only achieved a correlation of .097 ($p = .043$) with actual log k values.

Importantly, CT data also improved prediction above and beyond that achieved by cognitive predictors: adding CT scores to a model with cognitive performance as well as demographics improved the model fit ($F_{(403,422)} = 1.63, p = .047$). DD predicted from the reduced model with just demographics and cognition achieved a correlation of .31 ($p <$

.0001) between model-predicted and actual log k values, compared to a correlation of .40 ($p < .0001$) from the full model including CT data, cognitive data, and demographics.

Discussion

We examined associations between delay discounting and cortical thickness networks in a large adolescent sample. More impulsive preferences, as indexed by higher discounting, were associated with diminished CT in multiple networks. The strongest effects were found in OFC, vmPFC, temporal pole, and the TPJ. Associations between DD and brain structure did not vary over the age range studied, and could not be explained by confounding variables. Furthermore, consideration of structural networks improved prediction of DD above and beyond demographic and cognitive variables.

Structural covariance networks related to DD overlap with known functional networks

Greater discounting was associated with decreased cortical thickness in multiple structural networks. Relative to previous reports of both neurofunctional and neurostructural correlates of delay discounting (Bernhardt et al., 2014; Kable & Levy, 2015; Peters & Büchel, 2011), the effects we observed were fairly widespread across the brain. Notably, many of the regions encompassed by these networks correspond to findings from previous studies in adults, including functional networks known to be involved in DD. As hypothesized, we found associations between DD and CT in central elements of the valuation network, namely vmPFC (Bartra et al., 2013); the cognitive control network, including dlPFC (Peters & Büchel, 2011; Stanger et al., 2013); and the prospection network, involving the medial temporal cortex (Peters & Büchel, 2011).

While DD and CT relationships have not previously been evaluated in adolescents, one prior study documented diminished thickness in the ACC and medial PFC in association with greater DD in adults (Bernhardt et al., 2014). In addition to hypothesized effects, we also found associations between DD and CT in motor, somatosensory, and both early and higher-order visual cortices. Notably, when these effects were evaluated jointly in multivariate model, CT networks enhanced prediction of DD above and beyond demographic and cognitive variables. This result contributes to efforts in neuroeconomics to improve prediction of decision-making behavior using brain-based measures obtained independently of the behavior itself (Kable & Levy, 2015), and suggests that structural covariance networks may be a useful marker of impulsive choice in youth.

Results converge with data from lesion and neuromodulation studies

Although the negative associations between DD and CT were widespread and distributed, two structural covariance networks exhibited particularly strong associations with DD and robustness to all sensitivity analyses. Brain regions comprising these networks included vmPFC, OFC, temporal pole, and the TPJ. As mentioned above, our findings in vmPFC were expected based on substantial evidence from fMRI studies that this brain region is implicated in DD (Ballard & Knutson, 2009; Bartra et al., 2013; Kable & Glimcher, 2007). Furthermore, activity in vmPFC when merely thinking about the future predicts DD, such that lower discounters show greater activity when thinking about the far future (Cooper, Kable, Kim, & Zauberman, 2013). Finally, consistent with our results, a previous study in adults reported that diminished CT in that region was associated with higher DD (Bernhardt et al., 2014).

Beyond the vmPFC, there is evidence that regions including the OFC, temporal pole, and TPJ are both involved in and necessary for evaluating future outcomes in delay discounting. First, patients with medial OFC damage show greater discounting of both primary and secondary rewards, compared to healthy controls and non-frontal damage patients (Sellitto, Ciaramelli, & di Pellegrino, 2010), and this is the only region where lesions have been reported to increase discounting in humans. Notably, this relationship is dose-dependent, such that larger frontal lesions are associated with steeper discounting. Second, patients with semantic dementia, a disorder characterized by anterior temporal lobe atrophy, show greater discounting than controls (Chiong et al., 2015). Third, while the TPJ has typically been implicated in social cognition and theory of mind, recent data suggests it plays a role in both monetary and social discounting (Soutschek, Ruff, Strombach, Kalenscher, & Tobler, 2016; Strombach et al., 2015). Importantly, disrupting the TPJ in healthy adults using transcranial magnetic stimulation increased discounting (Soutschek et al., 2016). Collectively, this evidence suggests that the disruption of OFC, anterior temporal lobe, and TPJ may promote impulsive choice.

Associations with delay discounting are independent of age-related changes

While we replicated prior findings of association between lower discounting and higher IQ (Shamosh & Gray, 2008) and memory performance (Shamosh et al., 2008), we did not find significant associations between DD and age (Scheres et al., 2006; Steinberg et al., 2009). This may be due to differences in sample composition, including an older range being sampled and a dimensional rather than a stratified design that compared older and younger age groups. Notably, the association between brain structure and DD was

stable across the entire age range surveyed in our sample. This result is consistent with a prior study of DD in adolescents and white matter integrity assessed using diffusion imaging (Olson et al., 2009). Together, these results imply that individual differences in brain structure associated with impulsive choice do not emerge specifically during adolescence. These results may also suggest that such individual differences in brain structure may emerge early in development, consistent with literature describing the importance of structural brain development in utero, during the peri-natal period, and during early childhood (Di Martino et al., 2014; Thomason et al., 2013). While speculative, future research may reveal that individual differences in brain structure which emerge early in life may impact evolving patterns of value and cognitive control system function in adolescence which, in turn, may contribute to impulsivity during this critical period (Bjork, Smith, Chen, & Hommer, 2010; Casey et al., 2008).

Advantages of evaluating structural covariance networks in a large sample

The greater spatial extent of significant associations between brain structure and DD observed in our data compared to prior results may be due to several aspects of our study. First, the large sample size afforded greater statistical power, and thus greater sensitivity to detect effects in multiple brain networks. While the effect sizes of these associations were small, research documenting inflation of effect sizes in small studies suggests that the present results are more likely to be an accurate reflection of the true effect size than data from more modest samples (Button et al., 2013). Second, structural covariance networks defined by NMF provided a parsimonious summary of the high-dimensional imaging data that limited multiple comparisons. In contrast to anatomic

atlases based on sulcal folding patterns, NMF identifies structural brain networks based on patterns of covariance in the data itself. The concise summary of the data yielded by NMF limited multiple comparisons: we only evaluated 19 networks in our analyses, in contrast to the hundreds of thousands of voxels typically surveyed in mass-univariate VBM studies. This allowed us to use a rigorous FDR correction for all comparisons, rather than cluster-based inference that may produce substantial Type I error rates in many common implementations (Eklund et al., 2016).

Limitations

Certain limitations of this study should be noted. First, the observed effects were independent of age, suggesting that differences in brain structure associated with impulsive choice may emerge earlier than the examined age range. Future investigations should consider longitudinal designs including early childhood to precisely capture the emergence of these effects. Second, we were unable to directly test the associations with risky behaviors, such as tobacco, drug use and risky sexual behaviors. Subsequent work would benefit from the inclusion of such outcome measures and a direct evaluation of which specific DD-related networks predict increased risk-taking in adolescence. Third, we used hypothetical instead of real rewards in the DD task. However, prior studies have yielded similar results between real and hypothetical reward tasks in both behavioral (Johnson & Bickel, 2002) and functional neuroimaging paradigms (Bickel, Pitcock, Yi, & Angtuaco, 2009). Fourth, we cannot completely rule out potential confounding variables which may be correlated with DD. Previous studies have described associations between CT and SES in adolescence (Mackey et al., 2015), though importantly our

results remained largely unaffected after controlling for maternal education, a proxy of SES.

Conclusions and future directions

Understanding impulsive choice in adolescence is important because impulsivity is associated with a host of risky behaviors and outcomes, such as tobacco use (Audrain-McGovern et al., 2009; Reynolds, 2004), alcohol use (Fennie et al., 2013; Field et al., 2007), obesity (Fields, Sabet, & Reynolds, 2013) and early sexual initiation (Khurana et al., 2012), which lead to substantial morbidity and mortality during adolescence.





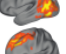
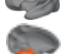

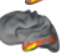




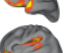
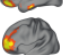





Leveraging a large developmental sample and advanced analytics, we found that individual variability in brain structure explains differences in DD in adolescence. Taken together, our results indicate that higher DD in youth is associated with reduced cortical thickness in multiple networks, including those known to be essential for valuation.

These results emphasize that risky behaviors in adolescents should be considered in the context of individual differences of structural brain networks that are present early in life.

Moving forward, such brain-based measures could potentially be used as biomarkers to identify youth at particularly high risk for negative outcomes, and aid in stratifying youth within targeted clinical trials that aim to reduce impulsivity and risk-taking behaviors during this critical period.

