The neurocomputational basis of self-control success and failure

Thesis by
Nicolette Jean Sullivan

In Partial Fulfillment of the Requirements for the degree
of
Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California
2015
(Defended May 5, 2015)

ACKNOWLEDGEMENTS

I'm indebted to a great many individuals for the position I am in today. Most of all, I am grateful for the unconditional love and support of my late mother, Bonnie Dunn Sullivan, and my father, Lucky Sullivan. They set very high standards for me, and although I didn't always appreciate that, or believe in my ability to meet those standards, this dissertation is physical proof of their persistence. My great aunts Marie Greider and the late Phyllis Hartroft were academics when that was difficult for women in science, and I am eternally grateful for their unique advice and emotional support. My cousins Katrina Fitzpatrick, Kathy and Kent Morris, and George Greider and Gayle Kranz were consistently encouraging. My aunts Linda Dunn and Connie Lee have always made my trips back to our family farm in Dunnesville much-needed opportunities to recharge my mental batteries.

The pioneering Kevin McCabe at George Mason University gave me my first research position in Neuroeconomics, and I will be eternally grateful for that opportunity. My mentors at the University of Chicago – Howard Nusbaum, John A. List, and Ali Hortacsu – provided me with unending support and tolerance, and are stellar examples of what it means to be a high-impact scholar. At Caltech, my advisor and mentor Antonio Rangel has been a touchstone for academic rigor. His uncanny scientific instinct and dedication set a high bar that I will continue to strive toward. I am a grateful benefactor of his fierce support and passion for his work. My dissertation committee – Colin Camerer, Ralph Adolphs, Steven Quartz, and John O'Doherty – have given me invaluable guidance, and the projects herein would not be what they are today without their input. I also thank my co-authors Adriana Galvan and Mara Mather for their availability and expert advice throughout the long slog that is lifespan research. Thanks also to Ella Young for her help with data collection.

I'm grateful for the advice and kindness of current and former fellow lab members, CNS graduate students and staff, and other victims I tricked into providing feedback on this manuscript. In no particular order, thank you to Amanda Brand, Leila Montaser, Jenny Neise, Peter "Bananas" Sokol-Hessner, John Clithero, Alison Harris, Todd Hare, Alice Liu (née Lin), Gabriela Tavares, Mike Tyszka, Vanessa Janowski, Ian Krajbich, Cendri Hutcherson, Geoff "Colorblind" Fisher, Damian Stanley, Laura Harrison, Tim Armstrong, Yazan Billeh, Stephanie Echols, Shuo Wang, and Dan McNamee. Also, thanks to Debbie Ma and Erin Dean for sending me delicious take-out food during the delirious final throws of thesis writing. This dissertation would also not have been possible without the generous support of the NSF, NIA, Institute for Humane Studies, and Kavli Foundation.

ABSTRACT

Individuals often have difficulty delaying gratification – that is, forging smaller sooner rewards in favor of larger rewards delivered at a delay. Common examples of this deficit in self-control are difficulties in saving for retirement, going to the gym, or eating healthy foods. Despite an extensive literature on the neural substrates of decision-making, relatively little is still known about the sources of underlying individual variation in the ability to successfully execute self-control. This manuscript presents three studies examining the sources of individual differences in delay of gratification, with the goal of obtaining a more complete understanding of the neural mechanisms underlying choice. The main question this thesis addresses is: *what features of the brain's decision process allow individuals to down-regulate the appeal of smaller sooner rewards, in order to forgo them in favor of greater future reward?*

In the first study, I present a novel method of measuring decision process dynamics, in which we harness the power of fine temporal resolution in recording computer mouse movements in dietary choices. We find that up to 39% of individual variation in dietary self-control can be explained by differences in the speed with which the decision-making circuitry processes basic attributes, such as tastiness, versus more complex, abstract, attributes, such as healthfulness.

In the second study, we extend this novel approach to a classic experimental economics paradigm, intertemporal choice. We found large individual variance in the speeds with which immediate and delayed reward values were processed. We found that about 25% of the individual differences can be explained by differences in the speed at which delayed rewards are processed. We also found that the relative speed at which immediate and delayed rewards are processed fluctuated across trials: subjects processed delayed rewards faster than immediate rewards when they made patient choices, but the order of processing speeds was reversed during impulsive choices. Together, these results show that a sizable fraction of variation in the ability to postpone gratification might be attributable to variables that affect the speed at which different types of rewards are processed, and not to differences on deep preference parameters like the temporal discount rate used by the brain's valuation systems.

Across the lifespan, self-control improves in many choice domains. The third and final study capitalizes on this phenomenon of behavioral change with age to investigate the neural mechanisms underlying improvements in self-control. I use functional magnetic resonance imaging (fMRI) to

examine the neural correlates of changes in discounting of future monetary rewards across the lifespan from adolescents (13 years old) to seniors (70 years old). We find that neural response to value in reward-related striatal brain regions dramatically decrease with age. In contrast, we find that the left dorsolateral prefrontal cortex, often found to be related to successful self-control, increases its functional connectivity to key valuation, reward, and future-thinking brain regions with age during very tempting trials. These results suggest a mechanism through which increased self-control is improved.

Taken together, these studies argue that individual features of the decision process have a large influence on the overall ability to exert self-control in both dietary and monetary choice domains. Specifically, we find that the speed with which abstract future attributes such as health information, relative to more concrete attributes such as taste, are processed have a large influence on individual self-control ability. We also find that decreased reward sensitivity, paired with increased effective connectivity between control and valuation regions specifically when control is required most, allow for increased ability to delay gratification with age.

TABLE OF CONTENTS

C h a p t e r 1

INTRODUCTION

1.1 Self-control in decision-making

Self-control deficits have long-term consequences for both the individual and for society. For example, dietary self-control choices present a trade-off between an immediate, pleasurable reward – like eating a delicious cookie – and a greater reward later, like forgoing the treat and having better health in the future. A large percentage of Americans struggle with these choices each year, and spend sizable amounts of money doing so; the diet business is currently a sixty-one billion dollar industry (Marketdata Enterprises Inc., 2014; Paeratakul, York-Crowe, Williamson, Ryan, & Bray, 2002) . Despite this, approximately 34% of U.S. adults are still overweight (Flegal, Carroll, Ogden, & Curtin, 2010). This not only has serious consequences for the individual, but also for society; it's estimated that obesity produces a 33% increase in cost to the health care system (Andreyeva, Sturm, & Ringel, 2004). Dietary self-control is not the only area in which individuals have difficulty forgoing tempting immediate rewards. The "revolving door" nature of drug, alcohol, and gambling rehabilitation facilities (Carey, 2008; Hunt, Barnett, & Branch, 1971) suggests that across many choice domains, individuals often fail at self-control despite costly efforts to succeed.

What accounts for this disparity between intention and outcome? I argue that much of the difference has to do with deficits in the ability to successfully exert self-control in everyday decisions – or, as philosophers term it, the uniquely human conflict between first and second order preferences (Brandt, 1979, 1998; Frankfurt, 1971). Philosophers call a lower or first order preference one that is a preference of the moment, such as the preference to satisfy a craving for fatty foods. A higher or second order preference is typically long-term focused, concerning the kind of person one desires to be – for example, a fit, healthy person who eats well and exercises. A contradiction between one's first and second order preferences presents conflict; making a choice in favor of one's second-order preference requires bringing one's choices more in line with long-term goals instead of fleeting first-order preferences. In this thesis I will term this the process of *self-control exertion*.

Decision-making requires the brain to calculate the subjective value[1] each option holds to the decision maker, so an individual can decide between those options. Successful self-control further requires the brain to make those values reflective not only of short-term rewards, but also of future long-term consequences. The main question of this thesis is: *What features of the decision process allow individuals to down-regulate the value of tempting first order options by incorporating information about their second-order preferences, in order to forgo them in favor of greater future reward?*

1.2 The neural computations of value processing and comparison

To make a motivated choice between different options, the brain must first compute the value of each option. A great deal of experimental work has been done to pinpoint the neural computations involved in valuation, and has converged on a network of brain regions that appear to be involved in processing value across many different types of tasks and stimuli. For example, the neural computations of value-based decision-making can be studied using a classic experimental economics paradigm, monetary intertemporal choice (e.g., Ainslie, 1975; Berns, Laibson, & Loewenstein, 2007; Green, Fristoe, & Myerson, 1994; Laibson, 1997; Loewenstein & Prelec, 1992). In these experiments, subjects are asked to decide between a small immediate amount of money, and a larger, but delayed, amount. This kind of financial choice is conceptually similar to the dietary choice trade-off discussed previously, and is another domain in which there are large, widespread deficits in self-control. For example, 31% percent of American adults do not have any retirement savings (Federal Reserve Board, 2014), preferring instead to spend all of their money maintaining their current lifestyle. The intertemporal choice task has another distinct advantage, that of external validity for predicting well-being. Low self-control in this task among aging populations predicts decreased well-being and higher mortality (Boyle, Yu, Gamble, & Bennett, 2013; Forstmeier, Drobetz, & Maercker, 2011).

An early functional magnetic resonance imaging (fMRI) study of monetary intertemporal choice found that activity in a specific set of brain regions, including the ventromedial prefrontal cortex (vmPFC), posterior cingulate cortex (PCC), and ventral striatum, correlates positively with value, (Kable & Glimcher, 2007). Since then, this network has been shown to increase in activity with

[1] Here, we define value as the importance or worth of an item – specifically, it is a quantifiable attribute assigned to an item that has high explanatory power for choice behavior. For example, taste is a value parameter in dietary choice.

increasing value for a wide range of goods such as foods, trinkets, and money (Chib, Rangel, Shimojo, & O'Doherty, 2009; Clithero & Rangel, 2014). There is even evidence that this network may convert such disparate stimuli types into a "common currency" of value for comparison before a choice can be made (McNamee, Rangel, & O'Doherty, 2013).

The literature suggests a reliable network that represents the value of choices across various studies and types of stimuli, but how are values constructed in the first place, and used to select one option? Decision making is a dynamic process that occurs over time, requiring the computation and comparison of various pieces of value information ("attributes"). Many computational models have attempted to describe the brain's value construction and comparison process, and prominent among them are accumulator models of choice such as Decision Field Theory (Busemeyer & Townsend, 1993), the Drift-Diffusion Model (Ratcliff & McKoon, 2008), and the Leaky Competing Accumulator Model (Usher & McClelland, 2001). In all of these models, decisions are made by dynamically integrating information for and against each option in a choice into a value signal. A decision is made when this relative value signal becomes sufficiently large or sufficiently small. This thesis will use one of these models in particular, the Drift Diffusion Model (DDM), to conceptualize this process.

To understand this process, imagine that an individual is deciding whether to eat a chocolate chip cookie or a less delicious snack like celery (Figure 1). At the beginning of the decision process, the relative value signal (RVS) has a value of zero. Over time, the RVS changes by accumulating information in the following manner:

$$RVS_t = RVS_{t-1} + \delta(V_1 - V_2) + \varepsilon_t.$$

The brain takes some amount of time (t) to calculate the value (V) each option holds for them; this is referred to as the non-decision time (NDT). Before this information is calculated, the RVS simply reflects any noise that exists in the decision circuitry (ε_t). After these values are calculated, the RVS begins accumulating a noisy weighted sum of the evidence in favor of each option according the equation above. Value information adds a certain amount, determined by the drift rate parameter (δ). A choice is made once the RVS becomes sufficiently strong in favor of one option, mathematically represented by the RVS crossing the boundary for that option.

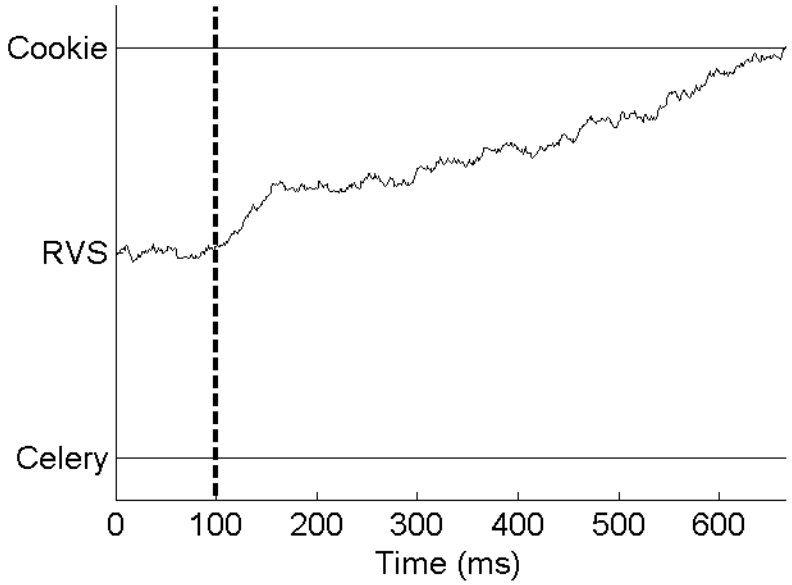

Figure 1. Illustration of a sequential integrator model of choice. In this example, an individual decides between eating a highly-valued cookie and a lower-valued bunch of celery. The dotted line indicates the time the brain takes to process the value of each option (100 ms). A choice is made by computing a relative value signal (RVS) that sequentially integrates a weighted linear sum of the value (once it is processed, at 100 ms). Random noise is also added to the RVS at each time point to represent noise in the decision circuitry. When the RVS crosses one of the two barriers (here, at approximately 680 ms), a choice is made.

Several questions arising from this model have not been addressed in the literature, and are examined in Chapters 2 and 3 of this thesis. Firstly, although previous research uses only one value, in reality there are vastly different types of information one must use to make most real-world choices. In even a simple food choice like the one introduced above between a cookie and some celery, there are various attributes one must compute before constructing each option's total value. For example, the taste of the foods is a distinctly different type of value information than their healthfulness. Individuals tend to place more weight on taste than health in their food choices, and taste, relative to health, has been found to be more strongly reflected in the vmPFC in individuals who are not successful at maintaining a diet to lose weight (T. Hare, C. F. Camerer, & A. Rangel, 2009). Moreover, these attributes are very different conceptually; taste is a more visceral, innate attribute, whereas health is more a complex, abstract attribute. More abstract values like healthfulness may be more difficult to process, and therefore take longer to compute by the brain's

valuation circuitry (Liberman & Trope, 2008). If taste (*T*) and health (*H*) information are truly distinctly different types of information, the relative value signal may evolve over time according to the following equation:

$$RVS_t = RVS_{t-1} + \delta_T(T_{Left}\text{-}T_{Right}) + \delta_H(H_{Left}\text{-}H_{Right}) + \varepsilon_t.$$

This equation implies that taste and health may contribute different amounts to the RVS at each time point. Additionally, the brain may require different amounts of time to calculate each type of information (i.e., have different NDTs). In the example depicted in Figure 2, an individual is again deciding between a chocolate chip cookie and celery. However, taste information is processed quickly, causing the RVS to initially trend toward selecting the tastier option. Health information is processed more slowly. In this example, health information is processed early enough and accumulates enough evidence to sufficiently bias the RVS in favor of the healthier option. However, one could easily imagine a case in which health information enters the value comparison process so late that it does not have an opportunity to influence this signal, leading to a self-control failure regardless of the true preference on how heavily to weight health in choice. In this example, even a milliseconds-later processing of health would have led to selection of the cookie instead.

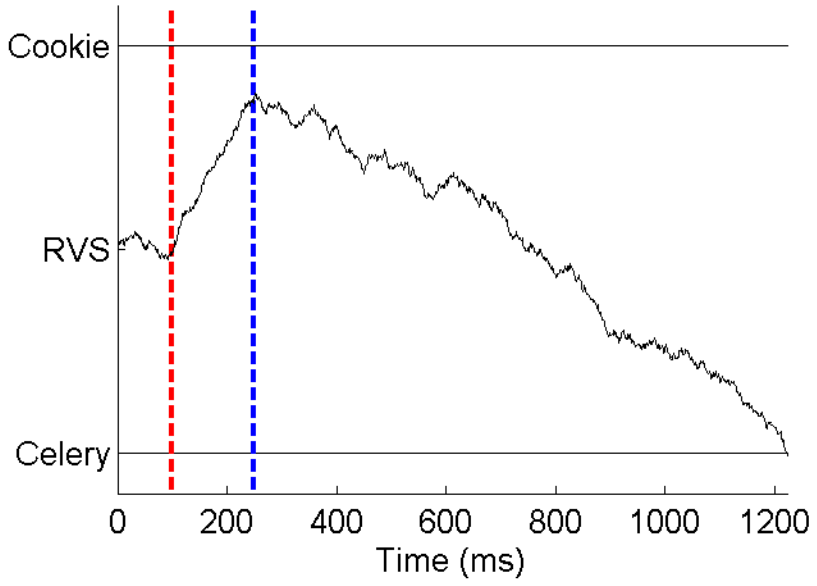

Figure 2. Illustration of how attribute value processing speed could influence choice. In this example, an individual decides between eating a tasty but very unhealthy cookie, and a very

healthy but not tasty bunch of celery. Suppose that it takes 100 ms to compute taste information (indicated by the red dotted line), whereas health information takes 250 ms to process (indicated by the blue dotted line). A choice is made by computing a relative value signal (RVS), which begins with a value of zero, and then sequentially integrates a weighted linear sum of the attributes represented at this time plus noise. A choice is made the first time the RVS crosses one of the two barriers.

Theory also suggests that, even if processing speeds for different choice features are the same, fast overall processing speeds may influence choice in favor of value types that are more weakly represented in the decision process (in this example, health information may be more weakly represented due to a general disregard for eating healthily). For example, imagine that taste and health actually are processed at the same time. However, due to a general deficit in preference for health, health value information contributes less to the RVS at every time point (that is, in the equation above, $\delta_H < \delta_T$). It is very possible that the longer that health has to influence the RVS, the more the RVS will reflect health's value. Additional elements could also be included in this model, such as inhibition or down-weighting of taste once health is processed.

These concepts are critical for building a more representative computational model of the decision process. Chapters 2 and 3 of this thesis will address these questions: *do some types of attribute values take longer to process? Do earlier-processed attribute values have a larger weight in choice? Do faster overall processing speeds lead to a better reflection of lower-weighted, longer-term goal focused, attribute values?* Next, I will discuss the brain circuitry involved in incorporating long-term goals, such as healthiness, into value calculations, allowing an individual to execute self-control.

1.3 Control over impulses

Bringing one's decision values more in line with one's long-term, likely more abstract, goals, instead of those offering only immediate gratification, is a difficult computational problem. A great deal of work has been done in the neuroscience literature to understand the neural computations underlying this process, and suggests that the regions discussed above involved in valuation may work in tandem with regions associated with cognitive control to allow an individual to down-regulate the value of tempting options. These studies suggested that one brain region in particular, the left dorsolateral prefrontal cortex (L DLPFC), is a crucial component of this process. For example, this region appears to be functionally connected to the vmPFC valuation region, and may

be responsible for down-weighting the value of tempting options, allowing more patient choices (T. Hare et al., 2009). Greater DLPFC cortical surface area and thickness have been associated with less temporal discounting – that is, less selection of the shorter, smaller monetary reward – in an intertemporal choice task (Drobetz et al., 2014). Further, inhibition of this region using transcranial magnetic stimulation leads to more impulsivity in a monetary intertemporal choice task (Figner et al., 2010). Together with the brain's valuation network discussed above, there is strong evidence to suggest there is a concerted system for decision-making when top-down regulation is required to bring one's actions in line with more long-term, abstract, goals.

To pinpoint the contribution of this control region on overall choice behavior, it is useful to measure simultaneous changes in this region's structure, function, and connectivity that are associated with changes in self-control. Such changes have, in fact, been seen across one dimension: age. Self-control increases with age across a number of choice domains, from experimental measures such as temporal discounting to alcoholism and gambling (Green et al., 1994; National Institute on Alcohol Abuse and Alcoholism, 2008; Read & Read, 2004; Sangrock, 2002; Spear, 2009). Additionally, the structure and function of the regions related to valuation and control discussed above change systematically across the lifespan (Eppinger, Nystrom, & Cohen, 2012; Eppinger, Schuck, Nystrom, & Cohen, 2013; Hare et al., 2008; Samanez-Larkin, Levens, Perry, Dougherty, & Knutson, 2012). For example, research suggests that children and adolescents have a much stronger response to reward in the ventral striatum than do adults (Galvan et al., 2006), and that older subjects have significantly less response to the valence of stimuli than adults (Mather et al., 2004). Moreover, lateral brain regions associated with control, such as the DLPFC, show marked change in their structure with age relative to other regions (Sowell et al., 2003).

This research suggests systematic improvements in the brain's self-regulation network across the lifespan that could lead to a better assignment of value to tempting options, resulting in better choices. Understanding how changes in the structure and function of the brain's decision-making circuitry allow for better reflection of second-order preferences will help us pinpoint the brain circuitry important for these choices in general. This brings us to the final question of this thesis: a*re there neural markers for improvements in self-control?* The study presented in Chapter 4 will address this question.

1.4 In this thesis

To address all of the questions raised above requires methods that can measure the neural processes involved in executing self-control. It also requires a way to measure the value comparison process at a fine temporal resolution. To do so, I will use converging evidence from several methodologies. These methodologies are mentioned below, and are described in more detail in the relevant chapters.

In Chapter 2, I address the question of how attribute value processing speeds can influence self-control behavior. As illustrated in the example in Figure 2, very small, millisecond-level timing differences may be all that stands between choosing a cookie and choosing the healthier option. To observe the decision process at such a fine temporal resolution is a challenge. In this chapter, I introduce a new method that harnesses the power of fine temporal resolution available when recording computer mouse movements to estimate the decision process in dietary self-control. I propose that variation across individuals in self-control abilities is partly due to differences in the speed with which the decision-making circuitry processes basic attributes, such as tastiness, versus more abstract attributes, such as healthfulness.

In Chapter 3, I apply this new method to the classic experimental economics choice paradigm of monetary intertemporal choice. Using the mouse tracking approach introduced in Chapter 2, I investigate how differences in value processing speeds influence ability to exert self-control in the financial domain. I extend this technique to understand how small differences in decision value processing speeds during a trial can lead a generally-patient subject to fail at self-control, and conversely allow an impatient subject to succeed.

Chapter 4 presents a lifespan study to understand the neural underpinnings of successful self-control. In this experiment, I harness the power of fMRI to examine the structure and function of neural systems associated with reward, valuation, future thinking, and control. I examine how these systems change through the lifespan to bring valuations more in line with long-term rewards, rather than short-term gains. In doing so, I also propose that there are not only age-dependent changes, but also age-invariant neural correlates of self-control.

Chapter 4 summarizes the research detailed in the previous chapters. In addition, I discuss implications of my findings on our understanding of the decision process, including the origins of self-control success and failure, and several open questions for the field of decision neuroscience that arise from this work.