Tables

Table 1. Association between delay discounting and NMF-derived structural covariance networks. β (unstandardized regression coefficient), SE (β 's standard error), t (t-value for testing β against 0, $dfs = 423$), p -value, and FDR-corrected p -value are obtained from separate general additive models run for each network. In this model, discount rate ($\log k$) predicts cortical thickness scores, controlling for age (fit as a penalized spline) and sex. In order to provide an estimate of the effect size, r is the partial Pearson's correlation coefficient between discount rate and CT scores in each network, while adjusting for linear age, quadratic age, and sex. FDR-significant p -values are indicated in bold.

| | Network | β | SE | t | p | FDR- p | r |
|---|---------|---------|--------|-------|---------|-------------------|--------|
|  | Ntwk 1 | -0.649 | 0.3946 | -1.64 | 0.101 | 0.137 | -0.080 |
|  | Ntwk 2 | -0.0138 | 0.4217 | -0.03 | 0.974 | 0.974 | 0.002 |
|  | Ntwk 3 | -1.5868 | 0.5606 | -2.83 | 0.005 | 0.019 | -0.136 |
|  | Ntwk 4 | -0.4337 | 0.6414 | -0.68 | 0.499 | 0.527 | -0.033 |
|  | Ntwk 5 | -0.9959 | 0.4811 | -2.07 | 0.039 | 0.062 | -0.100 |
|  | Ntwk 6 | -0.8277 | 0.5337 | -1.55 | 0.122 | 0.154 | -0.075 |
|  | Ntwk 7 | -1.1428 | 0.4359 | -2.62 | 0.009 | 0.024 | -0.126 |
|  | Ntwk 8 | -1.1598 | 0.4562 | -2.54 | 0.011 | 0.024 | -0.123 |
|  | Ntwk 9 | -0.7926 | 0.3580 | -2.21 | 0.027 | 0.047 | -0.110 |
|  | Ntwk 10 | -0.3748 | 0.3055 | -1.23 | 0.221 | 0.262 | -0.060 |
|  | Ntwk 11 | -1.1527 | 0.4669 | -2.47 | 0.014 | 0.027 | -0.119 |
|  | Ntwk 12 | -1.5839 | 0.6164 | -2.57 | 0.011 | 0.024 | -0.124 |
|  | Ntwk 13 | -1.173 | 0.4283 | -2.74 | 0.006 | 0.02 | -0.132 |
|  | Ntwk 14 | -2.019 | 0.4241 | -4.76 | <0.0001 | <0.0001 | -0.225 |
|  | Ntwk 15 | -1.257 | 0.3036 | -4.14 | <0.0001 | <0.0001 | -0.200 |
|  | Ntwk 16 | -0.4404 | 0.4371 | -1.01 | 0.314 | 0.351 | -0.050 |
|  | Ntwk 18 | -1.252 | 0.4305 | -2.91 | 0.004 | 0.018 | -0.140 |
|  | Ntwk 19 | -0.7172 | 0.3713 | -1.93 | 0.054 | 0.079 | -0.094 |
|  | Ntwk 20 | -0.8778 | 0.3014 | -2.91 | 0.004 | 0.018 | -0.140 |

!

Figures

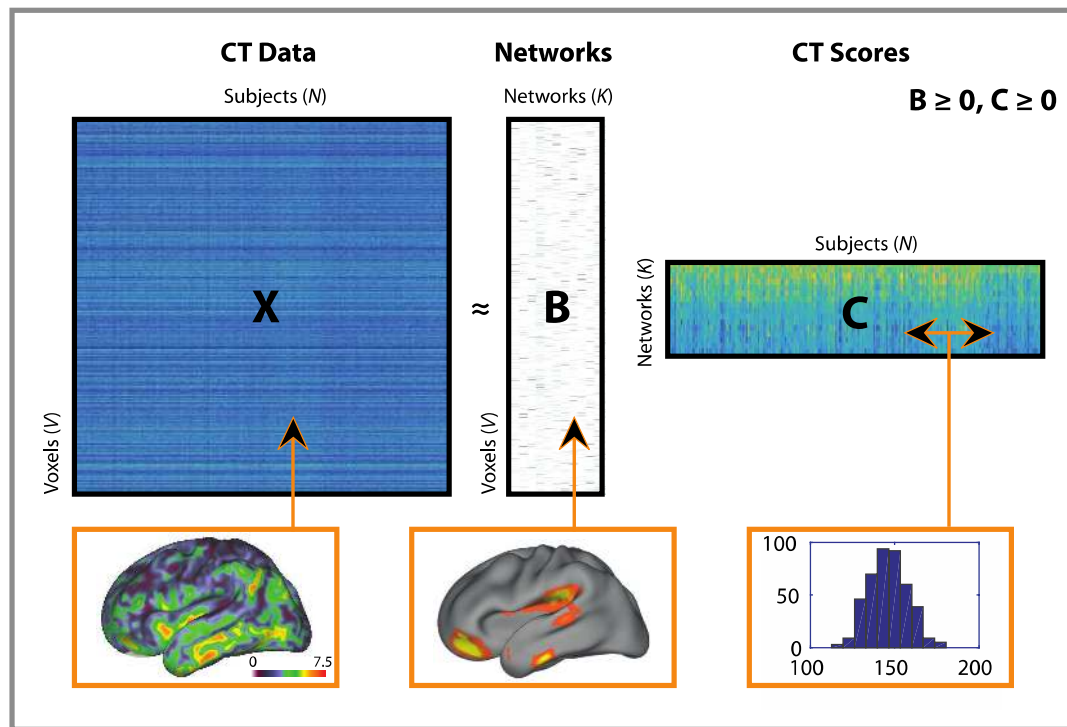


Figure 1. Schematic of non-negative matrix factorization and example data for each matrix. X is the original matrix with a sample of actual cortical thickness data; Sample CT map shows example CT data from one participant, and corresponds to a column in the X matrix; B is matrix with a sample of actual estimated networks and their loadings on each voxel; Loadings map shows example loadings from one network in NMF solution used in our analyses, and corresponds to a column in the B matrix; C is a matrix with actual subject-specific weights for each network; Histogram shows CT Scores in same sample network as visualized in B , and corresponds to a row in the C matrix. Matrix sizes are shown with following dimensions: V = number of cortical thickness voxels, N = number of participants; K = number of networks.




Figure 2. NMF reconstruction error identifies 20 cortical networks as the optimal parcellation resolution for cortical thickness data. Plot of reconstruction error gradient for NMF at multiple resolutions; the gradient is the difference in reconstruction error as the NMF solution increases by 2 networks. Blue circle indicates selected NMF solution of 20 networks.

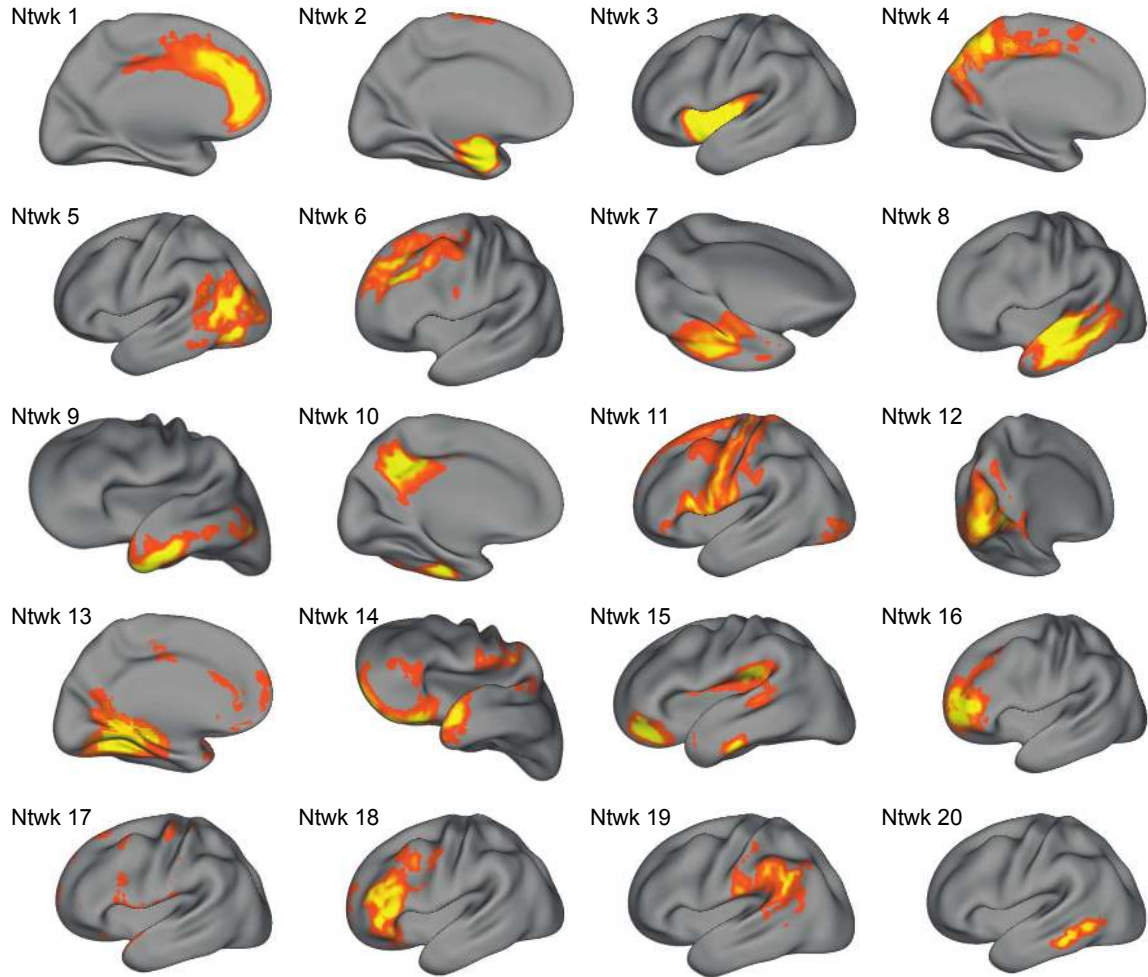


Figure 3. Structural covariance networks delineated by NMF. Visualization of structural covariance networks from the 20-network NMF solution. The spatial distribution of each network is indicated by loadings at each voxel in arbitrary units (from B matrix in NMF factorization); warmer colors represent higher loadings. For each network, we show one view that best captures the main area(s) of coverage. Approximate anatomical coverage of each structural covariance network: 1) medial prefrontal cortex and cingulate cortex; 2) medial temporal lobe; 3) insula; 4) medial posterior parietal cortex, including the precuneus; 5) temporo-occipital cortex; 6) dorsolateral prefrontal cortex (dlPFC); 7) fusiform gyrus; 8) lateral temporal lobe; 9) lateral temporal lobe and temporal pole; 10) posterior cingulate cortex and temporal lobe; 11) frontal and parietal cortex, including primary motor and somatosensory cortices; 12) occipital cortex; 13) medial temporal cortex, anterior cingulate cortex (ACC) and posterior cingulate cortex (PCC); 14) orbitofrontal cortex (OFC), frontal and temporal poles; 15) ventromedial prefrontal cortex (vmPFC), inferior temporal lobe, auditory cortex, temporoparietal junction (TPJ); 16) dorsal OFC; 17) the dura matter, a noise component that was not evaluated further; 18) dlPFC; 19) angular and supramarginal gyri; 20) posterior inferior temporal lobe.