Chapter 2

DIETARY SELF-CONTROL IS RELATED TO THE SPEED WITH WHICH HEALTH AND TASTE ATTRIBUTES ARE PROCESSED[2]

2.1 Introduction

A large fraction of the population struggles with dietary self-control. For them, a choice between an apple and a chocolate chip cookie is more difficult than a choice between an apple and an orange, and is also more likely to result in decision mistakes (Rangel, 2013). This leads to two basic questions: why is sustained dietary self-control so difficult for so many, and why do some individuals have better self-control than others?

These important questions have received considerable attention in several disciplines, including economics, psychology, and neuroscience. Behavioral economists conceptualize self-control failures as instances in which individuals over-discount the future consequences of their decisions (Ikeda, Kang, & Ohtake, 2010; Laibson, 1997; O'Donoghue & Rabin, 1999). Related to this idea, psychologists have proposed that dietary self-control is hard due to inherent limitations in how we process two different types of attributes (Liberman & Trope, 2008). In this view, basic attributes like the anticipated taste of foods are easily represented at the time of choice, and reliably weighted in decisions. In contrast, taking into account more abstract attributes like the healthiness of foods is thought to require effort, and as a result they are not weighted as reliably in decisions. This leads to a relative overweighting of taste compared to health information, which can result in poor dietary choices. Neuroscientists have identified neural mechanisms underlying the differences between how taste and health attributes are computed at the time of decision in people with good and poor self-control (T. Hare et al., 2009). In both groups, areas of ventromedial prefrontal cortex (vmPFC) represent the overall value assigned to the food. However, in individuals with poor self-control this area responds only to the taste of foods, whereas in individuals with good self-control it responds to both health and taste information. Furthermore, in those with good self-control, areas of the left dorsolateral prefrontal cortex seem to modulate value-related activity in vmPFC.

---

[2] Adapted with permission from Sullivan et al., 2015.

Here we propose that dietary self-control failures are partly due to differences in the speed with which the decision-making circuitry processes basic attributes like taste, versus more abstract attributes such as health. In particular, we test the following two explanations of the questions posed above. First, we hypothesize that dietary self-control is difficult for many because, on average, taste attributes are processed earlier than health attributes. Second, we hypothesize that individual differences in dietary self-control abilities are related to differences in the relative speed with which taste and health attributes are processed. In particular, we hypothesize that the ability to deploy self-control decreases with increasing delays of the start of health information processing, relative to taste.

These two hypotheses are inspired by a robust implication of several models of choice, including Decision Field Theory (Busemeyer & Townsend, 1993), the Drift-Diffusion Model (Ratcliff & McKoon, 2008), and the Leaky Competing Accumulator Model (Usher & McClelland, 2001). As illustrated in Figure 1, these models assume that the brain makes decisions by computing a dynamic signal that measures the value of two possible responses: in the example described above, choosing to eat a cookie (“Yes”) or declining it (“No”). At every instant, the signal estimates the relative value of “Yes” by adding a noisy measure of the value of the item, given by a weighted linear combination of the item’s attributes plus some Gaussian noise. A choice is made the first time one of two pre-established barriers is crossed: “Yes” if the upper one is reached first, and “No” in the other case. A critical feature of this model is that an attribute can affect the evolution of the value signal only if it is being represented at that instant. Thus, as illustrated in Figure 1, if taste information is processed starting at 500 ms, but health information is processed only starting at 1000 ms, then during the interim the value signal is influenced only by taste. This would move the signal closer to the choice associated with the tastier item and decrease the likelihood of choosing the healthy item. In fact, if the onset of health information processing is sufficiently delayed, a choice will be made before it can influence the decision at all.

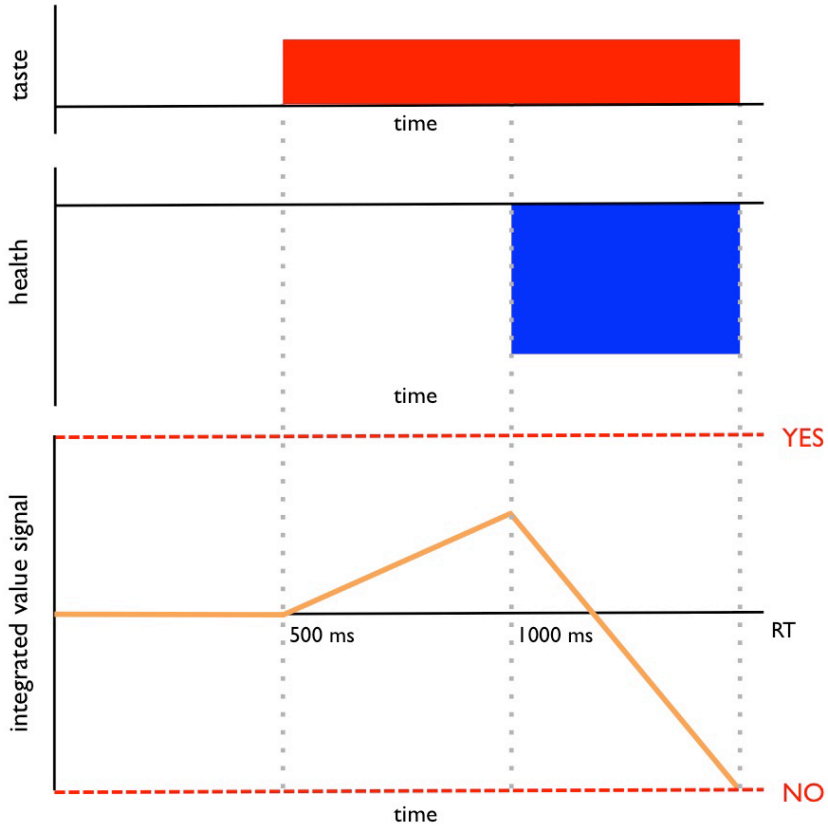

Figure 1. Illustration of how the time at which attributes are computed could affect choices in sequential integrator models of choice. Consider an example in which the subject needs to decide whether to eat a food that is somewhat tasty, but which is quite unhealthy. Suppose also that taste information is computed from 500 ms until a decision is made, whereas health information is computed only after 1000 ms. A choice is made by computing a dynamic value signal that starts at zero and sequentially integrates a weighted linear sum of the attributes represented at this time plus noise. These weights can be interpreted as the "true preferences" that the individual is trying to maximize. A choice is made the first time the integrated value signal crosses one of the two pre-established barriers. For simplicity, the integrated value signal is depicted without noise. The gap in processing times then implies that only taste information affects the value signal between 500 and 1000 ms, which increases the likelihood that an unhealthy choice is made.

These hypotheses are related to the previous findings discussed above, but posit a different and novel mechanism underlying difficulties in the ability to exercise self-control. The emphasis here is on the speed at which attributes are processed, and not on differences in the ability to process the attributes, or on how they would be weighted under the individual's true preference (i.e., if choices could be made without noise). In particular, differences in the speed of processing can generate

differences in the relative influence of health and taste on the final decision, but such weighting differences need not be due to differences in processing speed. This distinction is important because many contextual variables (e.g., time pressure, marketing) could have a direct effect on the relative speed at which different attributes are computed, without directly affecting how they influence choice.

In order to test these hypotheses, it is necessary to have a dynamic measure of the extent to which different attributes are integrated into the choice process while the decision is being made, and with sub-second temporal resolution. Building on the pioneering mouse-tracking literature (Dale, Kehoe, & Spivey, 2007; Dshemuchadse, Scherbaum, & Goschke, 2013; Freeman & Ambady, 2009; Freeman, Ambady, Rule, & Johnson, 2008; McKinstry, Dale, & Spivey, 2008; Song & Nakayama, 2008; Spivey, Grosjean, & Knoblich, 2005), we asked subjects to indicate their choice using a continuous computer mouse movement. By measuring the mouse positions along the trajectory, we can infer the state of the choice process prior to the final decision. Scherbaum, Dshemuchadse, and Goschke (2012) propose a related link between connectionist models and intertemporal choice.

2.2 Methods

Subjects. Twenty-eight students (25% female, 93% right-handed) participated in the experiment, which was approved by Caltech's Institutional Review Board. Subjects received $25 for their participation. They were asked to fast for four hours prior to the experiment. Compliance with this instruction was monitored by self-report.

Task. Subjects performed the following three tasks, always in the same order.

First, over three separate blocks, they rated 160 different foods in terms of 1) tastiness ("How tasty is this food?"), 2) healthiness ("How healthy is this food"), and 3) overall liking ("How much would you like to eat this food at the end of the experiment?"). Block and stimulus order were randomized across subjects. All ratings used a five-point scale (-2=very little to 2=very much). Based on pre-testing, we selected food stimuli in which health and taste had minimal correlation for a typical subject ($r=.2$, $p<.001$). The foods included fruits, candies, chips, and granola bars. On every rating trial, one color image was presented at the center of the screen (~ 5.6 x 4.3 inches), on a black background, and in high resolution (1680 x 1050 pixels) on 20.1-inch 96 DPI LCD monitors.

Second, subjects were asked to read a short excerpt from WebMD.com on the importance of healthy eating (Fig. S1). This was done to increase the frequency with which they exhibited dietary self-control in the food choice task described next.

Third, subjects made 280 binary choices among randomly-selected pairs of foods. Food pairs were selected for the choice task to include an equal number of trials for each combination of taste-health rating pairs, with foods rated as neutral on both taste and health excluded. Choices were made in 40-trial blocks, with short rests in between. There were two types of trials. In six randomly selected blocks, subjects used the computer mouse to indicate their choice ("mouse trials"). In the other block, they used the keyboard to enter their response ("keyboard trials"). Subjects were informed of the type of the trial at the beginning of each block.

Figure 2A describes a typical mouse trial. The trial began with the display of the Start Box at the bottom-center of a black screen. Subjects had to click in this box with the mouse to start the trial. This was followed by a blank screen of random duration drawn from a uniform distribution of 200-500 ms (mean duration 350 ms), during which the mouse cursor disappeared from the screen. At the end of this blank screen, the mouse cursor reappeared in the bottom-center of the screen, where the start box used to be, and pictures of two foods appeared on the top-left and top-right portions of the screen, surrounded by equal-size white boxes (174 x 131 pixels). Subjects were instructed to make a choice by using the mouse to move the cursor continuously to the location where their preferred food was displayed, either top-left or top right. The location of each food (left versus right) was randomized on each trial. Once the subject clicked inside one of the food boxes, a choice was recorded. Subjects were not constrained in how long they could take to enter their response, but were instructed to respond as quickly and accurately as possible. Importantly, the actual foods did not appear until a mouse movement was detected. This was done in order to promote fluid mouse movements during the choice process. Cursor velocity was significantly greater than zero for normalized time points 1 through 99 (two-sided t-test; $p<.01$, uncorrected for multiple comparisons), suggesting that movements were fluid on average. Trials were separated by a fixation of random duration (uniform random 400-700 ms; mean 550 ms).

Keyboard trials had an identical structure except for the following differences. Subjects kept two fingers over the keyboard buttons associated with left and right responses at all times, and did not

use the mouse. Trials were initiated by pressing the spacebar. Foods appeared automatically after the middle blank screen. Subjects indicated their choices by pressing the left or right keys.

Subjects cared about the choices because at the end of the experiment, their choice from a single randomly selected trial was implemented. Subjects were asked to remain in the lab following the choice task until they had eaten the food, or 30 minutes had expired, whichever occurred first. Since they did not know in advance which trial would be selected, their best strategy was to treat every choice as if it were the only one.

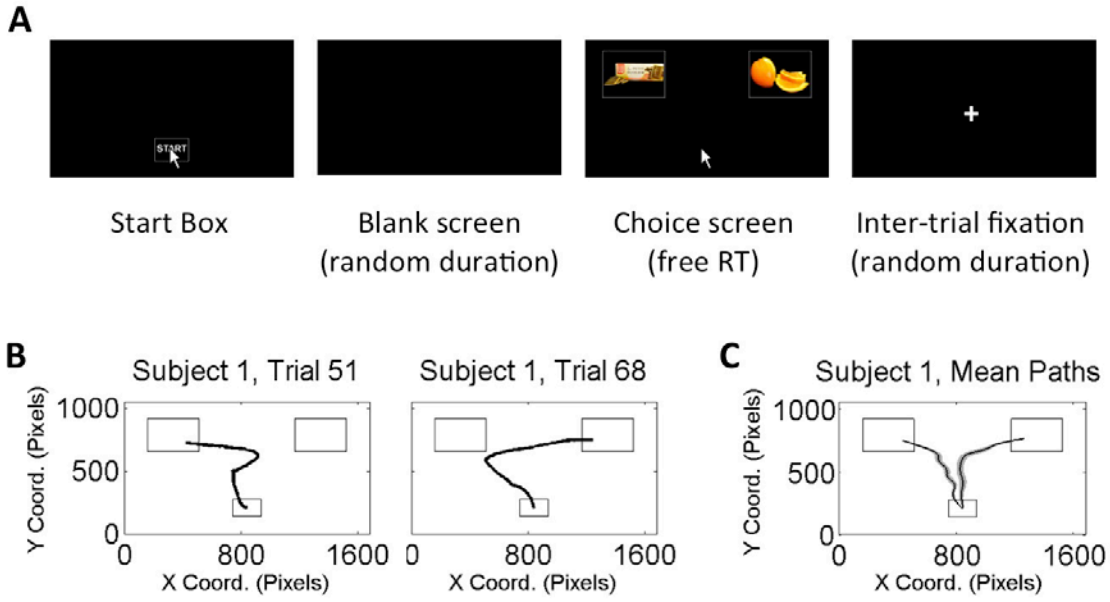

Figure 2. Task description. A) Structure of trials in the mouse tracking choice task. B) Representative paths with two (left) and one (right) changes of direction along the x-axis. C) Average path over all trials for the same subject, conditional on choosing left and choosing right.

Mouse tracking. During mouse trials, the cursor's X, Y position was tracked using Psychophysics Toolbox (Brainard, 1997) with a temporal resolution of 67 Hz. For the purpose of the analyses below, we shifted and normalized the coordinates so that the center of the Start Box (the location at which the cursor appears) is at (0,0), the pixel clicked to select the left option is at (-1,1), and the pixel clicked to select the right option is at (1,1).

Data pre-processing. To reduce some of the noise caused by trials where subjects seemed to have difficulty complying with the instruction to make continuous and rapid mouse movements, we applied two pre-processing steps to the mouse data. First, for each subject, mouse trials with reaction times (RTs) greater than two standard deviations above their mouse trial mean were excluded from further analysis (mean 13.3% of trials). For comparison, the mean RT in mouse trials was 1814 ms (S.D. = 524 ms). Second, we also removed a small fraction of trials in which the mouse trajectory crossed the y-axis more than three times (mean 3.3% of trials per subject; 26 of 28 subjects had at least one such trial). In contrast, as illustrated in Figure 2B, most trials involved one or two changes of direction. We emphasize that these criteria were selected *a priori*, are not based on their effect on the analyses below, and do not have a qualitative impact on the results.

Time normalization. In order to facilitate comparison across trials with large differences in reaction times, time was normalized in most of the analyses described below. In particular, the duration of every trial was sliced into 101 identical time bins, with $t = 1$ indicating the start position at (0,0), and $t = 101$ indicating the time a choice was recorded. Thus, the index t provides a measure of normalized time.

We have also carried out all of the analyses using absolute time and found similar but noisier qualitative results. We believe that this is due to sizable cross-subject differences in reaction times, which might reflect differences in underlying cognitive processing speeds. Such differences would render these analyses problematic, because absolute time then reflects different stages of cognitive processing in different subjects. In contrast, this issue does not arise in the time-normalized analyses. For this reason, here we work with normalized time.

Trajectory analyses. Below, we report various types of trajectory analyses. These analyses come in two different types.

First, we estimate several linear regressions of how the trajectory angle at every normalized time point is affected by the attributes of the foods. The trajectory angle at time t provides a measure of the direction of movement between t and $t+1$, and is normalized such that -45^o always indicates a direct movement towards the left item, 0^o always indicates a movement straight upwards, and $+45^o$ always indicates a direct movement towards the right item. These regressions are estimated at the individual level, and then the relevant estimated coefficients are pooled across subjects.

Second, we also use the results from these regressions to identify the earliest (normalized) time at which various properties of the foods have a significant and lasting influence on the trajectories.

This is done by carrying out a one-sided test of the hypothesis that the estimated regression coefficient of interest is significant at the 5% level, for each individual and time index. We use one-tailed tests because we are interested in when they become *positively* significant. The earliest time window at which the test is satisfied is then labeled as the earliest time at which the variable of interest has a significant positive impact on the mouse trajectory for that subject. Importantly, the test requires that the variable maintain its significance from that time point through the end of the trial. Group comparisons are then performed by carrying out two-sided t-tests on the relevant distribution of individual statistics.

For expositional reasons, further details on the trajectory analyses are provided in the relevant parts of the Results section.

2.3 Results

To test the hypothesis that the relative speed with which taste and health are computed affects dietary self-control, it is necessary to have a good measure of how the influence of these attributes evolves over the decision process. Here we constructed such a measure by asking subjects to indicate their choice using a continuous mouse movement, with positions along the trajectory serving to measure the state of the choice process prior to the final decision. This is a critical aspect of our experiment because it allows us to identify, prior to the final decision, the time at which the taste and health attributes begin influencing choice. The method is justified as long as movements of the computer mouse are correlated at every instant with changes in the integrated value signal illustrated in Figure 1. In this case, if the trajectory of the mouse at time *t* is affected by health information, then we can conclude that the underlying relative value signal of interest is also being modulated by health at that time.

To see the power of this technique, consider a hypothetical example. Suppose that we look at the mouse trajectories for a given subject and find that they have the following two properties. First, the relative position between the left and right options, as measured by x-position, is significantly related to the relative taste rating (given by $taste_R - taste_L$) after 300 ms, but not before. Second, the x-position is also related to the relative health rating of the foods (given by $health_R - health_L$) after 800 ms, but not before. Then we could conclude that the taste attributes were processed by the choice circuitry 500 ms earlier than the health attributes. Furthermore, the technique allows us to test whether differences in the earliest time at which health and taste information are processed could explain differences in dietary self-control across individuals.

Paradigm validation. We began the analysis by carrying out several tests designed to address concerns about the validity of the paradigm.

One natural concern is that requiring subjects to indicate their choices with a mouse, instead of more common methods such as a button press, might affect the choices or the computations of the choice circuitry. In order to address this concern, we asked subjects to make otherwise identical choices using either the keyboard or the mouse, and then compared the behavior generated by both conditions. Figure 3A depicts the average choice curve for the group, which relates the relative value of the left item (measured independently by $\text{liking}_L - \text{liking}_R$) to the probability of choosing left. We found that the two cases led to indistinguishable choice curves (mean logistic slope mouse = -0.37, mean logistic slope keyboard = -0.33; p=.22, t(27)=1.25, two-tailed). Figure 3B plots the reaction time curves as a function of choice difficulty (measured by $|\text{liking}_L - \text{liking}_R|$). Although mean reaction times were longer in the mouse condition (mouse mean = 1814 ms; keyboard mean = 1258 ms; t(27)=4.54, p=.0001, two-tailed), likely due to the additional motor complexity of the mouse-tracking task, in both cases RTs decreased with choice difficulty, as measured by the absolute value of the liking rating differences (mouse mean linear regression slope = -0.036; keyboard mean linear regression slope=-0.051; p=.18, t(27)=-1.36, two-tailed). This pattern is a common property of binary choice (Ashby, 1983; Luce, 1986; Milosavljevic, Malmaud, Huth, Koch, & Rangel, 2010; Ratcliff & Rouder, 1998; Rolls, Grabenhorst, & Deco, 2010), and a key prediction of the type of integrator models of decision-making discussed above. These results suggest that the nature of the choice process was not substantially altered by the use of mouse tracking.

Another natural concern is that the mouse trajectories might not provide a good measure of the extent to which the properties of the foods are dynamically incorporated in the decision process. This could happen, for example, if individuals always made their decision before moving the mouse, and then made a very rapid straight movement to the chosen option. The paths depicted in Figure 2B suggest that this was not the case, but we also conducted a more rigorous test, which was constructed as follows. For each normalized time window (from t = 1 to t =101), and each individual, the local angle trajectory was regressed against the relative value of the two options (as measured by $\text{liking}_R - \text{liking}_L$). The estimated slope, plotted in Figure 3C for each time window, provides a measure of the extent to which the value information is influencing the choice process at different times. We found that, across the group, the earliest time at which the value information

had a sustained, significant effect on the trajectories, without ever returning to being not-significant, was t = 56 for the group. Furthermore, the impact of value rose gradually over the course of the trial. In contrast, if individuals had made their choice before moving their mouse, and then had responded very rapidly with a straight trajectory, the time-course depicted in Figure 3C would look approximately like a step function. This suggests that the information about the choice options is incorporated gradually into the choice process, and that the magnitude of the effect could be measured with our methodology.

Finally, we tested if there was cross-subject variation in the ability to exercise dietary self-control. This is important since we are interested in looking at the extent to which individual differences in self-control can be explained by differences in the speed with which taste and health attributes are computed. To measure self-control success for each subject, we first divided trials into those that entailed a self-control challenge and those that did not. Challenge trials involved a choice between two foods in which one food had a higher health rating and the other had a higher taste rating. In this case, successful self-control occurred when the subject chose the healthier item, whereas choosing the tastier item constituted a self-control failure. In contrast, in non-challenge trials one of the items had a higher rating in both dimensions, and thus there was no need to exercise self-control. Given this, we defined a subject's self-control success ratio (SCSR) to be the fraction of challenge trials in which self-control was successfully exercised. Figure 3D depicts the distribution of SCSRs across the group, and shows that there was substantial variation across subjects (mean = 25%, range: 2-78%). There were no significant RT differences between challenge and non-challenge trials (challenge mean = 1782 ms; non-challenge mean = 1701 ms; $t(27)=.78$, $p=.44$, two-tailed), consistent with the hypothesis that both types of trials utilize a similar decision process.

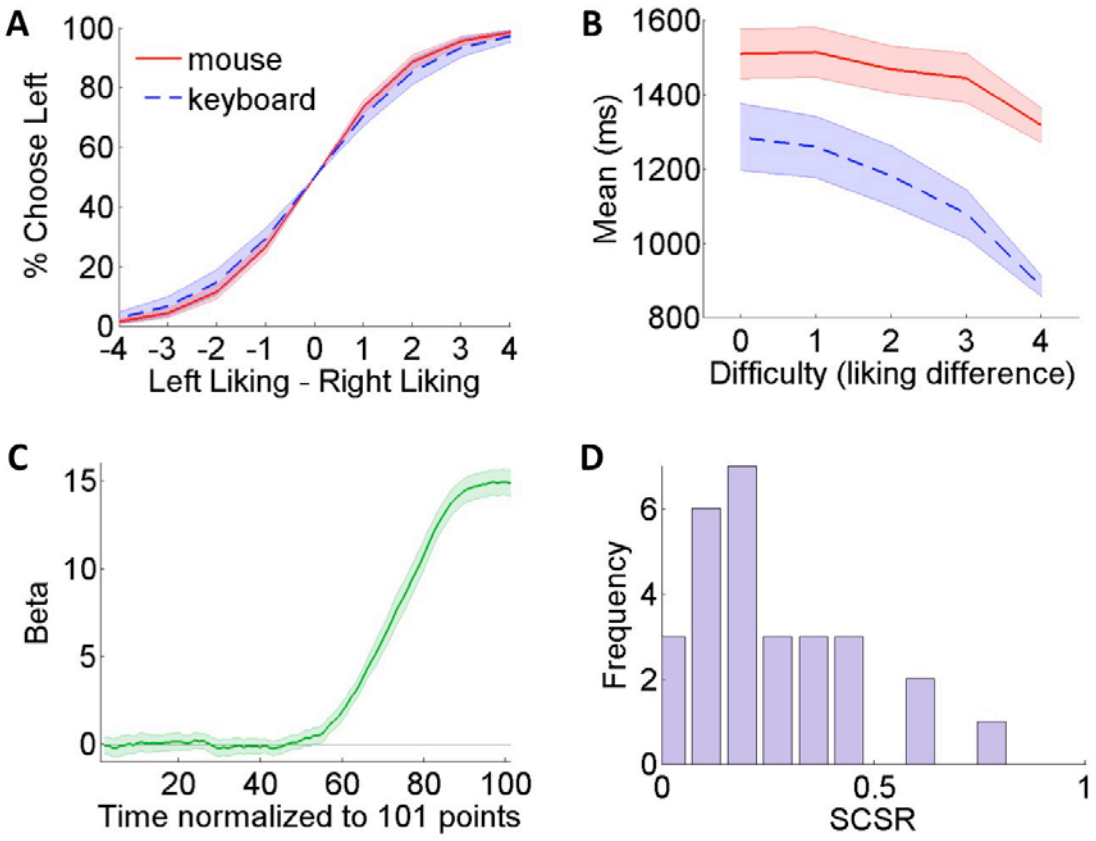

Figure 3. Paradigm validation. A) Psychometric choice curve for mouse and keyboard choice trials. B) Mean reaction time versus choice difficulty, as measured by $|liking_R - liking_L|$. A zero here indicates a difficult choice between two equally liked foods, and a four indicates an easy choice between foods with opposite liking ratings. C) Impact of relative food value (as measured by $liking_R - liking_L$) on the mouse's angle trajectory, measured at each normalized time window. Error bands denote standard errors computed across subjects. D) Distribution of self-control success ratio (SCSR) across subjects.

Taste is reflected in the choice process earlier than health. We now address the first hypothesis of the study: does taste information tend to influence the choice process earlier than health information?

To test this question, we examined how mouse angle trajectories across successive normalized time windows were weighted on taste and health information. For every subject and time window, we estimated a linear regression of the angle trajectory on relative taste ($taste_R - taste_L$) and relative health ($health_R - health_L$). The resulting estimated coefficients provide a measure of the extent to

which each attribute influences the choice process at that time window. Figure 4A summarizes the effect of these regressions by plotting the mean time course of the estimated coefficients, as well as the times at which the distribution of estimates for the group first became significantly larger than zero without ever returning to being not-significant. As hypothesized, mouse trajectories were influenced by taste at significantly earlier time slices compared to health (taste: t = 55; health: t = 67; $p<.05$ threshold, one-tailed). We carried out an additional test of the same hypothesis by estimating, separately for each individual, the earliest normalized time at which each attribute became significant and then comparing their distributions. We found that health never became a significant influence on the trajectory for 32% of subjects, whereas for taste this happened for zero subjects. Furthermore, as shown in Figure 4B we also found that the mean earliest time for a significant effect of health was later than the mean earliest time for taste (mean difference = 9.11, $p=.032$, $t(18)=-1.97$, one-tailed).

Together, these results suggest that, on average, taste information began to influence the choice process about 9% earlier than health information. Given the average duration of the trials, this corresponded to an onset of health attribute processing by the choice circuitry approximately 195 ms later than the onset for taste.

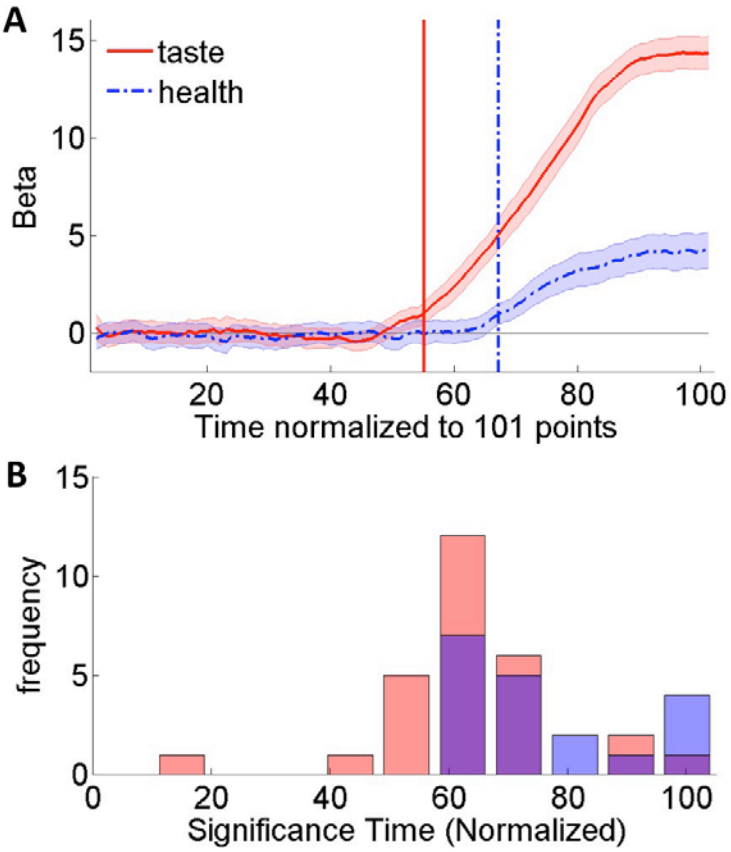

Figure 4. Average effect of taste and health attributes on mouse trajectories. A) Evolution of the effect of relative taste (given by $taste_R - taste_L$) and relative health (given by $health_R - health_L$) on the change in trajectory angle at each normalized time window. Lines denote group means. Error bands denote standard errors. Vertical lines are the earliest times at which each attribute becomes and remains significantly greater than zero (see text for details). B) Distribution of individual estimates of the earliest time at which each attribute significantly affects the trajectory.

Dietary self-control and the relative speed at which taste and health attributes are computed. We next addressed the second hypothesis of the paper: are individual differences in dietary self-control associated with differences in the relative speed at which the health and taste attributes are computed?

To answer this question, we began by comparing how taste and health information were reflected in the mouse trajectories for individuals with high versus low self-control, as defined using a median split of the SCSR statistic. As shown in Figure 5A, we found substantial differences across the groups. In particular, in the high self-control group, the paths for taste and health were quite similar

and became significantly greater than zero at approximately the same time (health: t = 67; taste: t = 60). In contrast, for the low self-control group, the latency at which taste became significant was similar to that of the high self-control group (t = 56), but that for health occurred much later (t = 85). As before, we tested for the significance of these differences by estimating at the individual level the earliest time at which health attributes become significantly greater than zero, and then comparing the distributions. The normalized time at which taste and health became significant did not differ for high self-control subjects (mean health: t = 69.50; mean taste: t = 68.54; mean difference = -2.9; t(9)=-0.46, p=.33, one-tailed), but they were significantly different in the low self-control group (mean health: t = 81.22; mean taste: t = 61.35; mean difference = -16, t(8)=-2.5, p=.02, one-tailed).

Figure 5B takes this analysis a step further. For each subject, we computed a measure of the "computational advantage" of taste, given by the earliest time at which taste significantly affected and continued to affect the trajectories minus the earliest time at which health did. We then estimated a linear regression of the SCSR measures of self-control against the individual measures of the 'computational advantage' for taste versus health. This regression tests the extent to which individual differences in self-control can be attributed to individual differences in the relative speed at which the health and taste attributes are processed. Since health never became significant for 32% of the subjects, we carried out the test in two different ways. In one case we assumed that, for the problem subjects, health first became significant at t = 102, one time point after the final time point, and estimated the regression using all of the subjects (linear regression slope=.005, p<.0004, R^2= 0.39). In the other case, we simply excluded these subjects, carrying out the regression using only the 67.9% of subjects for whom health became significant during the trial. This yielded a similar result (linear regression slope=.005, p=.02, R^2=.29).

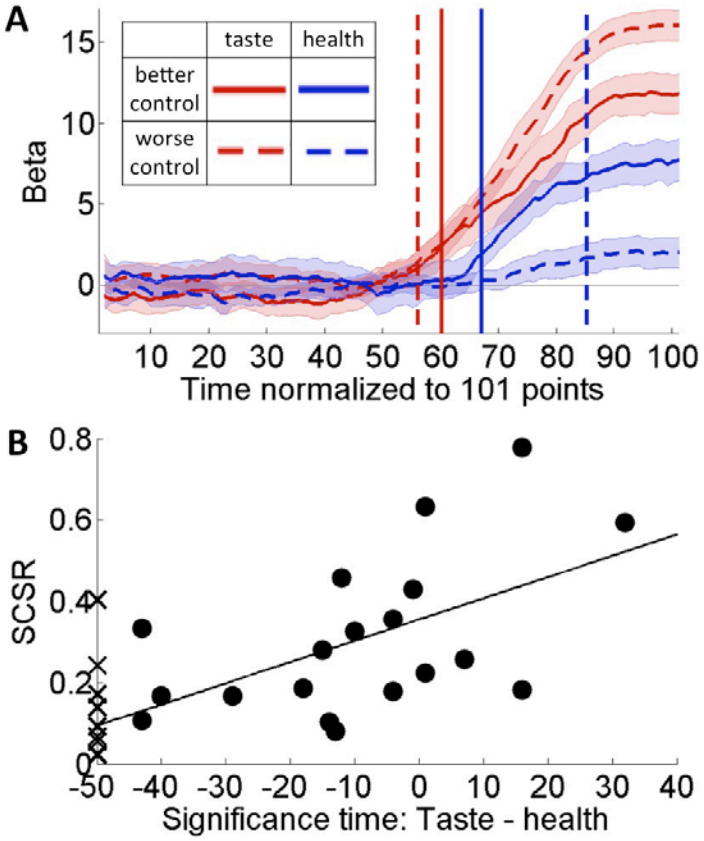

Figure 5. Individual differences in self-control. A) Evolution of the effect of relative taste and relative health on the mouse trajectory, computed separately for individuals with high and low self-control. Error bands denote standard errors. Vertical lines are the earliest times at which each attribute becomes significantly different than zero. Note that the health attribute never becomes significant for the low self-control group. B) Self-control success ratios against individual measures of the computational advantage of taste over health, as measured by the difference between the earliest time at which health becomes significant minus the earliest time at which taste becomes significant. Circles indicate subjects for whom both taste and health had a significant effect on trajectories before the end of the trial. Crosses indicate subjects for whom it did not. For those subjects, the value of the earliest time for health was set to $t = 102$. See text for details.

We also carried out an additional analysis to rule out a potential confound in these results. To understand the concern, suppose that health and taste begin to influence the true data generation process at the same instant, but that on average the weight given to taste in the instantaneous value signal being integrated is larger than the average weight given to health. In this case, if the noise associated with both signals is sufficiently similar, a test based on significance time will identify an

earlier entrance for taste than health based purely on the ratio of the weight to the noise, and may falsely overestimate the difference in significant times. Furthermore, this potential bias would be correlated with the relative "true" weights of health and taste, and thus with the SSCR measure, which could lead to a spurious correlation. Figure 6 depicts the results of an analysis designed to control for this problem. For every subject and for every percentile value *p* between 10% and 90% we use the previous regression results to identify the time at which each of the attributes reaches and never drops below *p*% of its final magnitude. This leads to the plot in Fig. 6A, which depicts the speed at which health and taste reach different thresholds. We then fit a cubic polynomial to the estimates of each subject, separately for health and taste, and interpolate its value to the y-axis, which provides a measure of the earliest time at which each of the attributes begins and continues to influence the mouse trajectories. Note that this measure is free from the confound described above, and is fairly robust to noise. As shown in Figure 6B, we find that the differences in earliest significant time identified with the previous analysis are highly correlated with those identified with this new method (R^2=.51, p=.006). Likewise, the alternative measure of the relative earliest times can still significantly predict a sizable fraction of the individual differences in self-control (Fig. 6C; R^2=.13, p=.059).

Together, these results suggest that a sizable fraction of individual differences in dietary self-control, on the order of 13-39%, might be attributable to differences in the relative speed at which the health and taste attributes are processed by the choice circuitry.

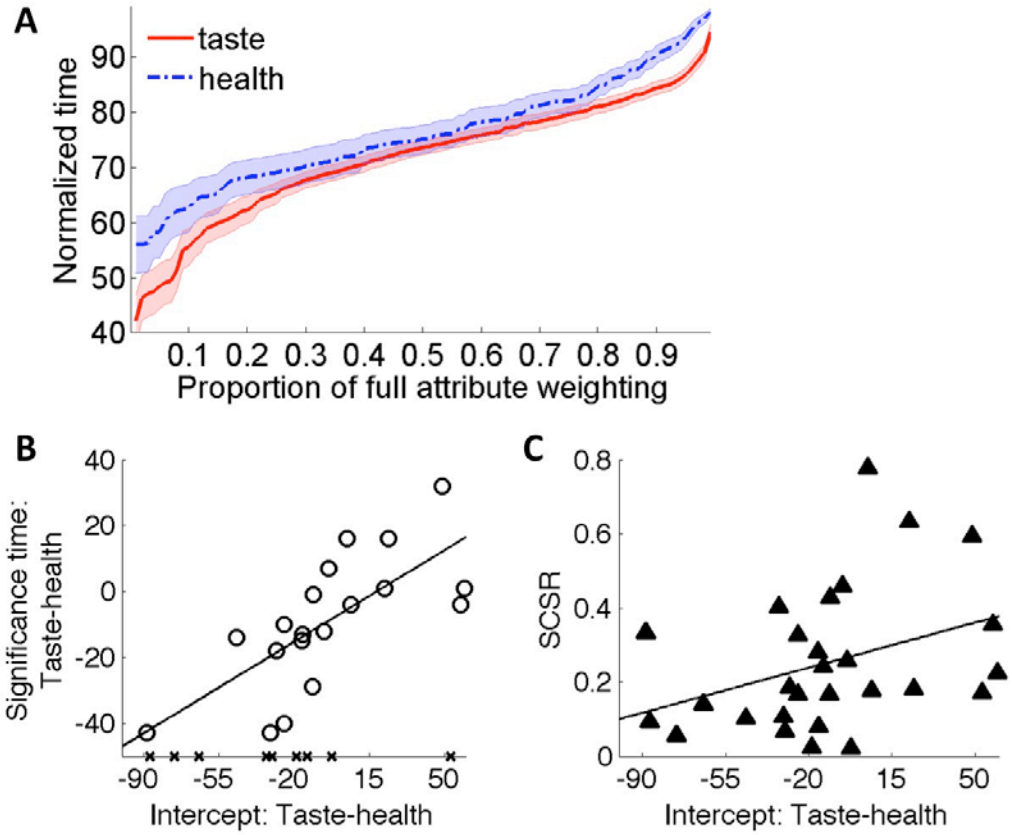

Figure 6. Control analysis. A) Estimated normalized time at which each attribute's weighting on cursor trajectory had accumulated a percent of its final value. Error bands denote standard errors. B) Cross-subject scatter plot showing of the relationship between the two measures of the time advantage of taste over health attributes. Crosses indicate subjects for whom health did not have a significant effect on trajectories before the end of the trial. C) Cross-subject scatter plot showing the correlation between of self-control success ratios and the intercept-based measure of the time advantage for taste.

Relationship between relative computation time and attribute decision weights. The previous results suggest that the relative speed at which health and taste are computed is related to the ability to exercise dietary self-control. We conclude our analyses by carrying out a *post hoc* analysis to investigate further the nature of the relationship. This was done in two steps.

First, Figure 7A depicts the extent to which differences in SCSRs are related to differences in the weights given to the taste and health attributes in the final decision. The plot was constructed as follows. For every subject we computed a logistic regression of the probability of choosing the right

item against relative taste ($taste_R - taste_L$) and relative health ($health_R - health_L$). This provided individual estimates of the weights that each attribute received on the final decision. We then plotted the individual estimates against the SCSR, and for each type of attribute we carried out a linear regression of the SCSR measure against the individual estimates. Both attributes were highly significant and, on their own, explained a sizable fraction of the individual differences in self-control (health: p=.0004, R^2=.39; taste: p=.0004, R^2=.38).

Second, we hypothesized that the final decision weights for taste and health are influenced by the speed with which the attributes were computed. This hypothesis is based on the idea that the earlier an attribute starts influencing the choice process, the longer it is integrated and thus the stronger its effect on the final choice. This hypothesis is consistent with the type of integrator decision models described in the Introduction. To test this hypothesis, we regressed the individual estimates of the final relative decision weight of health minus taste on the "computational advantage" of taste, as defined in the previous section. As shown in Figure 7B, we found that the two variables had a significant and sizable relation (slope = -0.02, p=.001, R^2=.487).

Given these results, we also hypothesized post-hoc that the main mechanism through which the speed of attribute processing influences SCSR is by modulating the weight assigned to each attribute in the final choice. To test this hypothesis, we carried out a mediation analysis. Using this method, we found that the final weight of taste was a significant mediator of taste's time-to-significance, reducing path strength by 79.14%, and making taste's time-to-significance not-significant for predicting SCSR (p=.60). This was also true for the health coefficient, reducing path strength by 88.26%, and also making health's time-to-significance non-significant for predicting SCSR (p=.75). See Figure S2 for details.

Together, these results suggest that attributes that are computed earlier have a stronger weight in the final decision, likely by increasing the amount of time during which they are integrated into the decision. In particular, this suggests that initial delays in the computation of an attribute are not made up later on during the choice process, for example by a heavier weighting closer to the time of decision.

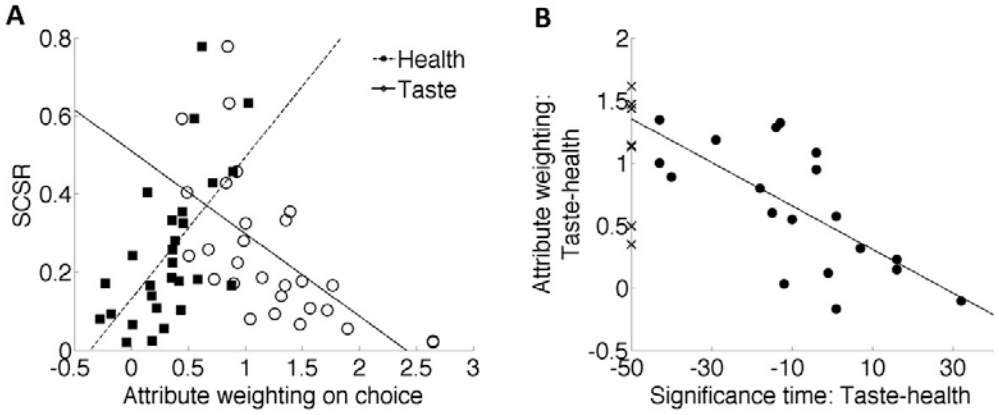

Figure 7. Relationship between final decision weights and speed of computation. A) Plot of the individual self-control success ratios versus the individual final decision weights for the health and taste attributes. Lines denote best linear fits for the regressions of SCSRs on each of the attribute weights, separately. B) Plot of relative final decision weights versus the "computational advantage" of taste, as measured by the earliest time at which taste significantly affects the trajectory minus the earliest time at which health does.

2.4 Discussion

In this study we propose that dietary self-control failures can be traced back to relative differences in the speed with which the decision-making circuitry processes basic attributes like taste versus more abstract attributes such as health. We find evidence that, on average, taste attributes are processed about 195 ms earlier than health attributes during the choice process. We also find that individual differences in dietary self-control are related to differences in the relative speed with which taste and health attributes are processed. In particular, we find that self-control decreases with increasing lag between the onset of taste and health information processing.

Since several alternative mechanisms are likely to be at work in dietary self-control (Baumeister & Vohs, 2004; Beaver et al., 2006; T. A. Hare, C. F. Camerer, & A. Rangel, 2009; Harris, Hare, & Rangel, 2013; Hutcherson, Plassmann, Gross, & Rangel, 2012; Ikeda et al., 2010; Loewenstein, 1996; Muraven & Baumeister, 2000), we were surprised to find that between 13% and 39% of observed individual differences in self-control ability could be explained by differences in the relative speed with which taste and health attributes are processed. By comparison, previous studies have found that personality traits such as impulsivity or I.Q. can only explain, respectively, around

13% and 5% of individual differences in self-control (de Wit, Flory, Acheson, McCloskey, & Manuck, 2007; Mitchell, 1999; Shamosh & Gray, 2008; Steinberg et al., 2009).

Our results are consistent with basic ideas from a wide class of integrator models of decision making, including Decision Field Theory (Busemeyer & Townsend, 1993), the Drift-Diffusion Model (Ratcliff & McKoon, 2008), and the Leaky Competing Accumulator Model (Usher & McClelland, 2001). Despite technical differences, in all of these models decisions are made by dynamically integrating information for and against the different options (Bogacz, 2007; Rangel & Clithero, 2013). At every instant during the choice process, the decision circuitry receives instantaneous information about the attributes associated with each stimulus under consideration, and integrates them into a dynamic value signal that measures the relative value of the two options. A decision is made when the relative value signal becomes sufficiently large or sufficiently small. These models are an important benchmark for understanding self-control because they have been shown to account, with high quantitative accuracy, for psychometric choice data. Importantly, they predict that the earlier an attribute is processed, the longer it is integrated into the relative value measure, and thus the larger the weight that the attribute receives in the final decision. Thus, these models provide computational foundations for our findings.

Our study borrows heavily from the pioneering mouse-tracking literature, which has been used to study intertemporal choice (Dshemuchadse et al., 2013), visual search (Song & Nakayama, 2008), stereotyping (Freeman & Ambady, 2009), and psycholinguistics (Spivey et al., 2005). However, we use the mouse-tracking data in a different way. In particular, whereas most mouse-tracking papers focus on comparisons between the average paths generated by various experimental conditions, here we use regression analyses to identify the time at which different stimulus attributes begin influencing choice. (See Dshemuchadse et al., 2013; Scherbaum, Dshemuchadse, Fischer, & Goschke, 2010; Scherbaum, Dshemuchadse, Leiberg, & Goschke, 2013 for related uses of mouse-tracking data.)

Our study also suggests a potential mechanism by which the difficulty and abstractness of attributes may influence the choice process (Liberman & Trope, 2008). Rather than simply resulting in a *weaker* representation of attributes like health, as is sometimes assumed, the abstractness of attributes might also affect the *speed* with which they are processed. One limitation of our study is that we cannot disentangle the influence that processing difficulty and processing speed have in

self-control. In fact, it is possible that in some contexts both effects might be correlated. Future work will be needed to tease apart the relative roles of these two mechanisms.