Figure 4. Higher discounting is associated with diminished cortical thickness in frontal, temporal, and parietal areas. Regions of FDR-significant association between $\log k$ and structural covariance networks. The composite network visualization was obtained by assigning each voxel to the network which had the highest loading for that voxel (from the B matrix), across all 19 networks. Maximal effects were observed in Networks 14 and 15, which included orbitofrontal cortex and ventromedial prefrontal cortex. Scatterplots for $\log k$ -CT association in these networks are shown, while adjusting for model covariates. Gray envelope represents the 95% confidence interval.

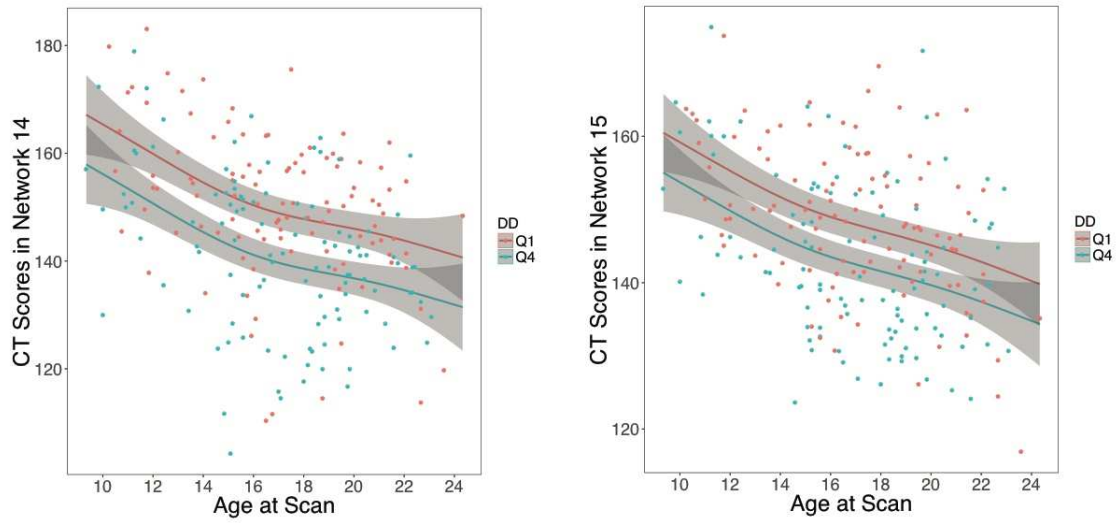


Figure 5. Association between cortical thickness and delay discounting is independent of age. Scatterplots for relationship between age and CT in networks 14 and 15, separated by top (Q4) and bottom (Q1) quartiles of $\log k$. The Q4 quartile group contains participants with the most impulsive preferences. For each quartile, the age-CT relationship is shown after adjusting for model covariates, and includes the 95% confidence intervals (gray envelopes).




Figure 6. CT data from structural covariance networks predicts delay discounting.

Scatterplot for relationship between actual log k values and predicted log k from multivariate CT prediction. Multivariate prediction is based on CT scores from all structural covariance networks plus demographic variables, sex, and age. Scatterplots include line of best fit for this association with a 95% confidence interval (gray envelope).

CHAPTER 4 — General discussion

Overall summary

The research presented in this dissertation furthers our understanding of individual differences in two types of decision-making: adaptive learning (Chapter 2) and delay discounting (Chapter 3). Results presented in Chapter 2 showed that individuals across the lifespan exhibit large variability in using normative learning strategies in a dynamic environment, and these differences relate to cognitive, but not personality, factors. Specifically, better memory performance was associated with more “adaptive” learning, in both young (age 18-35, Experiment 3) and elderly (age 60+, Experiment 2) participants. In addition, adaptive learning in elderly adults was less influenced by uncertainty, but also less influenced by a non-normative reward factor. Reliance on normative learning factors has previously been associated with neural activity in the ventromedial prefrontal cortex (vmPFC), lateral prefrontal cortex (LPFC), and medial temporal lobe, among other regions (McGuire, Nassar, Gold, & Kable, 2014). Next, we directly investigated the neural bases of delay discounting and found structural effects in some of these regions. Results presented in Chapter 3 showed that cortical thickness across multiple networks of brain regions was associated with individual differences in delay discounting in a large sample of adolescents. Specifically, diminished cortical thickness in frontal and temporal regions was associated with higher discounting, i.e., more impulsive choices. The strongest effects were found in regions typically implicated in delay discounting, such as the vmPFC and orbitofrontal cortex. Furthermore, cortical thickness data predicted differences in delay discounting above and beyond cognitive

variables typically found to correlate with delay discounting. Broader implications and future directions of this research are discussed below.

Influences of cognitive factors and arousal on adaptive learning

In Chapter 2, across three experiments, adaptive learning behavior, and specifically the degree of reliance on normative learning factors, was positively associated with cognitive abilities— memory and IQ— but not personality. Notably, there were no clear links between learning and personality traits related to affective arousal, e.g., anxiety and neuroticism, even though arousal has been shown to play a role in adaptive learning. Prior work with a simpler version of the predictive-inference task reported that arousal sensitivity to task statistics, as measured by pupillary responses, reflects behavioral sensitivity to normative learning factors (Nassar et al., 2012). Specifically, the degree to which pupil diameter tracks normative learning factors predicts learning rates across participants. Pupil dilation is considered a marker of affective arousal (Bradley, Miccoli, Escrig, & Lang, 2008; Partala & Surakka, 2003), and its response to emotionally arousing stimuli can be modulated by anxiety (Bertrand, Garcia, Viera, Santos, & Bertrand, 2013) and neuroticism (Prehn et al., 2008). However, in addition to being an indicator of affective arousal states, pupil diameter has also been linked to both state and trait cognitive characteristics. For example, pupil diameter tracks the amount of material being processed in memory (Kahneman & Beatty, 1966), and baseline pupil diameter is positively correlated with working memory capacity across individuals (Heitz, Schrock, Payne, & Engle, 2008). Furthermore, pupil dilation while performing cognitive tasks is associated with intelligence (Ahern & Beatty, 1979; Van

Der Meer et al., 2010), such that higher intelligence contributes to more efficient cognitive processing or better access to cognitive resources. Given pupil diameter's associations with cognitive abilities and processing, the link between our findings and previous findings of pupil-linked arousal's influence on adaptive learning (Nassar et al., 2012) is intriguing. Is the role of arousal in adaptive learning reflective of cognitive or affective influences on normative learning? Future work is needed to characterize the interplay of cognitive abilities and arousal in this task. Specifically, it would be interesting to measure cognitive abilities (memory and intelligence) and arousal in the same setting, and test whether cognitive abilities modulate arousal in response to environmental statistics. This work could also be replicated and extended by using alternative measures of arousal, such as galvanic skin response and heart rate variability.

The role of memory in value-based decision-making

The data presented in this dissertation point to memory as an important source of individual differences in decision-making. In the adaptive learning framework (Chapter 2), better memory performance, across different age groups and different memory tasks, was associated with higher reliance on normative learning factors. In the study of adolescent delay discounting (Chapter 3), we replicated a prior association with memory (Shamosh et al., 2008), while also using a different memory task (tapping episodic versus working memory). Recent integrative work has begun to focus on the relationship between memory and value-based decision-making— two research domains that have mostly been studied in isolation of each other (Euston, Gruber, & McNaughton, 2012; Shadlen & Shohamy, 2016; Weilbacher & Gluth, 2016). Episodic memory, including

encoding and retrieval of such memories, is most often associated with the medial temporal lobe, and specifically the hippocampus (Eichenbaum, 2004). In contrast, value-based decision-making is robustly associated, in part, with the vmPFC, where neural activity in fMRI experiments tracks the subjective value of choice stimuli (Bartra, McGuire, & Kable, 2013; Clithero & Rangel, 2013). A recent review article by Shadlen and Shohamy (2016) argues that many value-based decisions entail sequential sampling of evidence from memory until a decision threshold is reached, even for decisions where the choices need not be retrieved from memory. The authors hypothesize that this sequential process is implemented via an interaction between vmPFC and hippocampus, where the hippocampus “updates” a decision variable encoded in vmPFC. Furthermore, in their view, memory is particularly important to value-based decisions that depend either on the integration of distinct past events or on prospection about future events. Accordingly, the data presented herein fit into this framework in two ways.

First, in the adaptive learning study, a plausible explanation for the effect of memory ability on normative learning would be that retrieving past outcomes is necessary for gradually building a mental model of the environment (akin to model-based learning). Related to the above-proposed framework (Shadlen & Shohamy, 2016), uncertainty updating in the helicopter task is reliant on activity in both medial temporal lobe and vmPFC (McGuire et al., 2014). In healthy aging, both of these regions may become less “sensitive” because of age-related dopamine receptor loss (Mohr, Li, & Heekeren, 2010; Kaasinen et al., 2000) or brain volume atrophy (Coffey et al., 1992),

which may explain differences in uncertainty updating between young and old participants in this study.

Second, in delay discounting, episodic memory could play a role through its relation to prospection, i.e., episodic future thinking. Indeed, much research has shown that the cognitive processes of remembering past events and imagining possible future events are related, and are dependent on a similar network of brain regions, including the hippocampus and medial prefrontal cortex (for a review, see Schacter, Addis, & Buckner, 2008). In the context of delay discounting, previous work has shown an explicit link between past and future, where individuals discount events into the past to a similar degree as they discount events into the future, and both discounting phenomena exhibit similar characteristics (Yi, Gatchalian, & Bickel, 2006; Yi, Landes, & Bickel, 2009). Moreover, episodic thinking about the future has been identified as one of several cognitive processes implicated in delay discounting (Peters and Büchel, 2011). Similarly to the proposed mechanism of memory's role in value-based decision-making (Shadlen & Shohamy, 2016), prospection can reduce discounting through a functional interaction between hippocampus and medial prefrontal cortex (Peters & Büchel, 2010). Notably, although the connection between structure and function in these areas is unclear, cortical thickness in both medial prefrontal and, to a lesser extent, medial temporal areas was negatively associated with discounting, but we did not specifically test if episodic memory is at all related to these structural effects.