Several implications of these findings merit future investigation. First, these data suggest that slowing down decisions, even if only by adding a waiting period before choice, might increase the relative influence of abstract attributes like health on final choice. To see why, note that at longer decision times, there is a decrease in the fraction of time during which only the taste attribute is computed. Indeed, there is already a literature suggesting the influence of time pressure on choice (e.g., Reutskaja, Nagel, Camerer, & Rangel, 2011; Suter & Hertwig, 2011). Second, our results suggest that self-control could be improved by interventions that increase the relative speed with which health information is processed. Consistent with this, a previous study showed that simply cuing subjects to pay attention to health can improve dietary self-control (Hare, Malmaud, & Rangel, 2011). Third, these findings provide a rationale for regulating marketing practices that increase the relative ease with which abstract attributes such as health are processed. For example, prominently displaying health information such as calorie count may allow more rapid integration of health attributes. Fourth, they provide a potential mechanism through which differences in cognitive abilities might affect self-control. For example, it might be the case that individuals with higher I.Q. are relatively faster at processing abstract attributes, such as health, which increases their self-control. Fifth, the relative speed of attribute processing might also play a role in other types of decisions in which basic and abstract information need to be combined. Potential examples include altruism, where information about the self might be calculated faster than information about others, and any decision involving discounting, where information about immediate rewards might be processed earlier than information about future rewards.

2.5 Supplementary Figures

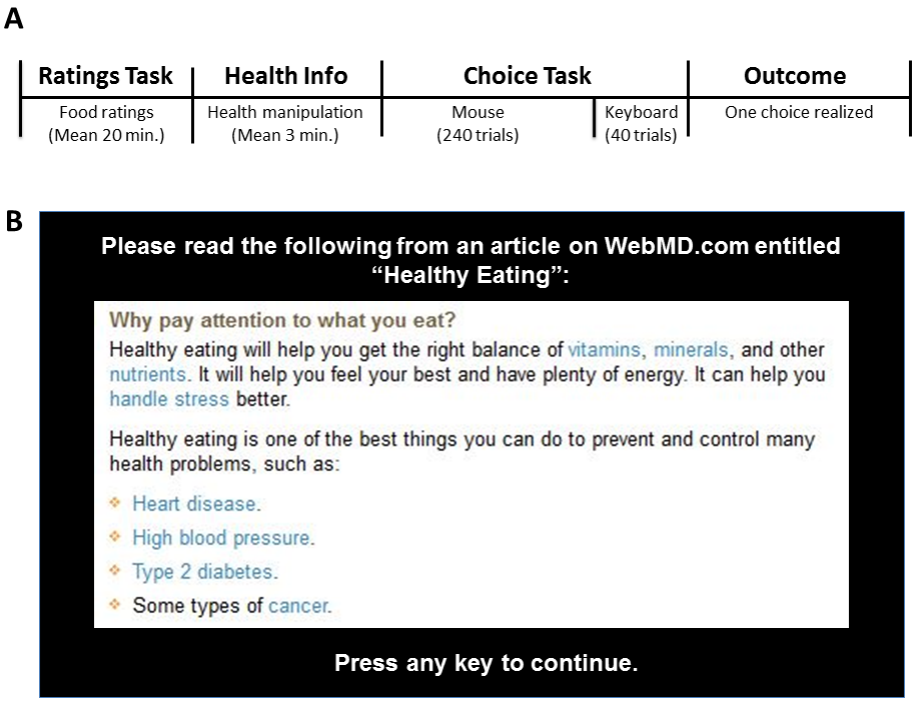

Figure S1. Task details. A) Flow of the experiment. B) Healthy eating manipulation. Mean view time was 1634 ms (min. 587 ms, max. 3409 ms, SD 587 ms).

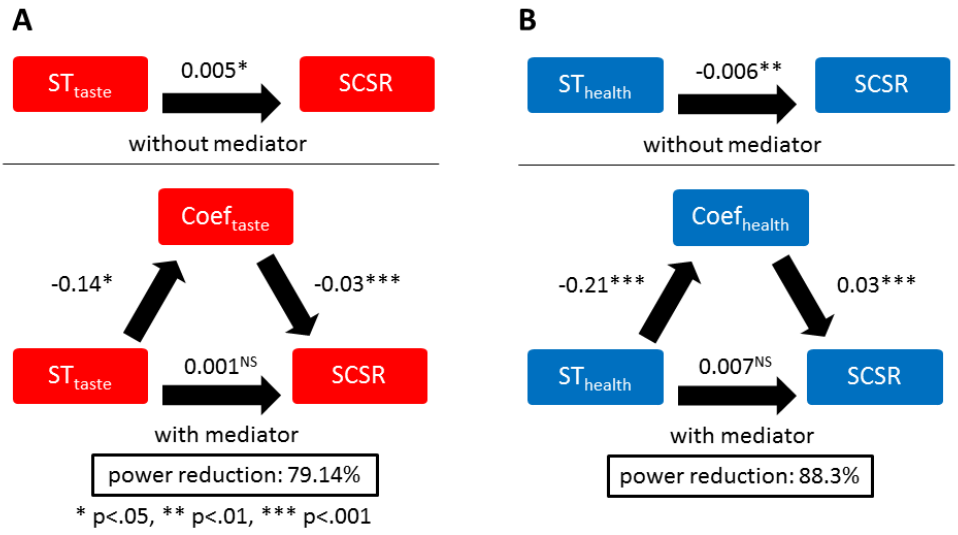

Figure S2. Mediation analyses. Analysis measured the mediating effect of final attribute weights on the association between self-control success ratios (SCSRs) and the earliest time at which the attribute becomes and remains predictive of angle of movement (ST) for A) taste and B) health.

Chapter 3

DELAY OF GRATIFICATION IN INTERTEMPORAL CHOICE IS RELATED TO THE SPEED WITH WHICH IMMEDIATE AND FUTURE REWARDS ARE PROCESSED[3]

3.1 Introduction

The ability to postpone gratification has a profound influence in health, education, financial outcomes, and varies greatly across individuals and contexts. This is powerfully illustrated by Walter Mischel's marshmallow experiments, which have shown that behavioral differences measured at age 4 are predictive of life-outcomes many decades later (Casey et al., 2011; Mischel, 2014; Mischel, Shoda, & Peake, 1988). Individual differences in temporal discounting have been linked to many important real-world outcomes, such as well-being, gambling, drug addiction, and mortality rates (Boyle et al., 2013; Forstmeier et al., 2011; Mischel et al., 1988). Research in psychology and behavioral economics has shown that observed temporal discount rates are inconsistent across time and domains (Ainslie & Monterosso, 2003; Berns et al., 2007; Chabris, Morris, Taubinsky, Laibson, & Schuldt, 2009; Charlton et al., 2013; Frederick, Loewenstein, & O'Donoghue, 2002; Green et al., 1994; Laibson, 1997; Loewenstein & Prelec, 1992).

Given this, a systematic characterization of the neurocomputational mechanisms responsible for this variation is of critical importance. Decades of work in psychology have identified a number of mechanisms that affect the ability to delay gratification during intertemporal choice. Examples include cognitive depletion (e.g., Baumeister & Vohs, 2004), memory (Peters & Buchel, 2010; e.g., Peters & Buchel, 2011; Weber et al., 2007) and cognitive biases (e.g., Fassbender et al., 2014; Kable & Glimcher, 2010; Liberman & Trope, 2008). A growing effort in decision neuroscience has also begun to identify the neurocomputational basis of some these processes (e.g., Ballard & Knutson, 2009; Hare, Hakimi, & Rangel, 2014; Harris et al., 2013; Kable & Glimcher, 2010; Muraven & Baumeister, 2000; Rodriguez, Turner, & McClure, 2014). However, most of the observed variation across individuals and contexts remains unexplained. For example, even non-mechanistic variables such as impulsivity personality traits or IQ can only explain

[3] In collaboration with Antonio Rangel.

around 13% and 5%, respectively, of individual differences in self-control (de Wit et al.; Shamosh & Gray, 2008; Steinberg et al., 2009), and they cannot account for contextual variation.

We propose that the relative speed at which immediate and delayed rewards are processed has a sizable impact on the ability to postpone gratification, and that differences in these speeds can explain an important fraction of the observed individual and contextual variation. As shown below in more detail, this prediction follows directly from the class of sequential integrator models of choice, such as the Drift-Diffusion Model (Ratcliff, 1978, 1980; Ratcliff, Cherian, & Segraves, 2003; Ratcliff & McKoon, 1982; Ratcliff & McKoon, 2008; Ratcliff & Rouder, 1998; Ratcliff & Smith, 2004; Ratcliff & Tuerlinckx, 2002), the attentional Drift Diffusion-Model (Krajbich, Armel, & Rangel, 2010; Krajbich, Lu, Camerer, & Rangel, 2012; Krajbich & Rangel, 2011), the leaky-accumulator model (Usher & McClelland, 2001), and Decision Field Theory (Busemeyer & Diederich, 2002; Busemeyer & Townsend, 1993; Roe, Busemeyer, & Townsend, 2001), which provide highly accurate algorithmic descriptions of the psychometrics of value-based choices, including in intertemporal monetary choice tasks (Dai & Busemeyer, 2014; Rodriguez et al., 2014). An important prediction of these models is that the faster information about a choice attribute is processed by the choice circuitry, the stronger the influence that it has on the choice algorithm (Ratcliff & McKoon, 1989; Sullivan, Hutcherson, Harris, & Rangel, 2015).

We test this hypothesis by combining an intertemporal monetary choice task, where subjects choose between smaller-sooner and larger-later money prizes, with methods pioneered in the mouse tracking literature (Dale et al., 2007; Dshemuchadse et al., 2013; Freeman & Ambady, 2009; Freeman et al., 2008; Lohse & Johnson, 1996; McKinstry et al., 2008; Song & Nakayama, 2008; Spivey et al., 2005). Subjects make their choices using continuous mouse movements, which provide a proxy measure of the state of the decision processes before a choice is actually made with millisecond temporal resolution. This, together with the analysis methodology developed in Sullivan et al. (2015), allow us to ask the following three basic questions. First, are immediate rewards processed faster than delayed rewards on average? Second, can individual differences in the processing speed of immediate and delayed rewards account for a sizable fraction of the cross-individual variation in the ability to delay gratification? Third, can fluctuations in the relative processing speed of immediate and delayed rewards explain contextual fluctuations within subjects?

3.2 Methods

Subjects. 36 subjects recruited from the Caltech community participated in the experiment (59.5% male, mean age 20.9 years old). Subjects were compensated with a $10 show-up fee, as well as the payment option selected in a randomly selected trial of the experiment, as described below. The experiment was approved by Caltech's IRB.

***Task*.** Subjects performed 280 binary choices in a monetary inter-temporal choice task, divided in 25 trial blocks separated by self-timed rests. On each trial, they chose between a constant reference offer of $25 paid at the end of the experiment (not shown in the screen) and a combination offer (shown in the screen) consisting of an amount (<$25) paid that day and another amount paid at a specified future date. Text was shown in white on a black background on 20.1-inch 96 DPI LCD monitors (1680x1050 pixels) using the Psychophysics Toolbox (Brainard, 1997).

As described in Figure 1A, each trial began with the display of a Start Box at the bottom-center of a black screen. Subjects had to click in this box with the mouse to start the trial. This was followed by a black screen of random duration (uniform on 200-500 ms). During the black screen the mouse cursor disappeared from the screen, and reappeared at its conclusion at the bottom-center of the screen. Once mouse movement was detected, "yes" and "no" options were displayed in equal-size white boxes (75 x 50 pixels) on the top-left and right of the screen, with their respective locations randomized for each trial. The combination offer also appeared on the screen at that time. Subjects were instructed to indicate their choices by moving the mouse cursor continuously to either response box. A choice was recorded as soon as the mouse was detected inside one of the boxes. Subjects were given a free response time, but were asked to make each choice as quickly and accurately as possible. Trials were separated by a fixation cross of random duration (uniform on 400-700 ms).

The number of days delay for the delayed payment portion of the combination offer could vary from 5 to 465 days in 10-day increments, and the delayed payment amounts varied from $5 to $65 in $5 increments. The immediate payment portion of the combination offer varied from $0 to $25 in $5 increments.

The first 45 trials consisted of an equal number of three types of trials, randomly intermixed. First, subjective values for all possible combination offers were calculated using a fixed hyperbolic discounting parameter k of 0.009, which was the mean k value of a pilot study. Next, 15 combination offers were randomly selected in which the reference item would be preferred (i.e. was of higher subjective value) over the combination offer by the average subject. An additional 15 offers were randomly selected in which the combination offer would be preferred. A final 15 offers were randomly selected in which the reference item and combination offer were within $1 of each other. After completing these 45 trials, the subject's individual hyperbolic discounting parameters were estimated (see Preference Parameter Estimation below for estimation procedure) and used to generate the remaining 235 choices such that that in 1/3 of trials the combination offer is of greater subjective value, in 1/3 of trials the off-screen reference option is of greater subjective value, and 1/3 of trials within $1 of the subject's point of indifference between the two options.

At the end of the experiment, one trial was selected at random, and the option chosen was implemented. All payments, delayed or not, were implemented using PayPal accounts to equate the associated transaction costs across both options. Subjects were offered, but none needed, assistance in signing up for PayPal. They were also given detailed instructions on how to collect and use the money. To reduce transaction costs, subjects received an email reminder whenever new funds were deposited in their PayPal account. The majority of subjects reported that using PayPal would be "easy" or "very easy" (69%).

Preference Parameter Estimation. Each subject's choices were fit to a standard hyperbolic model of discounting. The model assumes that subjects have linearly additive preferences over prizes received at different dates, and that the value of each such prize is given by

$$SV = \frac{X^{\alpha}}{1 + kD},$$

where X denotes the size of the prize, D represents the delay in days, *k* is the hyperbolic discounting parameter, and α represents the curvature of the utility function. The SV of the combination offers was computed by summing the SV of the immediate and delay payments, and denoted by SV_{Combo}. The SV of the constant reference point, denoted by SV_{Today}, equals $\$25^{\alpha}$. As is standard in the

literature (Chabris et al., 2009; Pawitan, 2001; Rodriguez et al., 2014), the model also assumes that subjects make probabilistic choices, with the probability of choosing the combo option given by

$$\frac{e^{\omega SV_{Combo}}}{e^{\omega SV_{Combo}} + e^{\omega SV_{Today}}},$$

where ω is a non-negative sensitivity parameter determining the influence of the SVs on the choices. Note that choices are random when ω=0, but become more responsive to the relative value of the two options as ω increases. For each subject, we estimated the α, k, and ω parameters that best fitted the choice data using maximum likelihood estimation.

Mouse trajectory preprocessing. On each trial, the mouse cursor's (x, y) pixel coordinates were recorded every 10 ms. To facilitate spatial comparison across trials, recorded pixel values were transformed and normalized so that the initial cursor position is (0,0), and the final cursor position (the point at which the cursor entered the response box) was recorded as either (1,1) if the choice was to the right option, or (-1,1) if the choice was to the left option.

Some outlier trials were excluded based on *a priori* defined criteria. Since our goal is to use mouse location as a proxy for the state of the decision process before the actual decision is recorded, it is important to exclude trials in which subjects did not comply with the instruction to move smoothly and continuously toward the chosen option. In particular, we excluded from further analysis trials with reaction times (RTs) greater than two standard deviations above the mean RT over all subjects (2842 ms; 3% of trials; all subjects had at least one such trial). We also excluded trials with overly complex trajectories, defined as those in which the mouse trajectory crossed the y-axis more than three times (mean 7% of trials per subject; 35 of 36 subjects had at least one such trial).

Trajectory regression analysis. For every subject, we computed the following linear regression analyses of the mouse trajectory, designed to measure how the influence of immediate and delayed payments on mouse positions changed over the course of the decision. A separate regression was estimated at each 10 ms time bin of the decision process. The dependent variable was the angle of the cursor trajectory at each time bin, which provides a measure of the direction of movement, with -45^o always representing a direct movement towards the left option, 0^o always

indicating a movement straight upwards, and +45^o always indicates a direct movement towards the right option. The independent variables were a constant, the relative value of the immediate payments, given by $SV^R_{Now} - SV^L_{Now}$, and the relative value of the delayed payments, given by $SV^R_{Del} - SV^L_{Del}$, with *R* and *L* denoting the right and left locations. For the reasons described in the main text, these values were computed for each trial using the mean *k* and α best fitting parameters in the group. Pooling the results from these regressions, we get a time-course of how the influence of immediate and delayed rewards on the mouse trajectories. Figure 4A shows the mean and SE of these paths over all subjects. Importantly, because reaction times varied within and between subjects, there was significant attrition of trials at late time points. For this reason, regressions were estimated for a subject from the first time point to the point at which fewer than 10 trials were still available for estimation.

Computation of NDTs. The results from the regressions could also be used to estimate the time at which each of the two attributes (immediate and delayed rewards) began to have a significant influence on the mouse trajectories, and remained so until the end of the trial. We refer to these times as NDTI, which stands for non-decision time for immediate rewards, and NDTD, which stands for non-decision time for delayed rewards.

A simple way of computing these measures would be to carry out a one-sided t-test against zero in each time bin, using the distribution of regression coefficients in the group. Unfortunately, these measure can exhibit substantial biases when (as is the case in this data, see Figure 4A), the time-course of coefficients reaches a different plateau for each type of reward.

To side-step this problem, we computed the NDT measures using the method developed in Sullivan et al. (2015), which does not exhibit this problem. The measures are constructed separately for each subject and reward type as follows. First, we identify the plateau reached by the time-course of the estimates; call it by β^{max}. Since the regression estimates of late time bins can be erratic, due the fact that only a small number of trials last that long, we set this value equal to 70% of the maximum estimated β over all time bins. Second, for each value in 0.1β to 0.99β (in 1% intervals), we identify the earliest time at which the time-course reaches that level and stays above it for the rest of the time. Each of these estimates defines a point in 2-dimensional space, as depicted in Figure 4B for a typical subject. Third, in order to take out estimation noise, we fit a cubic polynomial to the set of estimates (fit R^2 for NDTI mean = .900, SD=.141; for

NDTD mean = .903, SD=.085). The intercept of this cubic model with the y-axis is used as the measure of the NDT for that attribute. Figure 4C summarizes the computation of both NDTs in the entire group.

3.3 Results

Subjects made a series of binary choices between a constant off-screen reference option paid $25 on the day of the experiment, and an on-screen combination offer involving two monetary rewards, one delivered that day, and another delivered at a the future date (Fig. 1A, see Methods for details). Response options ("Yes" for the combo offer, "No" for $25 today) were displayed in white boxes at the upper left and right sides of the screen, with their location randomized across trials. Subjects were asked to indicate their choices by moving the location of the computer mouse from the bottom-center of the screen into the box surrounding their preferred option.

3.3.1 Behavioral data and paradigm validation

Estimated temporal discounting parameters. We used maximum likelihood to estimate subject-specific parameters in a hyperbolic discounting model that has been widely used to explain intertemporal choice (Ainslie, 1992; Laibson, 1997; Mazur, 1987; O'Donoghue & Rabin, 1999; Pine et al., 2009). A key element of this model is the discounting parameter k, which measures the strength with which later rewards are discounted. The distribution of estimated k values (Fig. 1B; mean=0.012, SD= .022) are similar to those found in previous studies (Chabris, Laibson, Morris, Schuldt, & Taubinsky, 2008; Frederick, Loewenstein, & O'Donoghue, 2002; Kable & Glimcher, 2007; O'Donoghue & Rabin, 1999). An alternative measure of patience is given by the fraction of trials in which subjects chose the patient option involving the combo offer (mean = 52.1%, SD = .09, max = 70%, min= 38%). Since the two measures were highly correlated (linear regression: slope =-0.04, p=5.1e-07, R2=0.59), we use them interchangeably in the analyses below. Note also that both measures exhibit substantial variation across subjects, which is important for investigating the speed of attribute processing in individual differences.

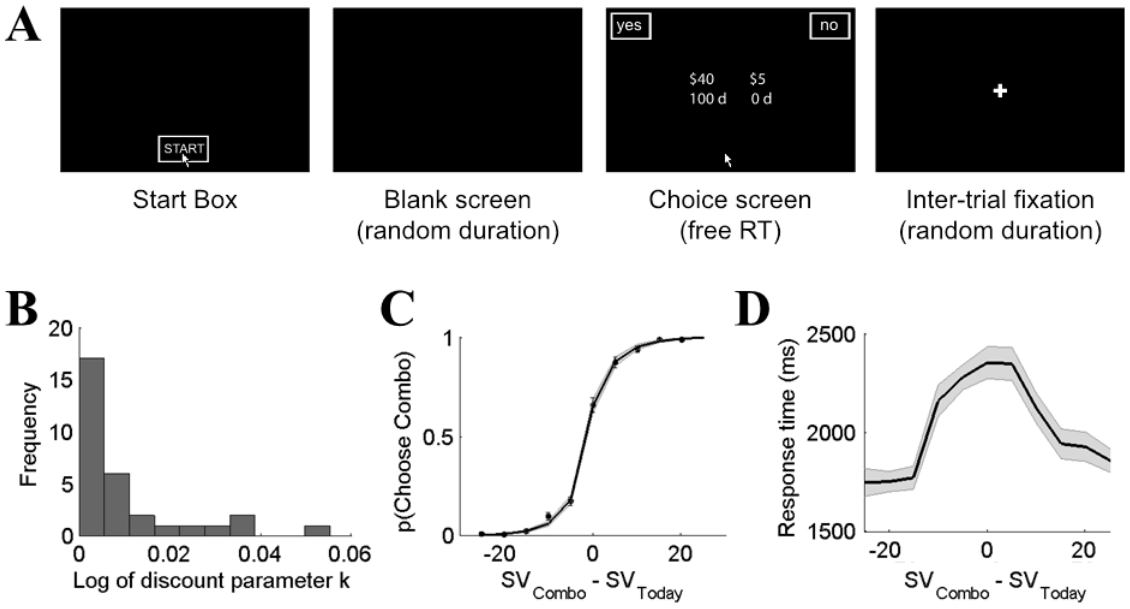

Figure 1. Mouse tracking choice task and behavior. (A) The figure shows an example trial in this task, in which participants used the computer mouse to choose between a fixed reference of $25 that day, or a combination offer displayed on the screen in each trial. This combination offer always featured one payment delivered that day, and another payment delivered in the future. (B) Distribution of the individual temporal discounting parameter k. (C) Average choice curve showing that as the relative value of the combination offer increased, so did its selection frequency. The line and shading represent the average and standard error of the fitted logistic function, and dots and their error bars represent average data and standard error. (D) Mean reaction times as a function of the difference in value between the choice options. The line and shading represent the average data and its standard error.

Basic psychometrics. For every subject and trial, we can compute the subjective value of the reference and combo options (SV_{Today} and SV_{Combo}) using the individually estimated discount parameters (see Methods for details). Using these value measures, we constructed the psychometric choice and reaction time (RT) curves shown, respectively, in Figures 1C and 1D. Both curves exhibit standard properties for this type of data: choices were a logistic function of SV differences (mixed effects logistic regression: slope = 46.97, $p=1.9x10^{-5}$), and RTs increased with choice difficulty, as measured by the SV of the best option minus the SV of the worse option (mixed

effects linear regression of RT on difficulty: slope=-0.0006, $p=1.2x10^{-8}$). Mean reaction time was 2161 ms (SD = 352), which is also similar to those found in the literature (e.g., Hare et al., 2014; Kable & Glimcher, 2007; Weber et al., 2007). Together, these results show that the task exhibited standard psychometric properties, despite the use of mouse-tracking to indicate the responses.

Properties of mouse trajectories. As in previous mouse-tracking studies (Dale et al., 2007; Dshemuchadse et al., 2013; McKinstry et al., 2008; Scherbaum et al., 2010; Spivey et al., 2005; Sullivan et al., 2015) subjects were instructed to indicate their choices by moving the computer mouse continuously towards their chosen option once the trial starts. This is an important part of the design because we used the mouse position of the cursor as a proxy measure for the state of the choice process prior to the final decision, which allow us to investigate the time at which immediate and delayed reward information begin influencing the choice process.

The validity of this proxy measure would be reduced if, for example, subjects simply moved the mouse quickly to start the trial, paused there while making the choice, and only moved the cursor in a direct path toward their choice once a choice was made. Clearly, in this case the mouse trajectory would not reflect much of the choice process. Fortunately, visual inspection of cursor paths suggested that this was not the case, as illustrated in the sample and mean trajectories depicted in Figure 2. In addition, for each subject we carried out a one-sided t-test at each time point within the trial to test if cursor velocity was greater significantly greater than zero at all times. Because reaction times varied widely between and within subjects, we carried this analysis using normalized time, constructed by interpolating movement in each trial to 100 equal time bins. For 94% of subjects, mouse movement was significantly greater than zero from t=1 to t=99 ($p<0.05$). For the remaining two subjects, cursor velocity was significantly greater than zero in, respectively, 97% and 94% of the time bins. Together, these results indicate that subjects generally complied with the request to move the mouse continuously to indicate their choice, and provide support for the use of cursor movements as a proxy for changes of the choice of the state process.

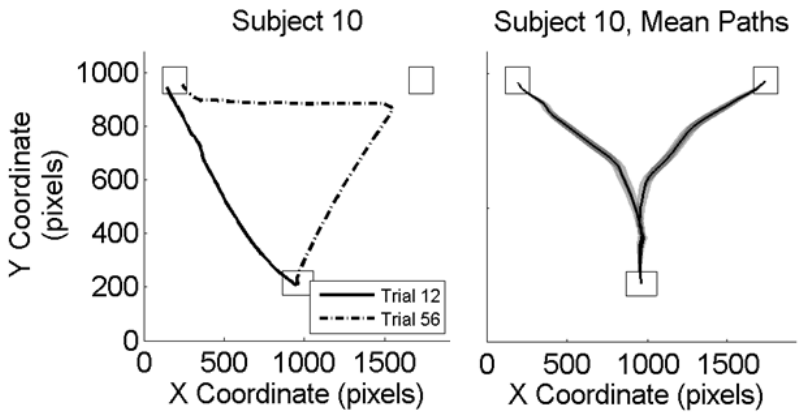

Figure 2. Example mouse trajectories. Two trials for the same subject are shown in the left panel, both featuring a selection of the left-hand choice. The path depicted with the dotted line featured a change of mind, and the path depicted with a solid line exhibited a direct movement toward the chosen option. Mean paths for the same subject are shown in the right panel with solid lines. Error bands denote standard error in both the x and y directions.

3.3.2 Theory

Before analyzing the mouse trajectories to address the main questions of the study, it is useful to look in more detail at a simple example of the class of sequential integrator frameworks that motivate this experiment. In these models, decisions are made by dynamically integrating instantaneous noisy samples of the true relative value different between the two options. In particular, consider the following simple extension of the Drift Diffusion Model (DDM) (Milosavljevic et al., 2010; Ratcliff & McKoon, 1989, 2008; Ratcliff & Rouder, 1998; Ratcliff & Smith, 2004; Ratcliff & Tuerlinckx, 2002). Decisions are made by computing a relative value signal (RVS) that estimates the relative value of the options associated with the left and right response boxes. A choice is made the first time the RDS crosses one of two pre-established barriers: one at +*B* associated with the left choice, and one at –*B* associated with the right choice. The RVS starts every trial at zero and then it changes in discrete time as follows:

$$RVS_t = RVS_{t-1} + \delta_{Now}(SV^L_{Now} - SV^R_{Now}) + \delta_{Del}\left(SV^R_{Del} - SV^L_{Del}\right) + \varepsilon_t$$

where ε_t denotes i.i.d. white Gaussian noise with a constant SD of σ, SV^L_{Now} and SV^R_{Now} denote the subjective associated with the immediate payment for the left and right options, SV^L_{Del} and

SV_{Del}^{R} denote the subjective value associated with the delayed payments for the left and right options, and δ_{Now} and δ_{Del} are positive weights that modulate the contribution of immediate and delayed rewards to the RDS.

In simple versions of the DDM, the weights δ_{Now} and δ_{Del} do not differ by attribute type, which implies that controlling for their magnitude, the immediate and delayed components of value have an identical impact in the decision process. Relaxing this assumption has profound implications for intertemporal choice, measured discount rates, and individual differences. In particular, suppose that δ_{Now} changes over the course of the decision, so that it equals zero for the first NDTI ms, and it equals a constant value *d* afterwards. Similarly, suppose that δ_{Del} changes over the course of the decision, so that it equals zero for the first NDTD ms, and it equals the same constant value *d* afterwards. Here, NDTI refers to non-decision time for the immediate rewards, and NDTD refers to non-decision time for the delayed rewards. Then, the probability of making a patient choice increases as NDTI increases, and decreases with as NDTD decreases (see the figure legend for details).

The intuition for this result is easily illustrated using a simple but robust set of simulations, depicted in Figure 3. All of the simulations assume that d = 0.01, σ = 0.01, and B = 1. These parameters are similar to those identified in previous work which has fitted the DDM to a variety of tasks with similar psychometric properties (e.g. Ratcliff, 1980). In each trial, subjects choose between an immediate option with a constant value subjective value of 1, and an option involving only a delayed payment, with a subjective value between 0 and 2. Each potential choice was simulated 1000 times for each combination of parameters.

Fig. 3A depicts a first of simulations in which NDTD was fixed to 500 ms and NDTI ranged from 0 to 1500 ms. The impact of delaying the processing of immediate rewards (i.e., increasing NDTI) is clear cut: the choice curve shifts to the left and individuals become more likely to choose the combo option and delay gratification. The forces at work are quite intuitive. Information about immediate reward favors the immediate option versus the combo, which includes a delayed payment. As NDTI increases, less information about immediate payoffs is integrated into the RVS, which decreases their influence on choices. The impact of changing NDTI on RTs was more complex.

Fig. 3B depicts the results of the opposite exercise: NDTI was fixed to 500 ms and NDTD ranged from 0 to 1500 ms. The results show that increasing NDTD has the opposite effect: it shifts the choice curve to the right, so that individuals become less likely to choose the combo option and delay gratification. Again, the changes on RTs were more complex.

The results of these two simulations provide a theoretical foundation for the hypotheses tested below; namely, that slower processing of immediate rewards and faster processing of delayed rewards increase the likelihood of making a patient choice. Interestingly, since the variation in NDTs explored in the simulation is a small fraction of the mean RTs, the simulations also show that small changes in NDTs can have a very sizable impact on intertemporal choices.

Importantly, since the underlying discount rate used to compute the subjective values is not changing across conditions, the simulations also show that the increases in NDTI lower the estimated discounting parameter, and that increases in NDTD increase the estimated discount rate. As a result, the subject looks more impulsive or myopic, not because the valuation system is increasing the rate at which it discount future rewards, but because this changes the relative weight that the two types of rewards received in the DDM algorithm.

At first sight, the results depicted in Figure 3 also suggest that changes in NDTI shift the psychometric choice curve, whereas changes in NDTD rotate it. This particular difference is an artifact of the simulation details, and can be easily changed by changing the simulation parameters.

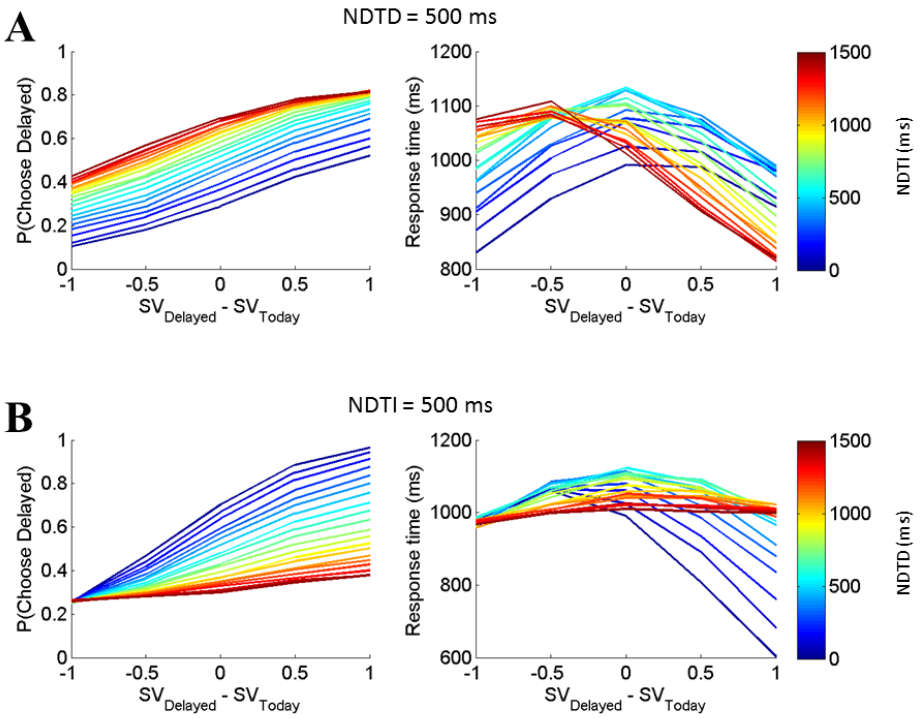

Figure 3. Simulated influence of non-decision time on choice. (A) Choices were simulated using a modified Drift Diffusion Model that varied the immediate value non-decision time (NDTI). Delayed value non-decision time (NDTD) was fixed to 500 ms. The left panel displays a psychometric curve for the simulated choices as a function of the difference in the today and delayed values. The right panel shows the influence of changing immediate value processing speed on reaction time. The color bar denotes NDTI values. (B) Choices simulated while varying NDTD and fixing NDTI to 500 ms. The color bar denotes NDTD values.

3.3.3 Average processing speed for immediate and delayed rewards

We next carried out several analyses of the mouse-tracking data to address the questions of the study. The first set of analyses where design to test if, on average, immediate rewards are processed faster than delayed ones.

To do this, we needed to measure how the influence of immediate and delayed values on the cursor trajectories changes over the course of the trial. In particular, we are interested in identifying the earliest time at which immediate and delayed reward information begins to affect the choice process, as captured by changes in the mouse trajectories. For example, if delayed rewards affect

the trajectories from time *T* until the final choice was made, but not earlier, this would provide evidence that this type of value information is not processed by the decision making circuitry before this time. This method was developed in a previous study, were we used it to show that the relative times at which taste and health begin to influence the choice process explain a large amount of individual variation in dietary self-control (Sullivan et al., 2015). (See also Dshemuchadse et al., 2013; Scherbaum et al., 2010; Scherbaum et al., 2013 for related approaches)

The analysis was carried out in two steps (see Methods for details). First, for every subject, we estimated a series of linear regressions, one for each 10 ms time bin, designed to investigate how the influence of immediate and delayed rewards on mouse trajectories evolves over the course of a decision. In each regression, the dependent variable was the angle of the cursor trajectory at each time bin, which provides a measure of the direction of movement, with -45^o always representing a direct movement towards the left option, 0^o always indicating a movement straight upwards, and +45^o always indicates a direct movement towards the right option. The independent variables were a constant, the relative value of the immediate payments, given by $SV_{Now}^{R} - SV_{Now}^{L}$, and the relative value of the delayed payments, given by $SV_{Del}^{R} - SV_{Del}^{L}$, with R and L denoting the right and left locations. Critically, the SVs used in these regressions were computed using the mean *k* and α best fitting parameters in the group, and thus any differences in the trajectories across subjects or trials cannot be attributed to differences in how the SVs are computed. Figure 4A depicts the time courses of the mean estimated coefficients, and their SE across subjects. Note that these paths suggest that neither type of reward has a significant effect on the mouse trajectory before 700 – 800 ms into the decision.

Second, we used the estimated coefficients from these regressions to compute subject specific measures of the earliest time at which immediate and delayed payoffs exhibit a significant influence on the mouse trajectories that lasted until the end of the trial. We refer to these measures as non-decision time (NDT), and estimate one for immediate rewards (NDTI) and one for delayed rewards (NDTD) for each subject. See the Methods and Figures 4B-C for a description of how the NDTs are computed. Figure 4D depicts the difference between these coefficients in the group: the two NDTs are not significantly different (NDTI: mean = 779, SD = 625; NDTD: mean = 847, SD = 564; difference: mean = -68, SD = 516; t(30)=-0.74, p=0.46), that the two decision times are highly correlated (linear regression: R^2= .39, slope=.57, p=1.7e-4), and that there are sizable differences across subjects.

Together, these results suggest that on average immediate and delayed rewards are processed at similar speeds. This result is important because, as shown in the theory above, it implies that the fact that the average discount factor in the population is positive and sizable cannot be attributed, at least in part, to a systematic difference in the speed with which both attributes are processed by the comparator system implementing DDM type algorithms.

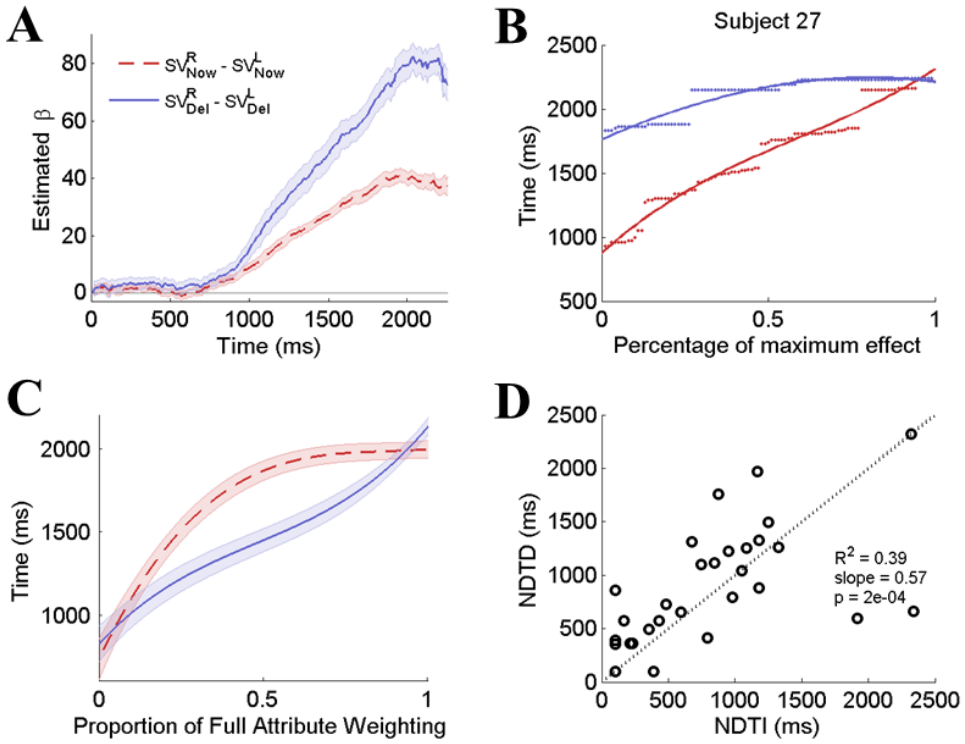

Figure 4. Effect of immediate and delayed value on mouse trajectories. (A) At each individual time sample, a linear regression was used to predict mouse angle using the relative (left – right) value of immediate and delayed subjective value. The coefficients of these attributes for each time point are displayed on the y-axis, and time is displayed on the x-axis from time point zero to a cut-off of 2250 ms. Error bands denote standard errors. (B) Each dot represents the time, on the y-axis, at which each attribute's estimated coefficient reaches a certain percentage of its maximum effect, on the x-axis. Solid lines represent the fitted cubic regression. The y-axis intercept of this line is used as an estimate of when an attribute begins to accumulate a proportion of its maximum value. (C) The time, on the y axis, at which each attribute's influence on mouse angle had accumulated a proportion of its maximum value. Solid lines represent the average fitted cubic function, and shading is their standard error. The y-axis dots indicate the average intercept of this function, and error bands are their standard error. (D) Delayed and immediate value significance times plotted against each other. The dotted

line represents a perfect correlation line, and solid line represents a linear regression fit line.

Differences across subjects. The second set of analyses was designed to test if some of the sizable individual differences in discount rates can be attributed, at least in part, to differences in the speed with which individuals process immediate and delayed rewards. To do this, we estimated a linear regression of the percentage of patience subjects on the NDTI and NDTD measures, and found that it could explain 26% of the individual differences. Interestingly, most of the effect was driven by differences in the processing speed of the delayed rewards (NDTD: slope = -0.0001, p=0.015; NDTI: slope = 7.1e-6, p=0.81; see also Figure 5, which depicts the univariate correlations and regression results). Together, these results suggest that about a quarter of individual differences in discount rates can be explained by differences on the speed at which the delayed rewards are processed.

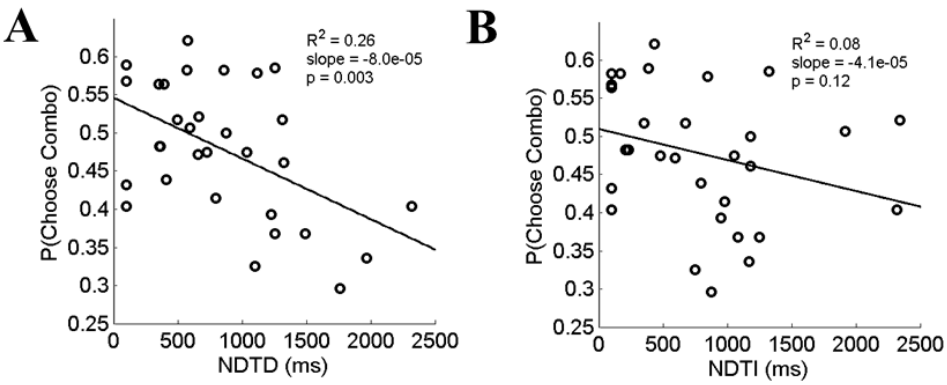

Figure 5. Effect of immediate and delayed value NDTs on choice. A) Delayed value significance time plotted against the percentage of patient combination offer choices by each subject. The solid line represents a least squares linear regression fit. B) Immediate value significance time plotted against the percentage of patient combination offer choices by each subject. The solid line represents a least squares linear regression fit.

Differences across trials. Finally, we investigated if random fluctuations of the relative speed at which immediate and delayed rewards are computed could induce changes in the ability to postpone gratification within subjects. To test this, we compare the distribution of NDTs in trials with patient and immediate choices. As shown in Figure 6, the processing speeds differed substantially across both types of trials. On patient trials, delayed rewards were computed substantially faster than

immediate rewards (NDTD: mean = 826, SD = 463; NDTI: mean = 1658, SD = 339; difference: mean = 832, SD = 545; t(29)=8.37, p=3e-9), whereas the opposite pattern was observed in impulsive trials (NDTD: mean = 2079, SD = 116; NDTI: mean = 1187, SD = 667; difference: mean = -892, SD = 684; t(30)=-7.26, p=4e-8). In addition, the average NDT for patient trials (1242 ms) was significantly earlier than the average NDT of impatient trials (1633 ms; t(29)=-4.45, p=0.0001). Interestingly, the fluctuation in NDTDs was much larger than the fluctuation of NDTIs (1254 ms vs. 499 ms; t(29)=4.03, p=0.0004). As in the previous analysis, this suggests that fluctuations in NDTDs might have a more profound impact on the ability to postpone gratification than fluctuations in NDTIs.

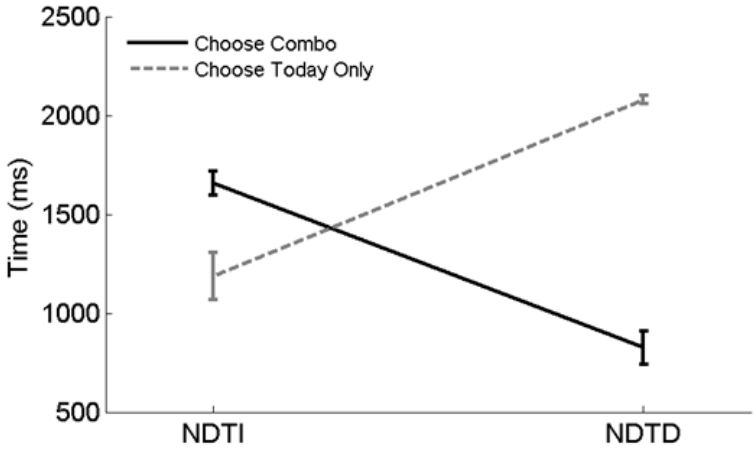

Figure 6. Trial-by-trial variation in significance time. Immediate and delayed value significance times by subject choice. Error bars denote standard errors.

3.4 Discussion

We have found supporting evidence for the hypothesis that changes in the speed with which immediate and delayed rewards are processed by the decision-making circuitry have a sizable impact in the ability to postpone gratification. First, we found that the average processing speed of both types of rewards in a canonical intertemporal monetary choice task was about 800 ms, and not significantly different between types. This result is important because it suggests that the fact that delayed rewards are almost always discounted cannot be attributed to a slower processing speed on average. Second, we found that about 25% of the individual differences in

discounting could be explained by differences on the speed with which delayed rewards were processed. Contrary to our priors, we did not find a similar effect for immediate rewards. Third, we found sizable differences on the relative processing speed of immediate and delayed rewards across trials: subjects processed delayed rewards about 800 ms faster than immediate rewards when they made patient choices, but the opposite was observed during trials that resulted in impulsive choices. Together, these results show that a sizable fraction of variation in the ability to postpone gratification might be attributable to individual variables that affect the speed at which different types of rewards are processed, and not to differences in deep preference parameters like the temporal discount rate used by the brain's valuation systems.

Several aspects of the results are worth highlighting.