More broadly, combined results from both studies contribute to a growing body of work showing that memory is important in value-based decision making. To the extent

that lower discounting and more adaptive learning in changing environments are desirable, better memory appears to confer additional benefits to individuals.

Memory versus IQ effects

In addition to memory effects, results from the adaptive learning study pointed to intelligence as another cognitive factor that predicts normative learning (although the effect was found in only one of the experiments). In the delay discounting study, higher intelligence was associated with lower discounting, as has been previously established (Shamosh & Gray, 2008). Intelligence is construed as a general factor that contributes to a variety of cognitive tasks, and is highly positively correlated with both episodic and working memory, but many have argued that intelligence and memory are distinct constructs (Ackerman, Beier, & Boyle, 2005; Conway, Kane, & Engle, 2003; Healey, Crutchley, & Kahana, 2014). Given this association and the finding that both constructs may contribute to decision-making, the question arises whether these effects are specific to one of the constructs or, alternatively, indicative of general effects of superior cognitive functioning. In the adaptive learning framework, we were unable to conclusively answer this question. Future experiments could address this distinction between specific memory versus intelligence effects or general cognitive effects in at least two ways.

A possible experiment with lesion patients can provide causal evidence. Performance of hippocampal patients with memory impairments can be compared to that of healthy controls, and a group with dorsolateral prefrontal cortex (dlPFC) lesions. dlPFC is a core

part of a network of regions supporting intelligence (Gläscher et al., 2010; Kane & Engle, 2002). If one of the lesion groups shows a larger impairment relative to the healthy controls, that would suggest that the related process has a primary influence on normative learning. Future work will also benefit from using samples of individuals who are more representative of the average population in terms of cognitive abilities, thus increasing variability and limiting ceiling effects. Understanding the specific cognitive processes that underlie adaptive learning could contribute to designing interventions targeted at improving decision-making in dynamic environments.

Brain-as-predictor approach

An important extension of the work described in Chapter 3 would be to use brain data to predict real-life outcomes; specifically, using cortical thickness networks related to delay discounting to predict risk-taking behavior in adolescents. As mentioned previously, high impulsivity, operationalized as steep discounting, has been linked to a host of maladaptive behaviors and outcomes, in both adolescents and adults. It is thus plausible that structure of brain regions implicated in delay discounting may be associated with these behaviors. This approach of utilizing fundamental neuroscience findings that identify neural bases of psychological processes for the prediction of real-life outcomes that are related to these processes has been termed the “brain-as-predictor” approach (Berkman & Falk, 2013). More broadly, improved prediction of human economic behavior from neural data has been one of the goals of the field of neuroeconomics (Kable & Levy, 2015).

Using neural data directly to predict real-life outcomes can be beneficial to the extent that it (1) explains variability that has not been accounted for by behavioral or self-report measures, especially for measures that are subject to considerable measurement noise; and (2) elucidates specific sub-processes or mechanisms that are associated with the real-life outcome (Berkman & Falk, 2013). For example, we found that cortical thickness in several brain regions associated with distinct cognitive processes is associated with delay discounting, but it is unclear if these processes contribute to risk-taking behavior equally. If successful, a future study showing that brain structure in these networks of regions directly predicts risk-taking in adolescence offers the promise of a fundamental neuroscience contribution to the field of public health. Prediction can potentially be further improved by adding other brain measures that carry a signal related to delay discounting, and potentially combining structural and functional data, as has previously been done in predicting intelligence (Choi et al., 2008).

Decision-making across the lifespan

Age is an important factor in the research presented in this dissertation. Stages of life such as adolescence and old age are accompanied by significant changes in brain structure (Buckner, 2004; Coffey et al., 1992; Giedd et al., 1999; Lenroot et al., 2007; Raz et al., 2005) and cognitive abilities (Blakemore & Choudhury, 2006; MacPherson, Phillips, & Della Sala, 2002; McArdle, Ferrer-Caja, Hamagami, & Woodcock, 2002), which could play a role in individual differences in decision-making. Results from the adaptive learning study (Chapter 2) showed that older adults used learning strategies that were less influenced by a normative learning factor related to uncertainty, compared to

younger adults. This finding contributes to a large literature showing age-related differences in both laboratory (Mata, Josef, Samanez-Larkin, & Hertwig, 2011; Samanez-Larkin & Knutson, 2015; Tymula, Belmaker, Ruderman, Glimcher, & Levy, 2013) and real-life decision-making (Finucane, Slovic, Hibbard, Peters, Mertz, & MacGregor, 2002; Thornton & Dumke, 2005). More specifically, since increased learning from uncertainty in the helicopter task is in part associated with higher BOLD activity in lPFC (McGuire et al., 2014), our findings are consistent with a report that older adults are impaired on tasks dependent on dlPFC function (MacPherson et al., 2002). The delay discounting study (Chapter 3) was specifically focused on adolescents because of the important neurostructural changes during this developmental period and the increased risk of impulsive behaviors and outcomes that are associated with higher discounting. In a large adolescent sample, cortical thickness associations with delay discounting showed some similarities with findings in adults, including neurostructural and neurofunctional findings in vmPFC, OFC, and ACC (Bernhardt et al., 2014; Bjork, Momenan, & Hommer, 2009; Kable & Glimcher, 2007; Sellitto, Ciaramelli, & di Pellegrino, 2010), among other regions. Yet, future studies will be required to directly test if the same cortical thickness networks are associated with delay discounting in adults.

Conclusion

Decision-making is ubiquitous in life and humans show large variability in choices they make and what decision-making strategies they use. The research presented here focused on two domains of decision-making and investigated factors that explain

differences between individuals. Across two studies, evidence showed that cognitive function, age, and brain structure are important factors to consider when studying decision-making. Understanding the factors that drive differences in decision-making is important, so as to design training for more advantageous decision-making.

BIBLIOGRAPHY

- Ackerman, P. L., Beier, M. E., & Boyle, M. O. (2005). Working memory and intelligence: The same or different constructs?. *Psychological Bulletin*, 131(1), 30-60.
- Ahern, S., & Beatty, J. (1979). Pupillary responses during information processing vary with Scholastic Aptitude Test scores. *Science*, 205(4412), 1289-1292.
- Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological Bulletin*, 82(4), 463-496.
- Alexander-Bloch, A., Giedd, J. N., & Bullmore, E. (2013). Imaging structural covariance between human brain regions. *Nature Reviews Neuroscience*, 14(5), 322-336.
- Amlung, M., Petker, T., Jackson, J., Balodis, I., & MacKillop, J. (2016). Steep discounting of delayed monetary and food rewards in obesity: a meta-analysis. *Psychological Medicine*, 46(11), 2423-2434.
- Anderson, J., Burks, S., DeYoung, C., & Rustichini, A. (2011). *Toward the integration of personality theory and decision theory in the explanation of economic behavior*. Unpublished manuscript.
- Aroian, L. A. (1947). The probability function of the product of two normally distributed variables. *The Annals of Mathematical Statistics*, 18, 265-271.
- Audrain-McGovern, J., Rodriguez, D., Epstein, L. H., Cuevas, J., Rodgers, K., & Wileyto, E. P. (2009). Does delay discounting play an etiological role in smoking or is it a consequence of smoking?. *Drug and Alcohol Dependence*, 103(3), 99-106.
- Avants, B. B., Tustison, N. J., Song, G., Cook, P. A., Klein, A., & Gee, J. C. (2011a). A reproducible evaluation of ANTs similarity metric performance in brain image registration. *NeuroImage*, 54(3), 2033-2044.
- Avants, B. B., Tustison, N. J., Wu, J., Cook, P. A., & Gee, J. C. (2011b). An open source multivariate framework for n-tissue segmentation with evaluation on public data. *Neuroinformatics*, 9(4), 381-400.
- Ballard, K., & Knutson, B. (2009). Dissociable neural representations of future reward magnitude and delay during temporal discounting. *NeuroImage*, 45(1), 143-150.

- Barrick, M. R., & Mount, M. K. (1991). The big five personality dimensions and job performance: a meta-analysis. *Personnel Psychology*, 44(1), 1-26.
- Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: a coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, 76, 412-427.
- Batty, G. D., Deary, I. J., & Gottfredson, L. S. (2007). Premorbid (early life) IQ and later mortality risk: systematic review. *Annals of Epidemiology*, 17(4), 278-288.
- Becker, A., Deckers, T., Dohmen, T., Falk, A., & Kosse, F. (2012). The relationship between economic preferences and psychological personality measures. *Annual Review of Economics*, 4(1), 453-478.
- Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9), 1214-1221.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1), 289-300.
- Berkman, E. T., & Falk, E. B. (2013). Beyond brain mapping using neural measures to predict real-world outcomes. *Current Directions in Psychological Science*, 22(1), 45-50.
- Bernhardt, B. C., Smallwood, J., Tusche, A., Ruby, F. J., Engen, H. G., Steinbeis, N., & Singer, T. (2014). Medial prefrontal and anterior cingulate cortical thickness predicts shared individual differences in self-generated thought and temporal discounting. *NeuroImage*, 90, 290-297.
- Bertrand, A.L., Garcia, J.B.S., Viera, E.B., Santos, A.M., & Bertrand, R.H. (2013). Pupillometry: the influence of gender and anxiety on the pain response. *Pain Physician*, 16(3), E257-E266.
- Bickel, W. K., Miller, M. L., Yi, R., Kowal, B. P., Lindquist, D. M., & Pitcock, J. A. (2007). Behavioral and neuroeconomics of drug addiction: competing neural systems and temporal discounting processes. *Drug and Alcohol Dependence*, 90, S85-S91.

- Bickel, W. K., Pitcock, J. A., Yi, R., & Angtuaco, E. J. (2009). Congruence of BOLD response across intertemporal choice conditions: fictive and real money gains and losses. *Journal of Neuroscience*, 29(27), 8839-8846.
- Bjork, J. M., Momenan, R., & Hommer, D. W. (2009). Delay discounting correlates with proportional lateral frontal cortex volumes. *Biological Psychiatry*, 65(8), 710-713.
- Bjork, J. M., Smith, A. R., Chen, G., & Hommer, D. W. (2010). Adolescents, adults and rewards: comparing motivational neurocircuitry recruitment using fMRI. *PloS One*, 5(7), e11440.
- Blakemore, S. J., & Choudhury, S. (2006). Development of the adolescent brain: implications for executive function and social cognition. *Journal of Child Psychology and Psychiatry*, 47(3-4), 296-312.
- Bobova, L., Finn, P. R., Rickert, M. E., & Lucas, J. (2009). Disinhibitory psychopathology and delay discounting in alcohol dependence: personality and cognitive correlates. *Experimental and Clinical Psychopharmacology*, 17(1), 51-61.
- Bradley, M. M., Miccoli, L., Escrig, M. A., & Lang, P. J. (2008). The pupil as a measure of emotional arousal and autonomic activation. *Psychophysiology*, 45(4), 602-607.
- Browning, M., Behrens, T. E., Jocham, G., O'Reilly, J. X., & Bishop, S. J. (2015). Anxious individuals have difficulty learning the causal statistics of aversive environments. *Nature Neuroscience*, 18(4), 590-596.
- Burks, S. V., Carpenter, J. P., Goette, L., & Rustichini, A. (2009). Cognitive skills affect economic preferences, strategic behavior, and job attachment. *Proceedings of the National Academy of Sciences*, 106(19), 7745-7750.
- Bush, R. R., & Mosteller, F. (1951). A mathematical model for simple learning. *Psychological Review*, 58(5), 313-323.

- Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5), 365-376.
- Casey, B. J., Jones, R. M., & Hare, T. A. (2008). The adolescent brain. *Annals of the New York Academy of Sciences*, 1124(1), 111-126.
- Chabris, C. F., Laibson, D., Morris, C. L., Schuldt, J. P., & Taubinsky, D. (2008). Individual laboratory-measured discount rates predict field behavior. *Journal of Risk and Uncertainty*, 37(2-3), 237-269.
- Chao, L. W., Szrek, H., Pereira, N. S., & Pauly, M. V. (2009). Time preference and its relationship with age, health, and survival probability. *Judgment and Decision Making*, 4(1), 1-19.
- Chiong, W., Wood, K. A., Beagle, A. J., Hsu, M., Kayser, A. S., Miller, B. L., & Kramer, J. H. (2016). Neuroeconomic dissociation of semantic dementia and behavioural variant frontotemporal dementia. *Brain*, 139(2), 578-587.
- Cho, S. S., Pellecchia, G., Aminian, K., Ray, N., Segura, B., Obeso, I., & Strafella, A. P. (2013). Morphometric correlation of impulsivity in medial prefrontal cortex. *Brain Topography*, 26(3), 479-487.
- Choi, Y. Y., Shamosh, N. A., Cho, S. H., DeYoung, C. G., Lee, M. J., Lee, J. M., ... & Lee, K. H. (2008). Multiple bases of human intelligence revealed by cortical thickness and neural activation. *Journal of Neuroscience*, 28(41), 10323-10329.
- Clithero, J. A., & Rangel, A. (2014). Informatic parcellation of the network involved in the computation of subjective value. *Social Cognitive and Affective Neuroscience*, 9(9), 1289-1302.
- Coffey, C. E., Wilkinson, W. E., Parashos, L. A., Soady, S. A. R., Sullivan, R. J., Patterson, L. J., ... & Djang, W. T. (1992). Quantitative cerebral anatomy of the aging human brain A cross-sectional study using magnetic resonance imaging. *Neurology*, 42(3), 527-527.

- Coffey, S. F., Gudleski, G. D., Saladin, M. E., & Brady, K. T. (2003). Impulsivity and rapid discounting of delayed hypothetical rewards in cocaine-dependent individuals. *Experimental and Clinical Psychopharmacology*, 11(1), 18-25.
- Conway, A. R., Kane, M. J., & Engle, R. W. (2003). Working memory capacity and its relation to general intelligence. *Trends in Cognitive Sciences*, 7(12), 547-552.
- Cools, R., Frank, M. J., Gibbs, S. E., Miyakawa, A., Jagust, W., & D'Esposito, M. (2009). Striatal dopamine predicts outcome-specific reversal learning and its sensitivity to dopaminergic drug administration. *Journal of Neuroscience*, 29(5), 1538-1543.
- Cooper, A. J., Duke, É., Pickering, A. D., & Smillie, L. D. (2014). Individual differences in reward prediction error: contrasting relations between feedback-related negativity and trait measures of reward sensitivity, impulsivity and extraversion. *Frontiers in Human Neuroscience*, 8(248), 40-50.
- Cooper, N., Kable, J. W., Kim, B. K., & Zauberman, G. (2013). Brain activity in valuation regions while thinking about the future predicts individual discount rates. *Journal of Neuroscience*, 33(32), 13150-13156.
- Costa Jr, P. T., & McCrae, R. R. (1992). *Neo personality inventory-revised (neo-pi-r) and neo five-factor inventory (neo-ffi) professional manual*. Odessa, FL: Psychological Assessment Resources.
- Courtney, K. E., Arellano, R., Barkley-Levenson, E., Gálvan, A., Poldrack, R. A., MacKillop, J., ... & Ray, L. A. (2012). The relationship between measures of impulsivity and alcohol misuse: an integrative structural equation modeling approach. *Alcoholism: Clinical and Experimental Research*, 36(6), 923-931.
- Craik, F. I., & Bialystok, E. (2006). Cognition through the lifespan: mechanisms of change. *Trends in Cognitive Sciences*, 10(3), 131-138.

- Crowe, S. F. (1998). The differential contribution of mental tracking, cognitive flexibility, visual search, and motor speed to performance on parts A and B of the Trail Making Test. *Journal of Clinical Psychology*, 54(5), 585-591.
- Daly, M., Harmon, C. P., & Delaney, L. (2009). Psychological and biological foundations of time preference. *Journal of the European Economic Association*, 7(2-3), 659-669.
- Das, S. R., Avants, B. B., Grossman, M., & Gee, J. C. (2009). Registration based cortical thickness measurement. *NeuroImage*, 45(3), 867-879.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8(12), 1704-1711.
- Daw, N. D., & Tobler, P. N. (2013). Value learning through reinforcement: the basics of dopamine and reinforcement learning. In P. W. Glimcher & E. Fehr (Eds.), *Neuroeconomics: decision making and the brain*, 2nd edition (pp. 283-298). San Diego, CA: Academic Press.
- Deary, I. J., Penke, L., & Johnson, W. (2010). The neuroscience of human intelligence differences. *Nature Reviews Neuroscience*, 11(3), 201-211.
- DeYoung, C. G., & Gray, J. R. (2009). Personality neuroscience: explaining individual differences in affect, behavior, and cognition. In P. J. Corr & G. Matthews (Eds.), *The Cambridge Handbook of Personality Psychology* (pp. 323-346). New York, NY: Cambridge University Press.
- DeYoung, C. G., Hirsh, J. B., Shane, M. S., Papademetris, X., Rajeevan, N., & Gray, J. R. (2010). Testing predictions from personality neuroscience brain structure and the big five. *Psychological Science*, 21(6), 820-828.

- Di Martino, A., Fair, D. A., Kelly, C., Satterthwaite, T. D., Castellanos, F. X., Thomason, M. E., ... & Milham, M. P. (2014). Unraveling the miswired connectome: a developmental perspective. *Neuron*, 83(6), 1335-1353.
- Dohmen, T., Falk, A., Huffman, D., & Sunde, U. (2010). Are risk aversion and impatience related to cognitive ability?. *The American Economic Review*, 100(3), 1238-1260.
- Dombrovski, A. Y., Siegle, G. J., Szanto, K., Clark, L., Reynolds, C. F., & Aizenstein, H. (2012). The temptation of suicide: striatal gray matter, discounting of delayed rewards, and suicide attempts in late-life depression. *Psychological Medicine*, 42(06), 1203-1215.
- Eaton, D. K., Kann, L., Kinchen, S., Shanklin, S., Flint, K. H., Hawkins, J., ... & Wechsler, H. (2012). Youth risk behavior surveillance-United States, 2011. *Morbidity and Mortality Weekly Report. Surveillance Summaries*, 61(4), 1-162.
- Eichenbaum, H. (2004). Hippocampus: cognitive processes and neural representations that underlie declarative memory. *Neuron*, 44(1), 109-120.
- Eklund, A., Nichols, T. E., & Knutsson, H. (2016). Cluster failure: why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences*, 113(28), 7900-7905.
- Eppinger, B., Hämmerer, D., & Li, S. C. (2011). Neuromodulation of reward-based learning and decision making in human aging. *Annals of the New York Academy of Sciences*, 1235(1), 1-17.
- Eppinger, B., Kray, J., Mock, B., & Mecklinger, A. (2008). Better or worse than expected? Aging, learning, and the ERN. *Neuropsychologia*, 46(2), 521-539.
- Eppinger, B., Mock, B., & Kray, J. (2009). Developmental differences in learning and error processing: evidence from ERPs. *Psychophysiology*, 46(5), 1043-1053.