First, since multiple alternative mechanisms are likely to affect the ability to delay gratification, we were surprised by the sizable amount of variation explained the differences in processing speeds. To put these results in perspective, previous studies have found that personality traits such as impulsivity or IQ only explain around 13% and 5%, respectively, of individual differences in self-control (de Wit et al.; Shamosh & Gray, 2008; Steinberg et al., 2009). Understanding the relationship between individual traits like impulsivity and IQ is an important open question for future research. However, it is important to emphasize that these types of variables cannot entirely account for our results even if they were highly correlated with processing speeds, because they explain a much smaller fraction of the individual differences and because they cannot account for within-subject changes across trials.

Second, the comparison of processing speeds across trials suggests that they can change rapidly, often by a factor of 100%. Although our design cannot address the source of these fluctuations, it is natural to speculate that they result from fluctuations on how relative attention or cognitive effort change across trials. Identifying these sources is an important open question for future research. Regarding the exact mechanisms at work, our data suggests that processing speeds might be highly responsive to changes in contextual variables, such as how information is presented, which suggests that it might be possible to increase self-control using simple "nudge" style interventions (Camerer, Issacharoff, Loewenstein, O'donoghue, & Rabin, 2003; Harris et al., 2013; Thaler & Benartzi, 2004; Thaler & Sunstein, 2003). These interventions would affect self-control by changing the speed at which delayed information is processed, for example by re-

directing attention to these types of rewards (T. A. Hare et al., 2011). In fact, the association between attention and contextual effects has also been used to explain preference reversals (Busemeyer & Townsend, 1993; Kim, Seligman, & Kable, 2012; Roe et al., 2001), and context effects in risky choice (Johnson, Schulte-Mecklenbeck, & Willemsen, 2008).

Third, the results suggest that changes in the processing speed of delayed rewards are more much more important than changes in the processing speed for immediate rewards. Consistent with this, the processing speed of delayed rewards fluctuate across trials about 250% more than for immediate rewards. We were surprised by this finding since the simple versions of the DDM that motivated our hypothesis predict that the processing speed of both types of rewards should affect self-control: delaying the processing of future rewards hurts, whereas delaying the processing of immediate rewards helps. A potential post hoc explanation is based on the work of Liberman and Trope (2008), which has argued that more abstract rewards are more difficult to process. Based on this, if immediate rewards are easier to process than delayed rewards, they might be less sensitive to changes in processing speed and duration than immediate rewards. Adding this force to the theoretical framework proposed here would then produce the observed asymmetry. We emphasize, however, that this is a post hoc and untested explanation.

Fourth, the mechanism proposed here follows directly from adding differences in processing speeds to any of the sequential integrator algorithmic models of choice that are popular in the literature, such as such as the Drift-Diffusion Model (Ratcliff, 1978, 1980; Ratcliff et al., 2003; Ratcliff & McKoon, 1982; Ratcliff & McKoon, 2008; Ratcliff & Rouder, 1998; Ratcliff & Smith, 2004; Ratcliff & Tuerlinckx, 2002), the attentional Drift Diffusion-Model (Krajbich et al., 2010; Krajbich et al., 2012; Krajbich & Rangel, 2011), the leaky-accumulator model (Usher & McClelland, 2001), and Decision Field Theory (Busemeyer & Diederich, 2002; Busemeyer & Townsend, 1993; Roe et al., 2001). This is important because their plausibility as an algorithmic description of the choice process is supported by a large and rapidly growing body of behavioral and neural data (Basten, Biele, Heekeren, & Fiebach, 2010; Britten, Shadlen, Newsome, & Movshon, 1993; Gold & Shadlen, 2000; Hare, Schultz, Camerer, O'Doherty, & Rangel, 2011; Jocham, Hunt, Near, & Behrens, 2012; Mulder, Wagenmakers, Ratcliff, Boekel, & Forstmann, 2012; Park, Kahnt, Rieskamp, & Heekeren, 2011; Ratcliff et al., 2003). All of these models could make a similar qualitative prediction: slowing down the speed at which an attribute is processed reduces its influence in the choice that is made. This idea has been previously used to explain

why the speed of information retrieval can vary by information type (Ratcliff & McKoon, 1989) and to explain differences in the ability to exercise dietary self-control (Sullivan et al., 2015). The main contribution of this study is to show that this mechanism plays a sizable role in basic intertemporal choice, and thus is possibly at work in all of the domains in which individuals have to choose between sooner-smaller and later-larger rewards.

Fifth, the mechanism identified here is connected to several other mechanisms that have been shown to influence self-control. Weber et al. (2007) proposed that some of the asymmetric treatment of immediate and delayed rewards is due to biases that make memories associated with immediate rewards easy to process (Peters & Buchel, 2010). Although our study is silent on this point, these memory biases could be responsible for some of the differences in processing speeds. Both the attentional DDM (Krajbich et al., 2010; Krajbich et al., 2012; Krajbich & Rangel, 2011) and Decision Field Theory (Busemeyer & Diederich, 2002; Busemeyer & Townsend, 1993; Roe et al., 2001) predict that increasing the relative amount of time devoted to processing delayed rewards would increase patience. This is connected to the mechanism studied here since increasing the amount of time that it takes to begin processing the delayed rewards affects the overall share of attention that they receive. The key difference with this previous work is that here, all of the variation in processing time comes at the beginning of the trial, whereas in those previous studies it results from fluctuations in attention at any point in the decision process, such as those resulting from changes in which attribute is fixated on at any given point. Cognitive effort and cognitive load have been shown to be powerful modulators of self-control (Baumeister & Vohs, 2004; Muraven & Baumeister, 2000). These effects might operate, at least in part, by changing the speed at which the delayed rewards are processed. Consistent with this, several recent studies have shown that the extent to which the stimulus value signals in the ventromedial prefrontal cortex, which are widely thought to guide choices (Bartra, McGuire, & Kable, 2013; Clithero & Rangel, 2014), reflects self-control is modulated by activity on areas of dorsolateral pre-frontal cortex that are known to be involved in cognitive control (Figner et al., 2010; T. Hare et al., 2009; Hare et al., 2014; Harris et al., 2013).

Finally, in order to test the type of hypotheses proposed here, it is critical to have a sufficiently good measure of how the decision process evolves before a choice is made. To do this, we borrow heavily from the pioneering mouse-tracking literature (Dale et al., 2007; Lohse & Johnson, 1996; McKinstry et al., 2008; Scherbaum et al., 2010), which has been used to study

visual search (Song & Nakayama, 2008), stereotyping (Freeman & Ambady, 2009), psycholinguistics (Spivey et al., 2005), and other related applications on intertemporal choice (Dshemuchadse et al., 2013). However, to accomplish our goals we used the mouse-tracking data in a different way, originally developed in Sullivan et al. (2015). Whereas most mouse-tracking studies focus on comparisons between the average paths generated by various experimental conditions, we use regression analyses to identify the time at which different stimulus attributes begin influencing choice. It is natural to speculate that the large individual and contextual differences observed in many other canonical problems, such as risky or altruistic choice, might also be partly attributable to variation in the processing speed of the associated attributes. For example, more risk averse individuals might process information about risk and/or losses faster than their less risk averse counterparts. Applying the conceptual framework and methods developed here to these other domains is an exciting avenue for future research.

Chapter 4

NEURAL CHANGES ACROSS THE LIFESPAN ARE ASSOCIATED WITH DECREASED TEMPORAL DISCOUNTING[4]

4.1 Introduction

Extensive research suggests that individuals discount the value of future rewards as a function of their delay. Moreover, increasing evidence suggests that these deficits in self-control have large, long-term consequences for both the individual and for society (Boyle et al., 2013; Forstmeier et al., 2011; Mischel et al., 1988). Given the importance of such decisions, there is great interest in understanding the sources underlying an individual's ability to delay gratification.

Several tasks are often used to study trade-offs between immediate and larger future rewards. One classic paradigm is intertemporal choice, in which subjects must decide between a smaller amount of money now, and a larger amount later (Ainslie, 1975; Berns et al., 2007; Green et al., 1994; Kable & Glimcher, 2007; Laibson, 1997; Loewenstein & Prelec, 1992). Another commonly used paradigm to assess decision making during delayed gratification is dietary choice. For example, choosing between a snack of ice cream and celery presents a trade-off between eating something delicious now and having better health in the future.

Studies using these paradigms suggest that a concerted valuation system, including the ventromedial prefrontal cortex (vmPFC) and ventral striatum (vStr), appear to encode the value of choice options. For example, in intertemporal choice, the vmPFC and vStr increase in response in accordance with increased in the value of the delayed option (Kable & Glimcher, 2007). In dietary choice, the vmPFC encodes the tastiness of food options in good and poor dieters alike (T. Hare et al., 2009). This valuation system appears to extend to a wide range of other decision contexts as well (Bartra et al., 2013; Clithero & Rangel, 2014), and may translate different value types into a "common currency" to facilitate their comparison (Kahnt, Heinzle, Park, & Haynes, 2011; McNamee et al., 2013).

[4] In collaboration with Mara Mather, Adriana Galvan, and Antonio Rangel.

In intertemporal choice, self-control is required to bring one's valuations more in line with larger future, rather than smaller immediate, rewards. Converging evidence from many studies suggests that the left dorsolateral prefrontal cortex (L DLPFC) enables long-term focused choices by modulating the value of choice options, making them more reflective of long-term benefits. In one experiment, in which subjects were choosing between smaller, sooner monetary rewards and larger later ones, the L DLPFC increased in effective connectivity to the vmPFC at the time of choice. Its connectivity further increased when a delayed choice was made (Hare et al., 2014). Additionally, individuals with better dietary self-control exhibited connectivity between the L DLPFC and vmPFC when self-control was required to, for example, say no to a tempting food option in favor of a helathier one. Worse dieters did not (T. Hare et al., 2009). Evidence further suggests that connectivity between the vmPFC and L DLPFC can be modified; simply instructing participants to focus on the health of foods encouraged increased effective connectivity between these regions (T. A. Hare et al., 2011). This is also a causal link between the L DLPFC and impulsivity in intertemporal choice; inhibition of this region using transcranial magnetic stimulation results in more impatient choices (Figner et al., 2010). Together, this research suggests a concerted system in which the L DLPFC enables the vmPFC to encode not only the value of immediate, but also more long-term consequences of a choice.

To pinpoint the contribution of this system to self-control behavior, it is useful to measure how changes in this system's structure, function, and connectivity are associated with changes in self-control. Conveniently, one important phenomenon surrounding self-control is that the incidence of many pathological behaviors related to self-control deficits decrease with age (Arnett, 1992; Green et al., 1994; Gullone & Moore, 2000; National Institute on Alcohol Abuse and Alcoholism, 2008; Read & Read, 2004; Sangrock, 2002; Spear, 2009). In this study, we capitalize on this phenomenon to estimate how changes in the brain's valuation and control systems change to allow for better choices.

Throughout the lifespan, distinct, systematic changes occur in the regions of the brain previously associated with reward and successful self-control discussed above (Samanez-Larkin & Knutson, 2015). For example, vStr response to monetary reward dramatically decreases from adolescence into adulthood (Galvan et al., 2006). Similarly, stimulus value signals in the amygdala are down-modulated with age (Mather et al., 2004). Increased effective connectivity between the vmPFC and DLPFC has been seen from younger to older children, and is related to increased ability to resist

temptation (Steinbeis, Haushofer, Fehr, & Singer, 2014). Dramatic changes in the structural composition of the brain also occur from childhood through old age; the frontal lobe isn't fully developed until adulthood, and lateral brain regions including the DLPFC are particularly compromised in elder years (Sowell et al., 2003).

Together, this research demonstrates that behavioral and neural changes occur in tandem throughout the lifecycle, suggesting improvements in the brain's self-regulation network. These age-related neural changes lead to a better assignment of value to tempting options by representing long-term goals more strongly, leading to more patient choices. We test this hypothesis by measuring brain activity while subjects from 13 to 70 make choices between smaller-sooner and larger-later rewards in an intertemporal choice task. In doing so, we hope to pinpoint the changes in information processing responsible for the improvements in self-control with age.

Specifically, we expect hypersensitivity to reward in our adolescent subjects' valuation system, which we expect to decrease considerably with age. This decreased value sensitivity may be related to a decreased need to exert control, resulting in more patient choices. We also predict increases in effective connectivity between valuation and control regions with age to further enable the evidenced increased ability to delay gratification.

4.2 Methods

Subjects. Fifty-one subjects spanning three age groups were recruited for this study. The adolescent age group included 13-16 year olds (mean 14.9 years old, N=18), the middle group included 30-44 year olds (mean 35.0 years old, N=20), and the older group included 65-70 years olds (mean 66.4 years old, N=13). Adolescent subjects were recruited through co-author Adriana Galvan's UCLA adolescent subject database. Middle group subjects were recruited through co-author Mara Mather's USC Healthy Minds subject database and Caltech's BrainScience recruitment website. Older subjects were recruited through the USC Healthy Minds recruitment website. Seven subjects are not included in the total subject count due to excessive head movement, inability to estimate discount parameters, early withdraw, or equipment failure.

Screening. All subjects were screened for MRI contraindications, psychiatric medication use, neurological conditions, and brain injury. Participants in the older group were screened for cognitive impairment and dementia using the Telephone Interview of Cognitive Status (de Jager,

Budge, & Clarke, 2003), which has been validated to have high sensitivity and specificity even in educationally and ethnically diverse populations (Manly et al., 2011). Adolescents were screened for impulsivity and developmental disorders.

Task 1: Discount parameter estimation task. Participants' hyperbolic discounting parameters were estimated using a Bayesian adaptive design (see "Behavioral data analysis" for details on the hyperbolic model). Participants made a series of choices between receiving one amount within a certain number of days, and another amount some time later. Amounts varied between $5 and $100 in $5 increments, and delays varied between 0 and 200 days in 4 day increments. The two options were presented on the left and right sides of the screen, with the more delayed choice randomly on the right or left. Participants responded using the keyboard. The program began with a flat prior across discounting parameters (k, α, and ω), and updated with each response. The task ended when estimated parameters confidence exceeded 85%, or when 15 minutes had expired. The discount parameter k estimated from this procedure was highly correlated with the parameter estimated using fMRI task choice data collected using the task described below (r=0.87, p=1.8e-19).

Task 2: Intertemporal choice task. In the scanner, participants made 300 binary choices between $25 the day of the experiment and another amount some time later (Fig. 1A). Amount and delay pairs ranged from $25 to $54, and 7 to 200 days. The pre-estimated discount parameters were used to generate a choice set in which the subject was likely to accept the delayed offer to 100 questions and the today offer to 100 questions. The remaining 100 questions were within $1 of the subject's point of indifference. Since the immediate reference option was constant, only the delayed choice option was displayed on the screen. Participants had three seconds to responded "yes" or "no" to the delayed option using their index and middle fingers. After responding, participants saw a .5 s reflection of their choice ("yes" or "no") in white. A red "x" was displayed if no response was received. Subjects then saw a white fixation cross of jittered duration (2 to 6 s, uniform random). Any difference in reaction time and the three-second choice time limit was added to the fixation cross display time. The 300 trials were randomly ordered and spread over three EPI runs. To avoid confounds related to the hour of testing, which might affect the age groups differently, half of each group was randomly assigned to either a morning or an afternoon fMRI session.

After the task, participants were paid $40 per hour and a $50 travel honorarium in cash. One trial was randomly selected for an additional payment, which was sent via PayPal either that day or in

the future, depending on participant choice. Participants unfamiliar with PayPal were guided through signing up, accepting, and retrieving funds and received a detailed instruction sheet to take home that included the delivery date and amount to be delivered.

Behavioral data analysis. Each subject's choices were fit to a hyperbolic model of discounting. This model assumes that subjects have linearly additive preferences over prizes received at different dates, and that the value of each such prize is given by

$$SV = \frac{X^{\alpha}}{1 + kD}$$

where X denotes the size of the prize, D represents the delay in days, *k* is a hyperbolic discounting parameter, and α represents the curvature of the utility function. Note that the SV of the constant reference point, denoted by SVT, equals \$25α.

As is standard in the literature (Chabris et al., 2009; Pawitan, 2001; Rodriguez et al., 2014), the model also assumes that subjects make probabilistic choices, with the probability of choosing the on-screen option given by

$$\frac{e^{\omega SV_D}}{e^{\omega SV_D} + e^{\omega SV_T}}$$

where ω is a non-negative parameter determining the influence of the SVs on the choices.

For each subject, we estimated α, *k,* and ω parameters that best fitted the choice data by maximizing the likelihood function calculated using this model.

In a further step, we calculated a measure of the area under the discounting curve (AUC), which has been suggested to be free of bias associated with imposing a hyperbolic model (Myerson, Green, & Warusawitharana, 2001). To calculate the AUC, we normalized the delay and subjective value for each data point and each subject. These values were plotted as x and y coordinates, and the area under the resulting curve was calculated and used as a secondary measure of delayed discounting.

Image acquisition. Data were collected using a Siemens 3.0 Tesla Trio MRI scanner and a 32 channel, phased array head coil. 285 echoplanar (EPI) volumes with blood oxygenation level-

dependent (BOLD) contrast were collected during the task with 2500 ms TR, 30 ms TE, 85º flip angle, 192 mm FOV, in-plane resolution of 3 mm x 3 mm, and 47 3 mm slices with a 0.3 mm gap using ascending acquisition were acquired. High-resolution T1-weighted anatomical images were also collected and co-registered with the EPI images to facilitate anatomical localizations. Phase and magnitude fieldmaps (3mm isotropic) were collected in-between EPI sessions.

Image pre-processing. Image data were pre-processed using FSL (Jenkinson, Beckmann, Behrens, Woolrich, & Smith, 2012). The T1 structural image was skull-stripped to facilitate registration to functional data, and was non-linearly registered to the Montreal Neurological Institute (MNI) template. Because the scanner began recording volumes after five TRs, only the first two recorded EPI volumes thereafter were deleted to account for saturation effects. Phase and magnitude fieldmaps were used to correct for distortions in the EPI data. Data were corrected for motion and realigned to the mean image. Realignments and motion outliers were collected and used as confounds in later analyses. EPI data were slice-time corrected, intensity normalized, spatially normalized to the MNI template, and spatially smoothed using a Gaussian kernel (FWHM = 6 mm) to reduce noise. Data were also temporally filtered using a filter width of 100 s to remove low frequency trends.

fMRI statistical analyses. All fMRI analyses presented here were performed using FSL (Jenkinson et al., 2012). GLM 1 estimated BOLD response from trial onset until a choice was made. There were two parametric modulators to this model: first, the demeaned subject-specific subjective value of the delayed option, and then an indicator for delayed choice (1 when a delayed choice is made, 0 otherwise). The delayed choice modulator was orthogonalized with respect to the SV modulator. GLM 2 was identical to GLM 1 except that SVs were calculated using fixed discounting and utility parameters for each subject. GLM 3 was identical to GLM 2, but separated trials into “hard” trials, which are within $1 of the subject’s indifference point, and all other trials.

Psychophysiological interaction 1 (PPI 1) used the significant L DLPFC cluster ($p<.001$ threshold) from the delayed choice indicator in GLM 1 as its region of interest (ROI). The ROI time course was extracted and used as the ROI seed. A subjective value modulator (calculate using fixed parameters) and an indicator for delayed choice were added to this model. Reported analyses are generated from the interaction between delayed choice and the physiological variable. We then ran an additional PPI model (PPI 2) to estimate regions with greater connectivity during delayed choice

when self-control was required to select the delayed option. We separated trials in to hard and easy according to the procedure for GLM 3. Otherwise, this PPI was identical to PPI 1.

Volume segmentation and voxel-based morphology. We used FAST (Zhang, Brady, & Smith, 2001) to segment T1 images into white matter, grey matter, and cerebrospinal fluid while correcting for bias field inhomogeneities. The resulting partial-volume estimates were nonlinearly warped to MNI space to facilitate cross-subject comparison, and then both whole-brain and ROI estimates of tissue composition were extracted.

Self-report and IQ measures. Subjects were administered the two-subtest Wechsler Abbreviated Scales of Intelligence (Wechsler, 2011) which took about 15 minutes. Subjects also filled out demographic and personality questionnaires: the Cornell Medical Index (Brodman et al., 1951), Beck Depression Inventory (Beck, Ward, Mendelson, Mock, & Erbaugh, 1961), BIS/BAS (Carver & White, 1994), PANAS (Thompson, 2007), STAI (Spielberger, Gorssuch, Lushene, Vagg, & Jacobs, 1983), and Peterson's Pubertal Development Scale (adolescents only; Petersen, Crockett, Richards, & Boxer, 1988)

4.3 Results

In this study (Fig 1A), we asked how age-related brain changes allow an individual increased ability to exert self-control. In doing so, we hope to understand how brain changes can facilitate older subjects' ability to bring valuations more in line with larger long-term, rather than short-term, rewards. To answer this question, we presented adolescent, adult, and older subjects with intertemporal choices between $25 that day, and another larger amount in the future, while recording their brain activity with fMRI.

4.3.1 Behavioral Results

Estimated Model Parameters. We estimated subject-specific discounting parameters using the widely used hyperbolic model (see Methods for estimation procedure), which captures two phenomena seen in human behavior: discounting of future rewards and decreased discounting with increased reward size (Ainslie, 1992; Laibson, 1997; Mazur, 1987; Pine et al., 2009). This model has been experimentally validated in a number of studies (Frederick et al., 2002; Green & Myerson, 2004; McKerchar et al., 2009). A key element of this model is the discounting parameter *k*, which

provides a measure of the present bias in subject's choices, with larger values indicating a stronger tendency to discount the future. The estimated discount parameters (k=0.011, SD=0.01; α=0.687, SD=0.10) were similar to those found in previous studies (e.g., Chabris, Laibson, Morris, Schuldt, & Taubinsky, 2008; Frederick et al., 2002; Kable & Glimcher, 2007). Later analyses will use log(k) as a measure of individual differences in discounting (mean=-4.9, SD=1.57), which was approximately normally distributed (Fig. 1B).

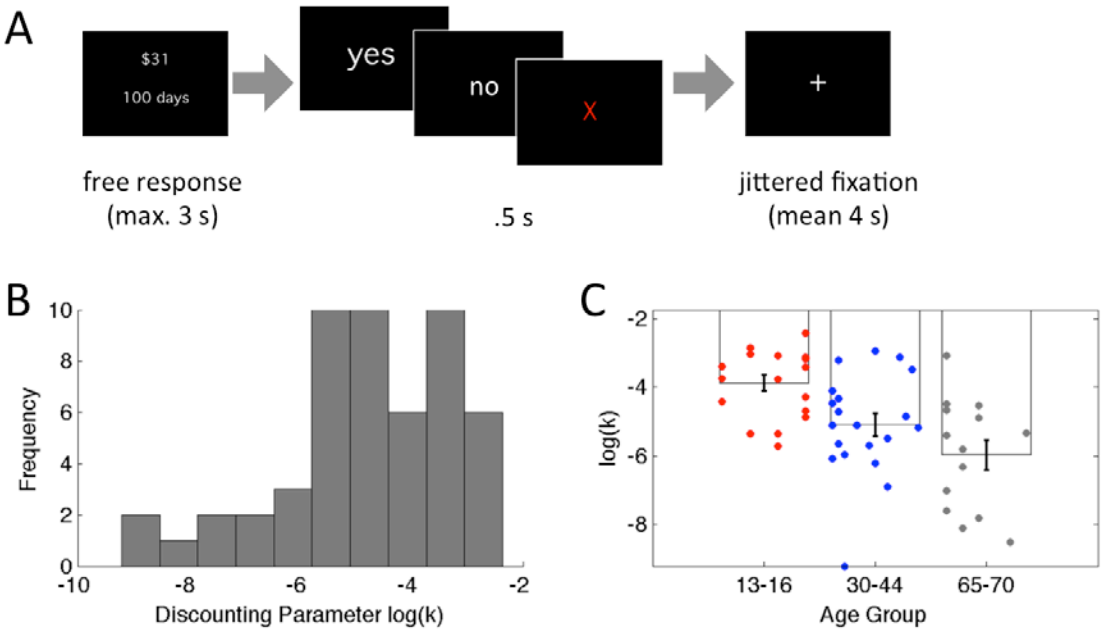

Figure 1. A) Stimuli in fMRI task. B) Distribution of discounting log(k) C) log(k) differs by age.

We compared changes with age, both with in discounting and other variables, in two ways. First, we use the ANOVA to measure mean differences between age groups. Next, to additionally capture changes within each age group, we also model linear changes in activation using linear least squares regressions with chronological age.

Discounting parameter log(k) decreases with age group, suggesting more patience with age (Fig 1C; ANOVA F=9.87, p=0.0002; R^2=0.26 β=-6.61 p=0.0001). However, the direction of change in temporal discounting with age is debated in the literature (Chao, Szrek, Pereira, & Pauly, 2009; Green, Myerson, Lichtman, Rosen, & Fry, 1996; Halfmann, Hedgcock, & Denburg, 2013; Read &

Read, 2004; Samanez-Larkin et al., 2011). Because of this, we validated our finding with an additional estimate of discounting by calculating the area under the discounting function (AUC; see Methods for calculation procedure). This is advantageous because it is free from some of the potential error in estimating parameters based on a hyperbolic discount function (Myerson et al., 2001). Using this AUC measure, we confirm that discounting decreases with age (ANOVA F=9.98, p=0.0002; R^2=0.25 β=42.02 p=0.0002). Additionally, AUC is highly correlated with the discounting parameter *k* (r=-0.88 p=1.3e-17), which justifies using this hyperbolic parameter in subsequent analyses.

Basic Psychometrics. Using the hyperbolic discounting function and each subject's fitted parameters (see Methods) we can estimate the subjective value (SV) of each option in each choice. Figure 2A illustrates that SVs reflect choice behavior (logistic slope greater than zero, mixed-effects T=5.07, p=5.6e-06). There were no differences between age groups in the curve's logistic slope (often seen as a metric of how noisy choices are) (ANOVA F=1.41, p=0.25; R^2=0.03 β=3.25 p=0.22). Similarly, reaction times were greatest when the today and delayed options were equal (Fig. 2B; mean 1284 ms, SD=204 ms; mixed effects T=-11.36, p=1.4e-15). There were no differences between age groups in value difference's influence on reaction time (ANOVA F=0.35, p=0.71; R^2=0.02 β=-281.11 p=0.35).

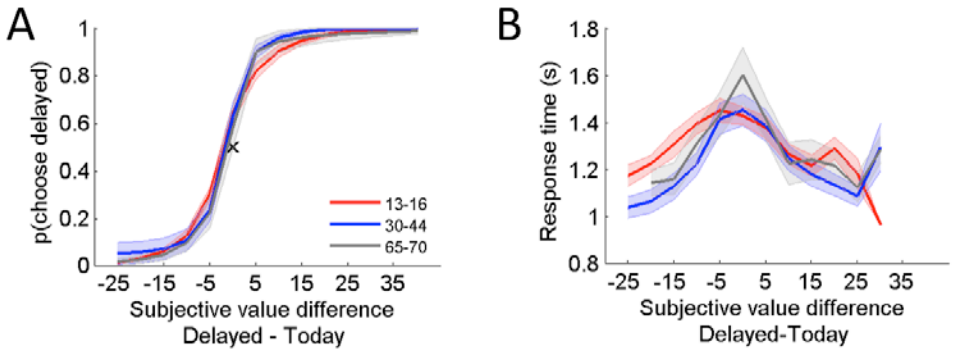

Figure 2. Choice data. A) Psychometric choice curve displaying the frequency of delayed choice selection as a function of the difference in subjective value between the today and delayed options. The "x" marks the frequency of .5 that corresponds to a value difference of 0. B) Reaction time plotted as a function of the difference in subjective value between the today and delayed options.

Demographic and IQ measures. Demographic and questionnaire measures were not related to self-control. Discount parameter log(*k*) was not related to gender (T=1.45, p=.15), household income (ANOVA F=1.26, p=0.30; R^2=0.05 β=0.0001 p=0.22), individual income (excluding adolescents who don't have full time jobs; ANOVA F=1.63, p=0.21; R^2=0.02 β=0.0001 p=0.46), or belief about household financial security (ANOVA F=2.75, p=0.07; R^2=0.07 β=8.5046 p=0.08). In fact, in controlling for these factors age alone remains predictive of discounting parameter log(*k*) (p=.04).

In this study, higher IQs are related to less discounting (R^2=0.22 β=-0.04 p=0.006), replicating previous results (Shamosh & Gray, 2008). IQ was related to age (ANOVA F=2.15, p=0.13; R^2=0.15 β=0.32 p=0.03) which has also been shown in the literature (Fry & Hale, 1996). Using mediation analysis, IQ reduces age's prediction of log(*k*) by 45.1%. A model with both age and IQ (R^2_{adj}=.24, p=.006) explains the same amount of variance in individual log(*k*) values as a model with only age (R^2_{adj}=.24, p=.00001), and more than one with IQ alone (R^2_{adj} =.19 p=.006).

4.3.2 fMRI Results

Localizing valuation regions. The goal of this study was to assess how the decision circuitry associated with valuation and self-control changed with age to facilitate improved choices. Based on overwhelming evidence from an extensive literature, we had a strong prior hypothesis that activity in the ventromedial prefrontal cortex (vmPFC), posterior cingulate cortex (PCC), and ventral striatum (vStr) would positively correlate with increases in subjective value (Bartra et al., 2013; Clithero & Rangel, 2014).

To test this in our data, we estimated a GLM with two parametric modulators: first, the demeaned SV of the delayed option, and then an indicator for delayed choice (GLM 1; see Methods for details). Because, in this analysis, we were interested specifically in localizing regions that increased in activity accordance with individual perceived value of the rewards offered, the SV modulator was calculated using each individual's estimated *k* and α parameters.

Pooling across age groups, activity in regions commonly associated with subjective value (including the vmPFC and vStr) were significantly correlated with value in our task (circled; Fig. 3A). Although the vmPFC cluster does not pass whole-volume cluster correction at a p<.05 threshold, based on our strong initial hypothesis we felt justified in using an ROI-based test of

significance. Doing so, we found that this region survives small volume correction ($p<.05$ threshold). Although this region is more anterior than the one many studies on mPFC response to value, it is actually very well in line with studies specifically on monetary decision making (Clithero & Rangel, 2014; Kable & Glimcher, 2007).

These results suggest that, across subjects, activity in regions previously associated with subjective value are also related to subjective value in our task, even in a subject pool that spans from ages 13 to 70.

Localizing regions active during delayed choice. Based on previous findings discussed in the introduction, across both intertemporal and dietary choice studies we expected that subjects would recruit the L DLPFC when choosing the more patient option. We investigated this by estimating activity related to the delayed choice indicator in GLM 1, pooling across all subjects. After controlling for subjective value modulation, the L DLPFC was more active when choosing the delayed option (Fig. 3B). This region overlaps with the L DLPFC region found when successfully executing dietary self-control (T. Hare et al., 2009), and survives small volume correction ($p<.05$).

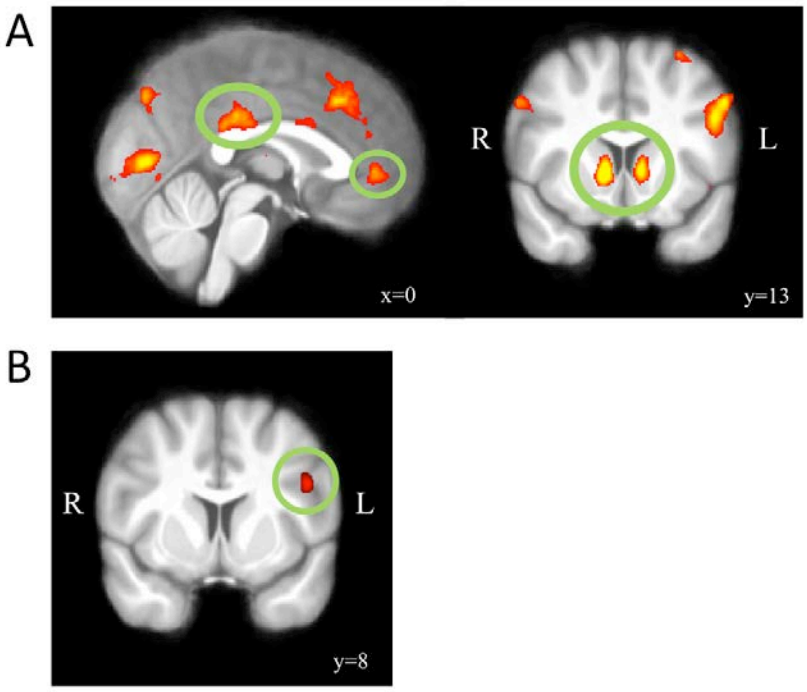

Figure 3. Whole brain results of GLM 1 pooled across age groups. A) Group-level response to subjective value in the brain. vmPFC, PCC, and vStr regions, commonly associated with encoding value, are circled in green. B) Group-level response to delayed choice, with the L DLPFC region previously found to be active when self-control is being exerted (circled). For illustration purposes only, these images are shown at a p<.001 threshold uncorrected. The background anatomical image is the mean structural from this study.

Effective connectivity between control and valuation regions. Previous work has found increased connectivity between the L DLPFC and vmPFC during value-based choices requiring self-control (T. Hare et al., 2009; Hare et al., 2014; T. A. Hare et al., 2011; Steinbeis et al., 2014). Therefore, we expected to find greater connectivity between the L DLPFC and vmPFC during delayed choice in our task. To estimate this in our data, we first created a region of interest (ROI) based on activation during delayed choice in GLM 1 at the p<.001 threshold. We then estimated a psychophysiological interaction (PPI) model using this L DLPFC ROI as the seed region to measure connectivity when the delayed choice was made (PPI 1; see Methods for details).

At the whole brain level, and pooling across subjects, the vmPFC and ACC are positively effectively connected to the L DLPFC when the delayed choice is made (Fig. 4A). Next we

specifically estimated the connectivity between L DLPFC and our value-related regions of interest, the vmPFC and vStr, using a ROIs created from the group-level subjective value modulation from GLM 1 at a p<.001 threshold. We find that both the vmPFC and vStr are significantly effectively connected to the L DLPFC when a delayed choice is made (T=1.97, p=0.03; vStr T=1.78, p=0.04).

Next, we estimate L DLPFC connectivity, specifically when self-control is required. To understand why not all delayed item choices require self-control, imagine that an individual is deciding between $25 today and a very large payment to be delivered after a very short wait (e.g., $500 in 1 day). In this extreme case, selecting the delayed offer is not an accurate representation of a self-control process for most individuals. To measure L DLPFC activity specifically during choices in which self-control was required to select the delayed option, we next estimated L DLPFC connectivity only when the delayed and today options were similar in subjective value for that individual, i.e., near their point of indifference (PPI 2; see Methods for details). Figure 4B displays regions with greater positive connectivity to the L DLPFC specifically in hard trials across all subjects, which included ACC and the hippocampus (Fig. 4B). Neither vmPFC nor vStr are significantly connected to the L DLPFC in this case (p>.09).

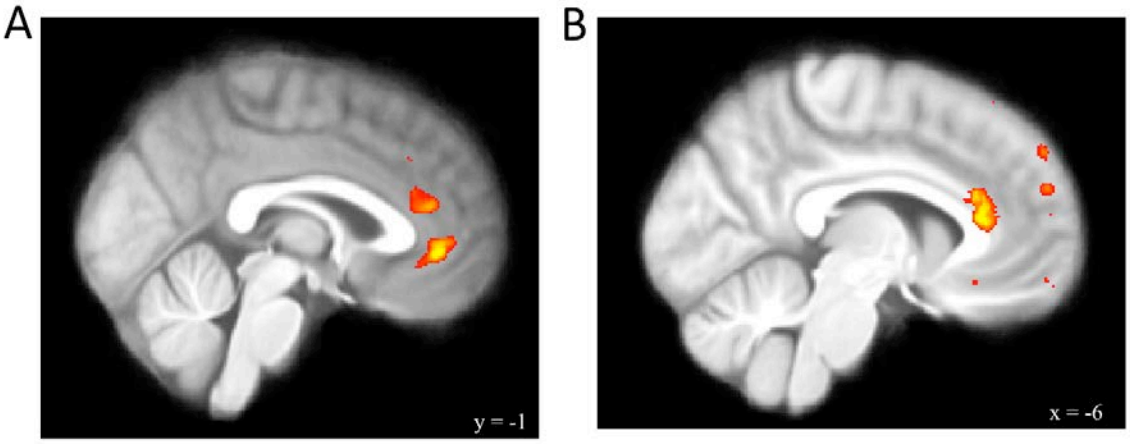

Figure 4. Whole brain estimates of positive connectivity to the L DLPFC during delayed choice. A) Connectivity to L DLPFC when delayed choice is made (PPI 1). B) Connectivity to L DLPFC when the delayed choice is made, in hard, relative to all other, trials (PPI 2). Results shown at p<.001 corrected for multiple comparisons at p<.05. The background anatomical image is the mean structural from this study.

Response to value changes with age. Next, we measure changes in value response across the lifespan. To compare response to value across age groups, we estimated a GLM identical to the one above, but with one important difference. We calculated subjective values using the group-level mean discounting and utility function parameters (GLM 2, see Methods for details; k=0.011, SD=0.01; α=0.687, SD=0.10). This is important because it facilitates comparisons across ages by putting all delayed values (which differ with age due to decreasing impulsivity) on a common scale. Importantly, by using SVs calculated with constant parameters, we still capture the shape of the value function. We used the vmPFC, vStr, and L DLPFC ROIs created for the previous analyses.

Across ages, vmPFC value response is preserved with age (Fig 5A; ANOVA F=0.28, p=0.76; R^2=0.02 β=-0.0009 p=0.32). However, value response in the vStr significantly decreased with age (Fig. 5B; ANOVA F=3.70, p=0.03; R^2=0.13 β=-0.002 p=0.01), suggesting that the vStr response becomes more "patient" with age. This differential change with age in vmPFC versus vStr value modulation occurs despite similar decreases in grey matter density with age in both regions (vmPFC: ANOVA F=19.11, p=7.9e-07; R^2=0.42 β=-0.003 p=2.3e-07; vStr: ANOVA F=5.24, p=0.01; R^2=0.09 β=-0.0009 p=0.03). The relationship between vStr value response and age holds when controlling for percentage of Vstr grey matter, which is not a significant predictor of this relationship (R^2=0.13; age β=-0.0022, p=.02, vStr grey matter β=.04, p=.89). This confirms previous findings that vStr, but not vmPFC, response to value decreases in younger individuals (Galvan et al., 2006), and extends this relationship to an older demographic.

Next we assessed whether specifically vStr value response change with age was a significant driver of impulsivity. First, we found that vStr value response is significantly related to log(k) even when controlling for decreases in vStr grey matter (R^2_{adj}=0.21 β=0.04 p=0.0009). We next performed a mediation analysis. vStr value response was as significant mediator of age's prediction of log(k) value response, reducing age's prediction of log(k) by 25.6%. This suggests that vStr response to value may be the driver of age-related changes in impulsivity.

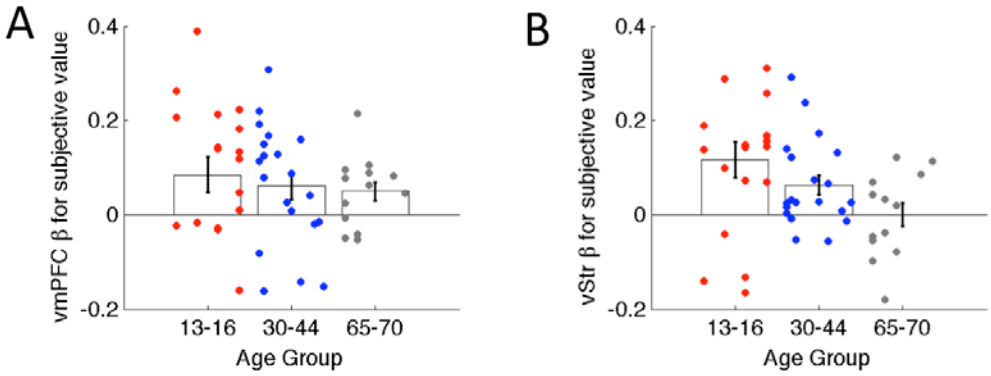

Figure 5. Change in value response with age (GLM 2). A) vmPFC value modulation by age. B) Striatal response to value with age. Bars represent the average response, with error bars representing standard error. Dots represent individual values for subjects (Red = adolescent, blue = middle, grey = older).

L DLPFC activity during delayed choice changes with age. We next investigated age-related changes in L DLPFC response during delayed choice. Previous work has suggested that individuals with better self-control exhibit increased L DLPFC activity during patient choice (Hare et al., 2009). If the increased ability to exert self-control with age is enabled by greater L DLPFC activity, we would expect greater activity in this region with age.

We were surprised to find a significant u-shaped relationship between L DLPFC response and age using GLM 2 (Fig. 6A; ANOVA F=4.89, p=0.01), in which L DLPFC response decreases the further the subject gets from the middle age group. To test this, we estimated a linear regression using absolute difference from the median age to predict L DLPFC response during delayed choice. We found that absolute difference from median age explained 18% of the variance in L DLPFC response during delayed choice (R^2=0.18 β=-0.003 p=0.002).

We hypothesized that decreased L DLPFC function in our older subject group results from their decreased vStr value response with age, shown in the previous section. Conversely, increased vStr sensitivity to value may result in more need for L DLPFC response. To test this, in a *post hoc* analysis we estimated the relationship between vStr value response and L DLPFC activity during delayed choice using GLM 2. We found that L DLPFC activity during delayed choice was positively correlated with vStr response to value (Fig. 6B; R^2=0.14 β=0.48 p=0.01). This

relationship remains significant when controlling for age, grey matter density, and log(*k*) (R^2=.17, p=.01). To verify that this relationship is specific to the vStr, we then estimated the same relationship with the vmPFC ROI, which was not significant (p=.08). These results suggest why our older subjects, who have much less vStr sensitivity to value, do not have strong L DLPFC response during delayed choice.

Next, we specifically look at L DLPFC activity during delayed choices where self-control was required, when the immediate and delayed payments were of similar subjective values (GLM 3). We found that although L DLPFC activity increased in the older group in hard versus easy trials, but this was not significant (Fig. 6C; ANOVA F=0.34, p=0.72; R^2=0.003 β=0.0004 p=0.71).

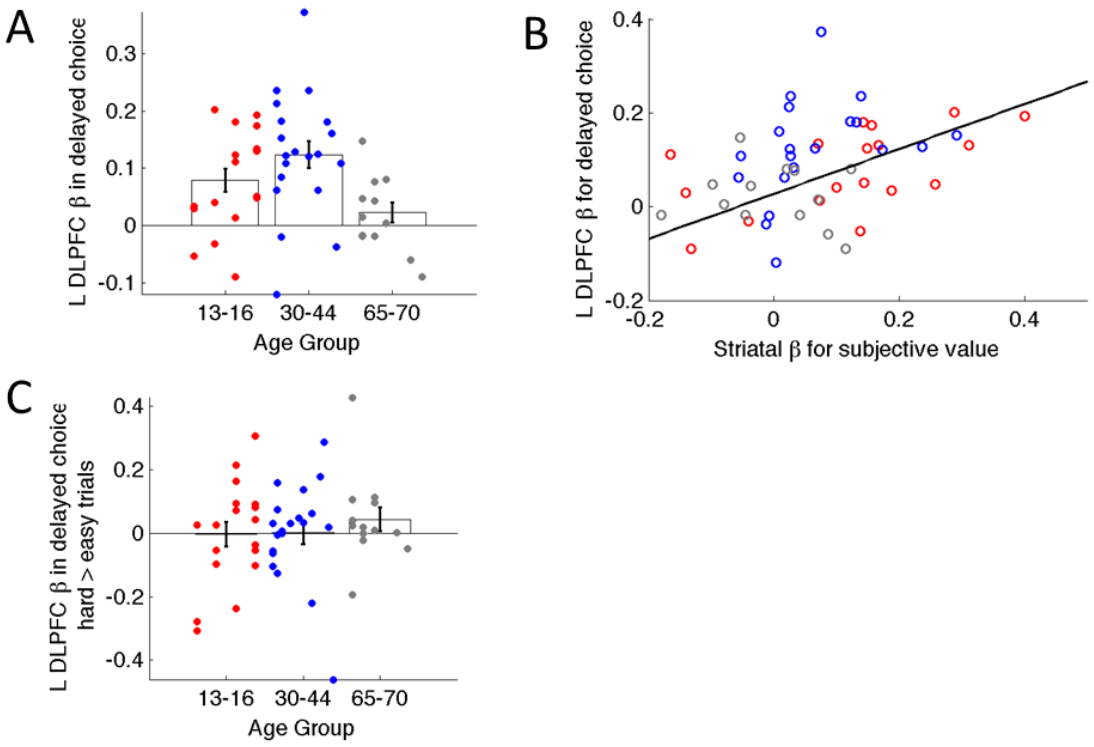

Figure 6. L DLPFC during delayed choice by age and value response. A) L DLPFC response during delayed choice. (GLM 2). B) L DLPFC activity during delayed choice as a function of vStr value response (GLM 2). C) L DLPFC response during delayed choice in hard, relative to all other, trials (GLM 3). Bars represent the average response, with error bars representing standard error. Dots represent individual values for subjects (Red = adolescent, blue = middle, grey = older).

L DLPFC connectivity changes with age. At the outset of this study, we hypothesized that age-related improvements in self-control could be linked to increased connectivity between the L DLPFC and the vmPFC. To test this in our study, we used PPI 1 to estimate changes in delayed-choice related connectivity across our age groups. Across all trials, we find that DLPFC-vmPFC connectivity slightly decreases with age (Fig 7A; ANOVA F=2.24, p=0.12; R^2=0.09 β=-0.0028 p=0.03). However, DLPFC-vStr connectivity does not change with age (Fig 7B; ANOVA F=1.88, p=0.16; R^2=0.06 β=-0.0019 p=0.08).