- Eppinger, B., Nystrom, L. E., & Cohen, J. D. (2012). Reduced sensitivity to immediate reward during decision-making in older than younger adults. *PloS One*, 7(5), e36953.
- Epstein, L. H., Richards, J. B., Saad, F. G., Paluch, R. A., Roemmich, J. N., & Lerman, C. (2003). Comparison between two measures of delay discounting in smokers. *Experimental and Clinical Psychopharmacology*, 11(2), 131-138.
- Ersner-Hersfield, H., Garton, M. T., Ballard, K., Samanez-Larkin, G. R., & Knutson, B. (2009a). Don't stop thinking about tomorrow: Individual differences in future self-continuity account for saving. *Judgment and Decision Making*, 4(4), 280-286.
- Ersner-Hersfield, H., Wimmer, G. E., & Knutson, B. (2009b). Saving for the future self: Neural measures of future self-continuity predict temporal discounting. *Social Cognitive and Affective Neuroscience*, 4(1), 85-92.
- Euston, D. R., Gruber, A. J., & McNaughton, B. L. (2012). The role of medial prefrontal cortex in memory and decision making. *Neuron*, 76(6), 1057-1070.
- Evenden, J. L. (1999). Varieties of impulsivity. *Psychopharmacology*, 146(4), 348-361.
- Fernie, G., Peeters, M., Gullo, M. J., Christiansen, P., Cole, J. C., Sumnall, H., & Field, M. (2013). Multiple behavioural impulsivity tasks predict prospective alcohol involvement in adolescents. *Addiction*, 108(11), 1916-1923.
- Field, M., Christiansen, P., Cole, J., & Goudie, A. (2007). Delay discounting and the alcohol Stroop in heavy drinking adolescents. *Addiction*, 102(4), 579-586.
- Fields, S. A., Sabet, M., & Reynolds, B. (2013). Dimensions of impulsive behavior in obese, overweight, and healthy-weight adolescents. *Appetite*, 70, 60-66.
- Finucane, M. L., Slovic, P., Hibbard, J. H., Peters, E., Mertz, C. K., & MacGregor, D. G. (2002). Aging and decision-making competence: an analysis of comprehension

and consistency skills in older versus younger adults considering health □ plan options. *Journal of Behavioral Decision Making*, 15(2), 141-164.

Fjell, A. M., McEvoy, L., Holland, D., Dale, A. M., Walhovd, K. B., & Alzheimer's Disease Neuroimaging Initiative. (2014). What is normal in normal aging? Effects of aging, amyloid and Alzheimer's disease on the cerebral cortex and the hippocampus. *Progress in Neurobiology*, 117, 20-40.

Folstein, M. F., Folstein, S. E., & McHugh, P. R. (1975). Mini-mental state. A practical method for grading the cognitive state of patients for the clinician. *Journal of Psychiatric Research*, 12(3), 189-198.

Gerber, A. S., Huber, G. A., Doherty, D., & Dowling, C. M. (2011). The big five personality traits in the political arena. *Annual Review of Political Science*, 14, 265-287.

Giedd, J. N., Blumenthal, J., Jeffries, N. O., Castellanos, F. X., Liu, H., Zijdenbos, A., ... & Rapoport, J. L. (1999). Brain development during childhood and adolescence: a longitudinal MRI study. *Nature Neuroscience*, 2(10), 861-863.

Gläscher, J., Rudrauf, D., Colom, R., Paul, L. K., Tranel, D., Damasio, H., & Adolphs, R. (2010). Distributed neural system for general intelligence revealed by lesion mapping. *Proceedings of the National Academy of Sciences*, 107(10), 4705-4709.

Glimcher, P. W. (2011). Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, 108(Supplement 3), 15647-15654.

Gottfredson, L. S. (1997). Why g matters: The complexity of everyday life. *Intelligence*, 24(1), 79-132.

Green, L., Fry, A. F., & Myerson, J. (1994). Discounting of delayed rewards: A life-span comparison. *Psychological Science*, 5(1), 33-36.

- Gur, R. C., Richard, J., Hughett, P., Calkins, M. E., Macy, L., Bilker, W. B., ... & Gur, R. E. (2010). A cognitive neuroscience-based computerized battery for efficient measurement of individual differences: standardization and initial construct validation. *Journal of Neuroscience Methods*, 187(2), 254-262.
- Gur, R. C., Richard, J., Calkins, M. E., Chiavacci, R., Hansen, J. A., Bilker, W. B., ... & Gur, R. E. (2012). Age group and sex differences in performance on a computerized neurocognitive battery in children age 8– 21. *Neuropsychology*, 26(2), 251-265.
- Haier, R. J., Jung, R. E., Yeo, R. A., Head, K., & Alkire, M. T. (2004). Structural brain variation and general intelligence. *Neuroimage*, 23(1), 425-433.
- Hämmerer, D., Li, S. C., Müller, V., & Lindenberger, U. (2011). Life span differences in electrophysiological correlates of monitoring gains and losses during probabilistic reinforcement learning. *Journal of Cognitive Neuroscience*, 23(3), 579-592.
- Healey, M. K., Crutchley, P., & Kahana, M. J. (2014). Individual differences in memory search and their relation to intelligence. *Journal of Experimental Psychology: General*, 143(4), 1553-1569.
- Heitz, R. P., Schrock, J. C., Payne, T. W., & Engle, R. W. (2008). Effects of incentive on working memory capacity: Behavioral and pupillometric data. *Psychophysiology*, 45(1), 119-129.
- Hirsh, J. B., Morisano, D., & Peterson, J. B. (2008). Delay discounting: Interactions between personality and cognitive ability. *Journal of Research in Personality*, 42(6), 1646-1650.
- Hoyle, R. H., Stephenson, M. T., Palmgreen, P., Lorch, E. P., & Donohew, R. L. (2002). Reliability and validity of a brief measure of sensation seeking. *Personality and Individual Differences*, 32(3), 401-414.

- John, O. P., Donahue, E. M., & Kentle, R. L. (1991). *The big five inventory: Versions 4a and 54*. Berkeley, CA: University of California, Berkeley, Institute of Personality and Social Research.
- Johnson, M. W., & Bickel, W. K. (2002). Within-subject comparison of real and hypothetical money rewards in delay discounting. *Journal of the Experimental Analysis of Behavior*, 77(2), 129-146.
- Kaasinen, V., Vilkmann, H., Hietala, J., Någren, K., Helenius, H., Olsson, H., ... & Rinne, J. O. (2000). Age-related dopamine D2/D3 receptor loss in extrastriatal regions of the human brain. *Neurobiology of Aging*, 21(5), 683-688.
- Kable, J. W. (2013). Chapter 10. Valuation, intertemporal choice, and self-control. In P. W. Glimcher & E. Fehr (Eds.), *Neuroeconomics: decision making and the brain*, 2nd edition (pp. 173-192). San Diego, CA: Academic Press.
- Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience*, 10(12), 1625-1633.
- Kable, J. W., & Levy, I. (2015). Neural markers of individual differences in decision-making. *Current Opinion in Behavioral Sciences*, 5, 100-107.
- Kahneman, D., & Beatty, J. (1966). Pupil diameter and load on memory. *Science*, 154(3756), 1583-1585.
- Kanai, R., & Rees, G. (2011). The structural basis of inter-individual differences in human behaviour and cognition. *Nature Reviews Neuroscience*, 12(4), 231-242.
- Kane, M. J., & Engle, R. W. (2002). The role of prefrontal cortex in working-memory capacity, executive attention, and general fluid intelligence: An individual-differences perspective. *Psychonomic Bulletin & Review*, 9(4), 637-671.
- Khurana, A., Romer, D., Betancourt, L. M., Brodsky, N. L., Giannetta, J. M., & Hurt, H. (2012). Early adolescent sexual debut: The mediating role of working memory ability, sensation seeking, and impulsivity. *Developmental Psychology*, 48(5), 1416-1428.

- Kim, B. K., & Zauberman, G. (2009). Perception of anticipatory time in temporal discounting. *Journal of Neuroscience, Psychology, and Economics*, 2(2), 91-101.
- Kirby, K. N., & Maraković, N. N. (1995). Modeling myopic decisions: Evidence for hyperbolic delay-discounting within subjects and amounts. *Organizational Behavior and Human decision processes*, 64(1), 22-30
- Kirby, K. N., & Petry, N. M. (2004). Heroin and cocaine abusers have higher discount rates for delayed rewards than alcoholics or non-drug-using controls. *Addiction*, 99(4), 461-471.
- Kirby, K. N., Petry, N. M., & Bickel, W. K. (1999). Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. *Journal of Experimental Psychology: General*, 128(1), 78-87.
- Kirby, K. N., Winston, G. C., & Santiesteban, M. (2005). Impatience and grades: Delay-discount rates correlate negatively with college GPA. *Learning and Individual Differences*, 15(3), 213-222.
- Klein, A., Ghosh, S. S., Avants, B., Yeo, B. T., Fischl, B., Ardekani, B., ... & Parsey, R. V. (2010). Evaluation of volume-based and surface-based brain image registration methods. *NeuroImage*, 51(1), 214-220.
- Lenroot, R. K., Gogtay, N., Greenstein, D. K., Wells, E. M., Wallace, G. L., Clasen, L. S., ... & Giedd, J. N. (2007). Sexual dimorphism of brain developmental trajectories during childhood and adolescence. *NeuroImage*, 36(4), 1065-1073.
- Lubinski, D. (2000). Scientific and social significance of assessing individual differences: "Sinking shafts at a few critical points". *Annual Review of Psychology*, 51(1), 405-444.
- MacPherson, S. E., Phillips, L. H., & Della Sala, S. (2002). Age, executive function and social decision making: A dorsolateral prefrontal theory of cognitive aging. *Psychology and Aging*, 17(4), 598-609.

- Mackey, A. P., Finn, A. S., Leonard, J. A., Jacoby-Senghor, D. S., West, M. R., Gabrieli, C. F., & Gabrieli, J. D. (2015). Neuroanatomical correlates of the income-achievement gap. *Psychological Science*, 26(6), 925-933.
- MacKillop, J., Amlung, M. T., Few, L. R., Ray, L. A., Sweet, L. H., & Munafò, M. R. (2011). Delayed reward discounting and addictive behavior: a meta-analysis. *Psychopharmacology*, 216(3), 305-321.
- Madden, G. J., Petry, N. M., Badger, G. J., & Bickel, W. K. (1997). Impulsive and self-control choices in opioid-dependent patients and non-drug-using control patients: Drug and monetary rewards. *Experimental and Clinical Psychopharmacology*, 5(3), 256-262.
- Marschner, A., Mell, T., Wartenburger, I., Villringer, A., Reischies, F. M., & Heekeren, H. R. (2005). Reward-based decision-making and aging. *Brain Research Bulletin*, 67(5), 382-390.
- Mata, R., Josef, A. K., Samanez-Larkin, G. R., & Hertwig, R. (2011). Age differences in risky choice: A meta-analysis. *Annals of the New York Academy of Sciences*, 1235(1), 18-29.
- Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analysis of behavior: Vol 5. The effects of delay and intervening events on reinforcement value*. (pp. 55-73). Hillsdale, NJ: Erlbaum.
- McAdams, K. K., & Donnellan, M. B. (2009). Facets of personality and drinking in first-year college students. *Personality and Individual Differences*, 46(2), 207-212.
- McArdle, J. J., Ferrer-Caja, E., Hamagami, F., & Woodcock, R. W. (2002). Comparative longitudinal structural analyses of the growth and decline of multiple intellectual abilities over the life span. *Developmental Psychology*, 38(1), 115-142.

- McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally dissociable influences on learning rate in a dynamic environment. *Neuron*, 84(4), 870-881.
- Meier, S., & Sprenger, C. D. (2012). Time discounting predicts creditworthiness. *Psychological Science*, 23(1), 56-58.
- Mischel, W., Shoda, Y., & Rodriguez, M. L. (1989). Delay of gratification in children. *Science*, 244(4907), 933-938.
- Mohr, P. N., Li, S. C., & Heekeren, H. R. (2010). Neuroeconomics and aging: neuromodulation of economic decision making in old age. *Neuroscience & Biobehavioral Reviews*, 34(5), 678-688.
- Moore, T. M., Reise, S. P., Gur, R. E., Hakonarson, H., & Gur, R. C. (2015). Psychometric properties of the Penn Computerized Neurocognitive Battery. *Neuropsychology*, 29(2), 235-246.
- Morris, J. C., Heyman, A., Mohs, R. C., Hughes, J. P., Van Belle, G., Fillenbaum, G. D. M. E., ... & CERAD investigators (1989). The consortium to establish a registry for Alzheimer's disease (CERAD): I. Clinical and neuropsychological assessment of Alzheimer's disease. *Neurology*, 39(9), 1159-1165.
- Morris, J. C., Weintraub, S., Chui, H. C., Cummings, J., DeCarli, C., Ferris, S., ... & Kukull, W. A. (2006). The Uniform Data Set (UDS): clinical and cognitive variables and descriptive data from Alzheimer Disease Centers. *Alzheimer Disease & Associated Disorders*, 20(4), 210-216.
- Nassar, M. R., Bruckner, R., Gold, J. I., Li, S. C., Heekeren, H. R., & Eppinger, B. (2016). Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nature Communications*, 7:11609.

- Nassar, M. R., Rumsey, K. M., Wilson, R. C., Parikh, K., Heasly, B., & Gold, J. I. (2012). Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience*, *15*(7), 1040-1046.
- Nassar, M. R., Wilson, R. C., Heasly, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *Journal of Neuroscience*, *30*(37), 12366-12378.
- Olson, E. A., Collins, P. F., Hooper, C. J., Muetzel, R., Lim, K. O., & Luciana, M. (2009). White matter integrity predicts delay discounting behavior in 9-to 23-year-olds: a diffusion tensor imaging study. *Journal of Cognitive Neuroscience*, *21*(7), 1406-1421.
- Olson, E. A., Hooper, C. J., Collins, P., & Luciana, M. (2007). Adolescents' performance on delay and probability discounting tasks: contributions of age, intelligence, executive functioning, and self-reported externalizing behavior. *Personality and Individual Differences*, *43*(7), 1886-1897.
- Ostaszewski, P. (1996). The relation between temperament and rate of temporal discounting. *European Journal of Personality*, *10*(3), 161-172.
- Otto, A. R., Raio, C. M., Chiang, A., Phelps, E. A., & Daw, N. D. (2013). Working-memory capacity protects model-based learning from stress. *Proceedings of the National Academy of Sciences*, *110*(52), 20941-20946.
- Otto, A. R., Skatova, A., Madlon-Kay, S., & Daw, N. D. (2014). Cognitive control predicts use of model-based reinforcement learning. *Journal of Cognitive Neuroscience*, *27*(2), 319-333.
- Partala, T., & Surakka, V. (2003). Pupil size variation as an indication of affective processing. *International Journal of Human-computer Studies*, *59*(1), 185-198.
- Pavlov, I. P. (1927) Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex. London: Oxford University Press.

- Peters, J., & Büchel, C. (2010). Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-mediotemporal interactions. *Neuron*, 66(1), 138-148.
- Peters, J., & Büchel, C. (2011). The neural mechanisms of inter-temporal decision-making: understanding variability. *Trends in Cognitive Sciences*, 15(5), 227-239.
- Petry, N. M. (2001a). Delay discounting of money and alcohol in actively using alcoholics, currently abstinent alcoholics, and controls. *Psychopharmacology*, 154(3), 243-250.
- Petry, N. M. (2001b). Pathological gamblers, with and without substance abuse disorders, discount delayed rewards at high rates. *Journal of Abnormal Psychology*, 110(3), 482-487.
- Polyn, S. M., Erlikhman, G., & Kahana, M. J. (2011). Semantic cuing and the scale insensitivity of recency and contiguity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(3), 766-775.
- Prehn, K., Heekeren, H. R., Blasek, K., Lapschies, K., Mews, I., & Van der Meer, E. (2008). Neuroticism influences pupillary responses during an emotional interference task. *International Journal of Psychophysiology*, 70(1), 40-49.
- Raz, N., Lindenberger, U., Rodrigue, K. M., Kennedy, K. M., Head, D., Williamson, A., ... & Acker, J. D. (2005). Regional brain changes in aging healthy adults: general trends, individual differences and modifiers. *Cerebral Cortex*, 15(11), 1676-1689.
- Read, D., & Read, N. L. (2004). Time discounting over the lifespan. *Organizational Behavior and Human Decision Processes*, 94(1), 22-32.
- Reimers, S., Maylor, E. A., Stewart, N., & Chater, N. (2009). Associations between a one-shot delay discounting measure and age, income, education and real-world impulsive behavior. *Personality and Individual Differences*, 47(8), 973-978.

- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current Research and Theory* (pp. 64-99). New York, NY: Appleton-Century-Crofts.
- Reynolds, B. (2004). Do high rates of cigarette consumption increase delay discounting?: A cross-sectional comparison of adolescent smokers and young-adult smokers and nonsmokers. *Behavioural Processes*, 67(3), 545-549.
- Reynolds, B. (2006). A review of delay-discounting research with humans: relations to drug use and gambling. *Behavioural Pharmacology*, 17(8), 651-667.
- Reynolds, B., Richards, J. B., Horn, K., & Karraker, K. (2004). Delay discounting and probability discounting as related to cigarette smoking status in adults. *Behavioural Processes*, 65(1), 35-42.
- Rhodes, R. E., & Smith, N. E. I. (2006). Personality correlates of physical activity: a review and meta-analysis. *British Journal of Sports Medicine*, 40(12), 958-965.
- Riemann, R., Angleitner, A., & Strelau, J. (1997). Genetic and environmental influences on personality: A study of twins reared together using the self- and peer report NEO-FFI scales. *Journal of Personality*, 65(3), 449-475.
- Roberts, B. W., Kuncel, N. R., Shiner, R., Caspi, A., & Goldberg, L. R. (2007). The power of personality: The comparative validity of personality traits, socioeconomic status, and cognitive ability for predicting important life outcomes. *Perspectives on Psychological Science*, 2(4), 313-345.
- Roberts, B. W., & Mroczek, D. (2008). Personality trait change in adulthood. *Current Directions in Psychological Science*, 17(1), 31-35.

- Samanez-Larkin, G. R., & Knutson, B. (2015). Decision making in the ageing brain: changes in affective and motivational circuits. *Nature Reviews Neuroscience*, 16(5), 278-289.
- Samanez-Larkin, G. R., Levens, S. M., Perry, L. M., Dougherty, R. F., & Knutson, B. (2012). Frontostriatal white matter integrity mediates adult age differences in probabilistic reward learning. *Journal of Neuroscience*, 32(15), 5333-5337.
- Samanez-Larkin, G. R., Mata, R., Radu, P. T., Ballard, I. C., Carstensen, L. L., & McClure, S. M. (2011). Age differences in striatal delay sensitivity during intertemporal choice in healthy adults. *Frontiers in Neuroscience*, 5, 126.
- Satterthwaite, T. D., Connolly, J. J., Ruparel, K., Calkins, M. E., Jackson, C., Elliott, M. A., ... & Gur, R. E. (2016). The Philadelphia Neurodevelopmental Cohort: a publicly available resource for the study of normal and abnormal brain development in youth. *NeuroImage*, 124, 1115-1119.
- Satterthwaite, T. D., Elliott, M. A., Ruparel, K., Loughhead, J., Prabhakaran, K., Calkins, M. E., ... & Gur, R. E. (2014a). Neuroimaging of the Philadelphia neurodevelopmental cohort. *NeuroImage*, 86, 544-553.
- Satterthwaite, T. D., Shinohara, R. T., Wolf, D. H., Hopson, R. D., Elliott, M. A., Vandekar, S. N., ... & Gur, R. E. (2014b). Impact of puberty on the evolution of cerebral perfusion during adolescence. *Proceedings of the National Academy of Sciences*, 111(23), 8643-8648.
- Schacter, D. L., Addis, D. R., & Buckner, R. L. (2008). Episodic simulation of future events. *Annals of the New York Academy of Sciences*, 1124(1), 39-60.
- Scheres, A., Dijkstra, M., Ainslie, E., Balkan, J., Reynolds, B., Sonuga-Barke, E., & Castellanos, F. X. (2006). Temporal and probabilistic discounting of rewards in children and adolescents: effects of age and ADHD symptoms. *Neuropsychologia*, 44(11), 2092-2103.
- Schönberg, T., Daw, N. D., Joel, D., & O'Doherty, J. P. (2007). Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making. *Journal of Neuroscience*, 27(47), 12860-12867.

- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593-1599.
- Schwartz, D. L., Mitchell, A. D., Lahna, D. L., Luber, H. S., Huckans, M. S., Mitchell, S. H., & Hoffman, W. F. (2010). Global and local morphometric differences in recently abstinent methamphetamine-dependent individuals. *NeuroImage*, 50(4), 1392-1401.
- Seeley, W. W., Menon, V., Schatzberg, A. F., Keller, J., Glover, G. H., Kenna, H., ... & Greicius, M. D. (2007). Dissociable intrinsic connectivity networks for salience processing and executive control. *Journal of Neuroscience*, 27(9), 2349-2356.
- Sellitto, M., Ciaramelli, E., & di Pellegrino, G. (2010). Myopic discounting of future rewards after medial orbitofrontal damage in humans. *Journal of Neuroscience*, 30(49), 16429-16436.
- Senecal, N., Wang, T., Thompson, E., & Kable, J. W. (2012). Normative arguments from experts and peers reduce delay discounting. *Judgment and Decision Making*, 7(5), 568-589.
- Shadlen, M. N., & Shohamy, D. (2016). Decision making and sequential sampling from memory. *Neuron*, 90(5), 927-939.
- Shamosh, N. A., & Gray, J. R. (2008). Delay discounting and intelligence: A meta-analysis. *Intelligence*, 36(4), 289-305.
- Shamosh, N. A., DeYoung, C. G., Green, A. E., Reis, D. L., Johnson, M. R., Conway, A. R., ... & Gray, J. R. (2008). Individual differences in delay discounting: relation to intelligence, working memory, and anterior prefrontal cortex. *Psychological Science*, 19(9), 904-911.
- Sobel, M. E. (1982). Asymptotic confidence intervals for indirect effects in structural equation models. *Sociological methodology*, 13, 290-312.
- Sotiras, A., Resnick, S. M., & Davatzikos, C. (2015). Finding imaging patterns of structural covariance via non-negative matrix factorization. *NeuroImage*, 108, 1-16.

- Sotiras, A., Toledo, J. B., Gur, R. E., Gur, R. C., Satterthwaite, T. D., & Davatzikos, C. (2017). Patterns of coordinated cortical remodeling during adolescence and their associations with functional specialization and evolutionary expansion. *Proceedings of the National Academy of Sciences*, 114(13), 3527-3532.
- Soutschek, A., Ruff, C. C., Strombach, T., Kalenscher, T., & Tobler, P. N. (2016). Brain stimulation reveals crucial role of overcoming self-centeredness in self-control. *Science Advances*, 2(10), e1600992.
- Sowell, E. R., Thompson, P. M., Leonard, C. M., Welcome, S. E., Kan, E., & Toga, A. W. (2004). Longitudinal mapping of cortical thickness and brain growth in normal children. *Journal of Neuroscience*, 24(38), 8223-8231.
- Specht, J., Egloff, B., & Schmukle, S. C. (2011). Stability and change of personality across the life course: the impact of age and major life events on mean-level and rank-order stability of the Big Five. *Journal of Personality and Social Psychology*, 101(4), 862-882.
- Spielberger, C. D., Gorsuch, R. L., Lushene, R., Vagg, P. R., & Jacobs, G. A. (1983). *Manual for the State-Trait Anxiety Inventory*. Palo Alto, CA: Consulting Psychologists Press.
- Stanger, C., Elton, A., Ryan, S. R., James, G. A., Budney, A. J., & Kilts, C. D. (2013). Neuroeconomics and adolescent substance abuse: individual differences in neural networks and delay discounting. *Journal of the American Academy of Child & Adolescent Psychiatry*, 52(7), 747-755.
- Steinberg, L., Graham, S., O'Brien, L., Woolard, J., Cauffman, E., & Banich, M. (2009). Age differences in future orientation and delay discounting. *Child Development*, 80(1), 28-44.
- Strombach, T., Weber, B., Hangebrauk, Z., Kenning, P., Karipidis, I. I., Tobler, P. N., & Kalenscher, T. (2015). Social discounting involves modulation of neural value signals by temporoparietal junction. *Proceedings of the National Academy of Sciences*, 112(5), 1619-1624.

- Thomason, M. E., Dassanayake, M. T., Shen, S., Katkuri, Y., Alexis, M., Anderson, A. L., ... & Romero, R. (2013). Cross-hemispheric functional connectivity in the human fetal brain. *Science Translational Medicine*, 5(173), 173ra24.
- Thornton, W. J., & Dumke, H. A. (2005). Age differences in everyday problem-solving and decision-making effectiveness: a meta-analytic review. *Psychology and Aging*, 20(1), 85–99.
- Tjur, T. (2009). Coefficients of determination in logistic regression models—A new proposal: The coefficient of discrimination. *The American Statistician*, 63(4), 366-372.
- Tustison, N. J., Avants, B. B., Cook, P. A., Zheng, Y., Egan, A., Yushkevich, P. A., & Gee, J. C. (2010). N4ITK: improved N3 bias correction. *IEEE Transactions on Medical Imaging*, 29(6), 1310-1320.
- Tustison, N. J., Cook, P. A., Klein, A., Song, G., Das, S. R., Duda, J. T., ... & Avants, B. B. (2014). Large-scale evaluation of ANTs and FreeSurfer cortical thickness measurements. *NeuroImage*, 99, 166-179.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297-323.
- Tymula, A., Belmaker, L. A. R., Ruderman, L., Glimcher, P. W., & Levy, I. (2013). Like cognitive function, decision making across the life span shows profound age-related changes. *Proceedings of the National Academy of Sciences*, 110(42), 17143-17148.
- Vandekar, S. N., Shinohara, R. T., Raznahan, A., Roalf, D. R., Ross, M., DeLeo, N., ... & Satterthwaite, T. D. (2015). Topologically dissociable patterns of development of the human cerebral cortex. *Journal of Neuroscience*, 35(2), 599-609
- Van Der Meer, E., Beyer, R., Horn, J., Foth, M., Bornemann, B., Ries, J., ... & Wartenburger, I. (2010). Resource allocation and fluid intelligence: Insights from pupillometry. *Psychophysiology*, 47(1), 158-169.
- Van Essen, D. C. (2005). A population-average, landmark-and surface-based (PALS) atlas of human cerebral cortex. *NeuroImage*, 28(3), 635-662.

- Van Essen, D. C., Drury, H. A., Dickson, J., Harwell, J., Hanlon, D., & Anderson, C. H. (2001). An integrated software suite for surface-based analyses of cerebral cortex. *Journal of the American Medical Informatics Association*, 8(5), 443-459.
- Van Leijenhorst, L., Moor, B. G., de Macks, Z. A. O., Rombouts, S. A., Westenberg, P. M., & Crone, E. A. (2010). Adolescent risky decision-making: neurocognitive development of reward and control regions. *NeuroImage*, 51(1), 345-355.
- Vuchinich, R. E., & Simpson, C. A. (1998). Hyperbolic temporal discounting in social drinkers and problem drinkers. *Experimental and Clinical Psychopharmacology*, 6(3), 292-305.
- Wang, L., Zang, Y., He, Y., Liang, M., Zhang, X., Tian, L., ... & Li, K. (2006). Changes in hippocampal connectivity in the early stages of Alzheimer's disease: evidence from resting state fMRI. *Neuroimage*, 31(2), 496-504.
- Wang, Q., Chen, C., Cai, Y., Li, S., Zhao, X., Zheng, L., ... & Xue, G. (2016). Dissociated neural substrates underlying impulsive choice and impulsive action. *NeuroImage*, 134, 540-549.
- Weber, E. U., Blais, A. R., & Betz, N. E. (2002). A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making*, 15(4), 263-290.
- Wechsler, D. (1997). *WAIS-III Administration and Scoring Manual*. San Antonio, TX: The Psychological Corporation.
- Wechsler, D. (1999). *WASI manual*. San Antonio, TX: Psychological Corporation.
- Weilbacher, R. A., & Gluth, S. (2016). The interplay of hippocampus and ventromedial prefrontal cortex in memory-based decision making. *Brain Sciences*, 7(1), 4.
- Wilson, R. C., Nassar, M. R., & Gold, J. I. (2010). Bayesian online learning of the hazard rate in change-point problems. *Neural Computation*, 22(9), 2452-2476.
- Wood, S. N. (2004). Stable and efficient multiple smoothing parameter estimation for generalized additive models. *Journal of the American Statistical Association*, 99(467), 673-686.
- Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(1), 3-36.

- Yi, R., Gatchalian, K. M., & Bickel, W. K. (2006). Discounting of past outcomes. *Experimental and Clinical Psychopharmacology*, 14(3), 311-317.
- Yi, R., Landes, R. D., & Bickel, W. K. (2009). Novel models of intertemporal valuation: Past and future outcomes. *Journal of Neuroscience, Psychology, and Economics*, 2(2), 102-111.
- Zaromb, F. M., Howard, M. W., Dolan, E. D., Sirotin, Y. B., Tully, M., Wingfield, A., & Kahana, M. J. (2006). Temporal associations and prior-list intrusions in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(4), 792-804.
- Zielinski, B. A., Gennatas, E. D., Zhou, J., & Seeley, W. W. (2010). Network-level structural covariance in the developing brain. *Proceedings of the National Academy of Sciences*, 107(42), 18191-18196.