Next we investigated how connectivity changes with age during delayed choice, specifically in hard trials using PPI 2. We find that, in contrast to connectivity in all trials, DLPFC-vmPFC connectivity increases with age (Fig. 7C; ANOVA F=2.41, p=0.10; R^2=0.08 β=0.0039 p=0.04). There were no changes in L DLPFC-vStr connectivity with age (Fig; 7D; ANOVA F=1.88, p=0.16; R^2=0.06 β=-0.002 p=0.08).

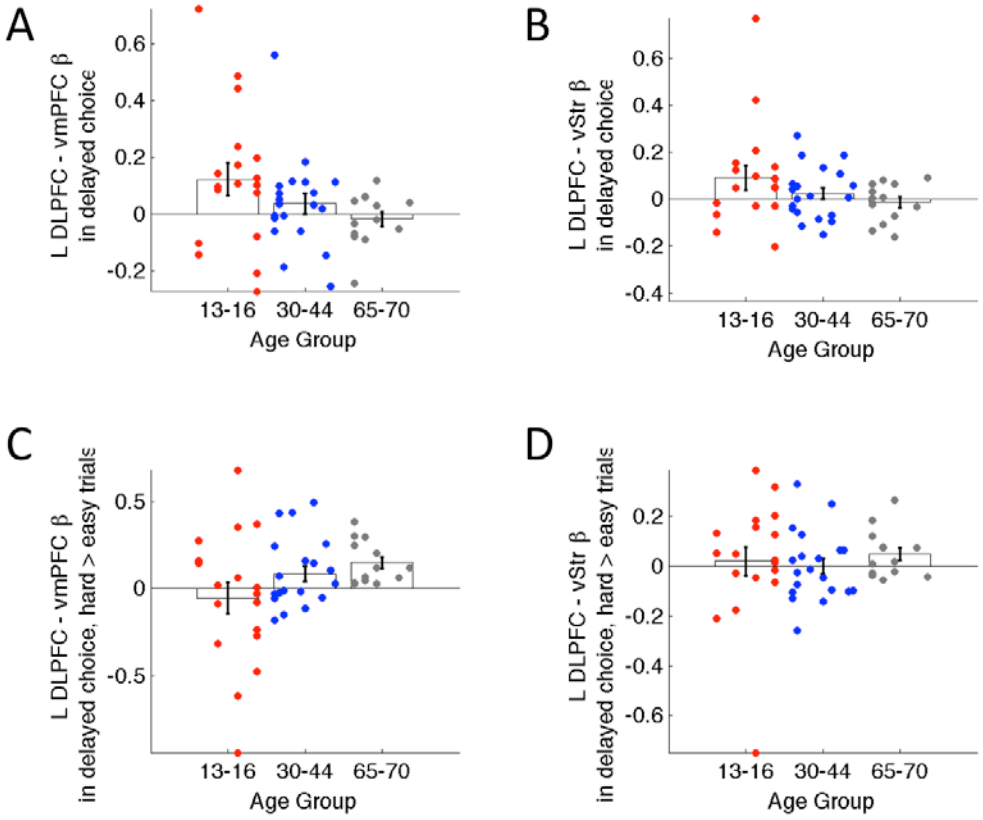

Figure 7. L DLPFC connectivity changes with age (PPI 2). A) Connectivity between L DLPFC and vmPFC during delayed choice by age B) Connectivity between L DLPFC and vStr during delayed choice by age C) L DLPFC-vmPFC connectivity during delayed choice in hard, relative to all other, trials by age. D) L DLPFC-vStr connectivity during delayed choice in hard, relative to all other, trials by age. Bars represent the average response, with error bars representing standard error. Dots represent individual values for subjects (Red = adolescent, blue = middle, grey = older).

4.4 Discussion

In this study, we set out to understand the sources underlying improvements in the ability to successfully exert self-control using an intertemporal choice task. To examine the neural underpinnings of self-control, we capitalized on a population-wide phenomenon of improvements in self-control: healthy human development and aging. By recording a proxy of brain activity, BOLD response in fMRI, while subjects made choices between a small amount of money today and a larger amount in the future, we attempted to assess whether specific neural substrates were linked

to improvements in self-control while controlling for environmental influences that also change with age. Specifically, we hypothesized that changes in the function and connectivity of value and control related regions would be linked to increases in the ability to delay gratification with age.

We found that discounting is tightly related to age in this task, even when controlling for demographic variables like financial security. Moreover, choice-related variables such as response times, and how strongly choices were predicted by preference (i.e., how noisy choices were), did *not* change with age. This allows us to propose that underlying neural variables, rather than environmental factors, may be an underlying source of lifespan changes in self-control.

Our behavioral results inform the on-going debate concerning the relationship between temporal discounting and ageing. For example, previous studies have found no change after 30 (Green et al., 1996), a curvilinear relationship between discounting and age (Read & Read, 2004), *increased* impulsivity (Green, Myerson, & Ostaszewski, 1999), no age-related change (Chao et al., 2009; Samanez-Larkin et al., 2011). Here, we echo other recent findings (Halfmann et al., 2013) that only when controlling for cognitive impairments such as dementia does temporal discounting decreases through the lifespan. This helps bring experimental results more in line with real-world evidence that pathological behaviors related to self-control deficits, such as gambling and alcoholism, decrease across the lifecycle (Gullone & Moore, 2000; National Institute on Alcohol Abuse and Alcoholism, 2008).

Because previous findings show that vStr, not vmPFC, changes from childhood through adolescence (Galvan et al., 2006), we hypothesized that there would be a similar pattern when looking from adolescence through senior years. We confirmed that sensitivity to value in the vStr, and not the vmPFC, decreased with age, despite significant grey matter decreases in both regions. Using mediation analysis, we showed that this decreased vStr value response seemed to mediate age-related decreases in discounting in our population, and the relationship between impulsivity and value response held even controlling for grey matter density in this region.

Neither adolescents nor adults showed increased L DLPFC activity when self-control was required to select the delayed option. The older group, however, had increased (but not significant) L DLPFC activity during these hard choices. Moreover, when self-control was required to select the delayed option, we found that connectivity between the L DLPFC and vmPFC increased with age, despite overall decreased connectivity with age in all trials. This suggests that older individuals may

be better able to identify cases in which regulation is required, and L DLPFC responds selectively to increase neural efficiency. We propose that this selective deployment of L DLPFC during very difficult trials, in addition to decreased response of vStr to reward, may underpin increases in self-control across the lifespan.

Here, we assume that self-control improvements arise from changes in brain composition and connectivity that occur as a natural part of the aging process. However, one could imagine instead that repeated, successful execution of self-control alters the connectivity and composition of the brain. Certainly we know that neural connectivity changes with experience (e.g., Hebbian learning, Hebb, 1949). Because experience is correlated with age, and because individuals with better monetary self-control have lower mortality rates (Boyle et al., 2013) and are therefore more likely to be part of our experimental sample, the observed experimental outcome of this scenario is the same as the result from our study. This question of causality is difficult to address with cross-sectional research. Longitudinal studies, or pre-post studies following a population that is learning or being trained to exert self-control over a period of time, may be able to address this question.

Temporal discounting is of great importance for well-being. Increased temporal discounting is related to decreased well-being (Forstmeier et al., 2011), poor academic performance (Mischel et al., 1988), and gambling and drug addiction (Reynolds, 2006). This study contributes to our growing understanding of how individual differences in brain structure and function improve one's ability to exert self-control, which is a necessary step toward developing interventions to improve choices.

Chapter 5

CONCLUSION

At the outset of this manuscript, I discussed some basic questions about how and when different types of value information are constructed, and about our struggle as humans to execute self-control using the value of long-term, rather than short-term, rewards. The big-picture question I attempted to address was: *what features of the decision process allow individuals to down-regulate the value of tempting options in order to forgo them in favor of greater future reward?* I presented converging evidence from both monetary and dietary choice domains and used fMRI, simulated data, and a new behavioral technique to answer this question. Below, I summarize my findings, discuss their significance, and pose some questions for the future of the field.

Do some types of attribute values take longer to process? We proposed that different kinds of value may vary in the ease with which they are calculated by the brain's decision circuitry, and therefore may be available to enter the decision process at different times. Evidence from the studies presented in Chapters 2 and 3 suggests that this depends on how different those attributes are. For example, the choice-relevant attributes in Chapter 2, tastiness and healthfulness, are distinctly different. One is visceral and easy to represent – for example, even very young children intuitively know how tasty foods are to them. Healthfulness, instead, is a fairly abstract, complex attribute that involves many pieces of information such as calorie, fat, and carbohydrate content. As predicted based on the different nature of these value types, we found that taste and health processing speeds were quite different and that, as discussed below, these times influence behavior. This has implications for accumulator models of choice which, until now, have assumed a unitary value.

In the experiment described in Chapter 3, individuals chose between different amounts of money delivered after varying delays. Evidence from that study suggests that value processing speeds for immediate and delayed monetary rewards were statistically indistinguishable at the group level. This similarity in processing speeds suggests that the brain constructs the subjective value of monetary rewards at similar times, perhaps even using the same process, regardless of their delay interval.

Together, the findings from the studies presented in Chapters 2 and 3 suggest that processing speeds do differ when the attributes are sufficiently distinct in their basic nature. Further investigation of this timing property promises to be a fruitful line of research.

Do earlier-processed attribute values have a larger weight in choice? Do faster overall processing speeds lead to a better reflection of lower-weighted attribute values? The accumulator models of choice introduced in Chapter 1 are the theoretical foundation for our hypothesis that earlier-processed values will have a larger influence on final choice than later-processed ones (see Chapter 1, Figure 2 for an illustration of this prediction). In Chapter 2, we termed this the "computational advantage" that the earlier-processed attribute has in the decision process. In that study, taste was processed earlier than health on average, and there was a substantial amount of individual variation in health vs. taste processing times for each individual. This allowed us to determine whether or not earlier processed values have a larger weight in choice using our experimental data. We find that earlier taste, relative to health, processing speeds led subjects to select the tastier, less healthy food more often. The strength of this timing property's prediction of dietary self-control was large; as much as thirty-nine percent of individual variance in self-control was explained by the computational advantage of the earlier-processed attribute alone.

As previously noted, the time it took participants to process the value of rewards delivered immediately and at a delay did not systematically differ. That made investigating the influence of processing speeds on monetary intertemporal choice a challenge to address using the same techniques as those used in dietary choice where these speeds differed widely. Instead, I paired our task's experimental data with simulated choices generated using the DDM, and found evidence that even slightly earlier processing times influence choice in favor of the earlier-processed attribute. I extend this question to find that earlier monetary subjective value processing speeds in general (rather than separated by immediate and delayed rewards) influence temporal discounting behavior. I propose that earlier value processing provides more time for delayed value information (weaker-represented, due to discounting) to accumulate, resulting in more patient choices. Conversely, very late value calculations result in noisier choices that more strongly reflect the more heavily-weighted attribute (typically immediate reward, due to impatience). Moreover, I found distinct trial-to-trial variability in processing speeds that result in uncharacteristic choices. When the immediate reward's subjective value is processed unusually early for a relatively patient subject, she often makes the more impulsive choice. Conversely, when an impatient subject processes delayed reward

value uncharacteristically early in a trial, that subject is enabled to succeed at self-control by selecting the larger, delayed reward. I propose that trial-to-trial changes in processing speeds could be due to momentary shifts in attentional focus or choice framing, and further investigation of this phenomenon could help us understand how to influence choice behavior.

Together, the results from both dietary and monetary intertemporal choice provide strong evidence for the theory I proposed in Chapter 1: attribute value processing speeds have a large influence on choice behavior. Earlier processing can allow information about an attribute we may weigh less heavily in choice (e.g. health information, because we discount its benefits for our future health) to contribute proportionally more to the value signal, and therefore to final choice. Several implications of this finding are worth future study. One particularly interesting extension of this work would develop interventions based on this theory to bring choices more in line with long-term goals. Several possibilities are suggested later in this conclusion.

This finding is a natural extension of accumulator models of choice discussed in the introduction and in more detail in Chapters 2 and 3 (e.g., Busemeyer & Townsend, 1993; Ratcliff, 1980; Ratcliff & McKoon, 2008; Ratcliff & Rouder, 1998; Usher & McClelland, 2001). The time it takes the decision circuitry to process value, parameterized in the DDM by the non-decision time (NDT), which has not been given much attention in the literature. However, this thesis provides evidence from both experimental and modeling methodologies that it has a large influence on choice, and further focus on it has great promise for providing additional insight into the dynamics of choice.

Are there neural markers for improvements in self-control? The focus of this thesis is the decision process underlying self-control successes and failures. To examine the neural underpinnings of self-control, we capitalized on a population-wide phenomenon of improvements in self-control: healthy human development and aging. By recording a proxy of brain activity, BOLD response in fMRI, while subjects made choices between a small amount of money today and a larger amount in the future, I hoped to discover what neural substrates are linked to improvements in self-control while controlling for environmental influences that also change with age. I found that age is a strong predictor of impulsivity in this task; even when controlling for income, patience increased from our adolescents to adult samples, and more improvements were seen into senior years. Moreover, I found that although grey matter decreased with age, connectivity between brain regions associated with valuation, reward, impulse control, and future-thinking actually increased with age. These

findings contribute to our growing understanding of how self-regulation occurs, and how it can improve.

A great number of unresolved questions arise from the findings reported in this manuscript. One question plagues the self-control research at large: how can researchers disentangle preference from self-control? It would be quite easy to argue that some individuals with what appears to be poor self-control could simply have different second order preferences, i.e., ideas about the kind of person they want to be. For example, an individual who always selects the chocolate chip cookie, or the smaller immediate monetary reward, may simply not desire more money or a healthier life. The research reported here imposes the assumption that our subjects have second order preferences for a healthier, wealthier life. That simply may not be the case. This question can partially be investigated in the future by developing and using techniques that can impose or measure and control for such preferences. Directly querying individuals on their second-order preferences and goals is problematic (Bowling, 2005; Presser & Stinson, 1998; Schwarz & Hippler, 1995) but may be a place to start.

Accumulator models of choice are incredibly useful for understanding a variety of simple binary choice problems, and a great deal of work has been done to understand their influence on choice. Although this thesis proposes an amendment to these models by including different attribute values, further improvements are required to bring these models more in line with the real cognitive process, which could improve their explanatory power. One open question, for example, is whether there are any inhibitory elements that enter the decision process. In dietary choice, we have shown evidence that health is processed later for most people. However, once health is processed, our model assumes that the value of taste continues to accumulate at the same rate. This may not be the case; it is equally possible that health contributes some additional inhibitory influence on the rate of taste value accumulation once it is processed. Indeed, Ratcliff (1980) proposed a model in which rate of value accumulation changes during the decision process. Alternatively, taste inhibition could be an entirely separate parameter that enters the relative value signal independent of other value types. The value construction process itself requires further work, for example to understand whether various elements of taste (sweetness, saltiness, etc.) and health (calories, fat, etc.) are calculated independently and then integrated into one taste or health value, which seems likely. The new technique reported in this manuscript has great promise to investigate these questions.

In Chapter 4, I provide evidence that age-related neural changes are associated with improvements in self-control. However, it is difficult to pinpoint the originator of this change. One could imagine that the natural changes that occur as the body ages include changes to the structure of the brain, including composition and connectivity, and that the changes we see in self-control arise from these natural brain changes. Here, I implicitly assume this is the case. However, one could imagine instead that repeated, successful execution of self-control changes the connectivity and composition of the brain, resulting in the effect we find. Certainly we know that neural connectivity changes with experience (e.g. Hebbian learning, Hebb, 1949). Because experience is correlated with age, and because individuals with better monetary self-control have lower mortality rates (Boyle et al., 2013) and are therefore more likely to be part of our experimental sample, the observed experimental outcome of this scenario is the same as the result from our study. Of course, the true answer may be some combination of the two, and this dilemma boils down to a classic chicken-or-the-egg question of causality that is difficult to address with cross-sectional research. Longitudinal studies, or pre-post studies following a population that is learning or being trained to exert self-control over a period of time, may be able to address this question.

Investigating discrete, one-trial-counts, binary choice to understand the decision process is very common and incredibly useful in the decision sciences. However, it does not reflect our day-to-day lives. For example, it may be quite easy for an individual to decline a delicious bowl of chocolate ice cream in the moment during one of our experiments. However, over the course of several hours, or even several days, the knowledge of ice cream (in addition to tens of other tempting items) sitting in the freezer, with easy access and few short-term negative consequences, presents constant demand on the decision circuitry required for self-control. That is the real world we live in, and developing research methodologies to bring our experimental paradigms more in-line with the constant demands required to live a healthy life would be extremely useful.

This body of work contributes to our understanding of why some people succeed, and others fail, at self-control. A natural extension of this research would be to develop interventions that actually improve an individual's self-control either in the moment, or gradually over time. In dietary choice, it is possible that with education, practice, more prominent health information, or simply better attentional focus, an individual could shift the time at which health information is processed; there is certainly no evidence to suggest that value processing speeds are fixed. One easily-implemented intervention that springs directly from our processing speed findings would be to allow a waiting

period before choice. This would, theoretically, give lower-weighted value information longer to contribute to the decision process. Indeed, some evidence already exists suggesting that extending decision time results in less impulsive choice (Dai & Fishbach, 2013). Further, there no evidence that the brain's reward responsivity or connectivity are set in stone even despite one's age; developing techniques to down-regulating one's response to reward, e.g. harnessing the power of psychological distancing (e.g., Wang, Lin, Huang, & Yeh, 2012) could have great promise for improving self-control. A large sector of the population struggles with self-control daily, and understanding the mechanisms underlying self-control is a crucial step toward developing interventions that could greatly improve quality of life for many.

BIBLIOGRAPHY

Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychol Bull, 82*(4), 463-496.

Ainslie, G. (1992). *Picoeconomics*. Cambridge, UK: Cambridge University Press.

Ainslie, G., & Monterosso, J. R. (2003). Building blocks of self-control: increased tolerance for delay with bundled rewards. *J Exp Anal Behav, 79*(1), 37-48.

Andreyeva, T., Sturm, R., & Ringel, J. S. (2004). Moderate and Severe Obesity Have Large Differences in Health Care Costs. *Obesity Research, 12*(12), 1936-1943.

Arnett, J. (1992). Reckless behavior in adolescence: A developmental perspective. *Developmental Review, 12*, 339-373.

Ashby, F. G. (1983). A biased random walk model for two choice reaction times. *Journal of Mathematical Psychology, 27*(3), 277-297.

Ballard, K., & Knutson, B. (2009). Dissociable neural representations of future reward magnitude and delay during temporal discounting. *NeuroImage, 45*(1), 143-150.

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.

Basten, U., Biele, G., Heekeren, H. R., & Fiebach, C. J. (2010). How the brain integrates costs and benefits during decision making. *Proceedings of the National Academy of Sciences, 107*(50), 21767-21772.

Baumeister, R. F., & Vohs, K. D. (Eds.). (2004). *Handbook of Self-Regulation: Research, Theory, and Applications*. New York: Guilford Press.

Beaver, J. D., Lawrence, A. D., van Ditzhuijzen, J., Davis, M. H., Woods, A., & Calder, A. J. (2006). Individual differences in reward drive predict neural responses to images of food. *J Neurosci, 26*(19), 5160-5166.

Beck, A. T., Ward, C. H., Mendelson, M., Mock, J., & Erbaugh, J. (1961). An inventory for measuring depression. *Arch Gen Psychiatry, 4*, 561-571.

Berns, G. S., Laibson, D., & Loewenstein, G. (2007). Intertemporal choice – toward an integrative framework. *Trends in Cognitive Sciences, 11*, 482-488.

Bogacz, R. (2007). Optimal decision-making theories: linking neurobiology with behaviour. *Trends in Cognitive Sciences, 11*(3), 118-125.

Bowling, A. (2005). Mode of questionnaire administration can have serious effects on data quality. *J Public Health (Oxf), 27*(3), 281-291.

Boyle, P. A., Yu, L., Gamble, K. J., & Bennett, D. A. (2013). Temporal Discounting Is Associated with an Increased Risk of Mortality among Community-Based Older Persons without Dementia. *PLoS ONE, 8*(6), e67376.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spat Vis, 10*(4), 433-436.

Brandt, R. B. (1979). *A Theory of the Good and the Right*. Oxford: Oxford University Press.

Brandt, R. B. (1998). The Rational Criticism of Preferences. In C. Fehige & U. Wessels (Eds.), *Preferences* (pp. 63-77). Berlin and New York: de Gruyter.

Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1993). Responses of neurons in macaque MT to stochastic motion signals. *Visual neuroscience, 10*(06), 1157-1169.

Brodman, K., Erdmann, A. J., Jr, Lorge, I., Wolff, H. G., & Broadbent, T. H. (1951). The cornell medical index-health questionnaire: Ii. as a diagnostic instrument. *Journal of the American Medical Association, 145*(3), 152-157.

Busemeyer, J. R., & Diederich, A. (2002). Survey of decision field theory. *Mathematical Social Sciences, 43*(3), 345-370.

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol Rev, 100*(3), 432-459.

Camerer, C., Issacharoff, S., Loewenstein, G., O'donoghue, T., & Rabin, M. (2003). Regulation for Conservatives: Behavioral Economics and the Case for" Asymmetric Paternalism". *University of Pennsylvania Law Review*, 1211-1254.

Carey, B. (2008). Drug Rehabilitation or Revolving Door? *The New York Times*,

Carver, C. S., & White, T. L. (1994). Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: The BIS/BAS Scales. *Journal of Personality and Social Psychology, 67*(2), 319-333.

Casey, B. J., Somerville, L. H., Gotlib, I. H., Ayduk, O., Franklin, N. T., Askren, M. K., et al. (2011). Behavioral and neural correlates of delay of gratification 40 years later. *Proceedings of the National Academy of Sciences, 108*(36), 14998-15003.

Chabris, C. F., Laibson, D., Morris, C. L., Schuldt, J. P., & Taubinsky, D. (2008). Individual Laboratory-Measured Discount Rates Predict Field Behavior. *National Bureau of Economic Research Working Paper Series, No. 14270*.

Chabris, C. F., Morris, C. L., Taubinsky, D., Laibson, D., & Schuldt, J. P. (2009). The allocation of time in decision-making. *Journal of the European Economic Association, 7*(2-3), 628-637.

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.

Charlton, S. R., Yi, R., Porter, C., Carter, A. E., Bickel, W., & Rachlin, H. (2013). Now for Me, Later for Us? Effects of Group Context on Temporal Discounting. *J Behav Decis Mak, 26*(2), 118-127.

Chib, V. S., Rangel, A., Shimojo, S., & O'Doherty, J. P. (2009). Evidence for a Common Representation of Decision Values for Dissimilar Goods in Human Ventromedial Prefrontal Cortex. *The Journal of Neuroscience, 29*(39), 12315-12320.

Clithero, J. A., & Rangel, A. (2014). Informatic parcellation of the network involved in the computation of subjective value. *Soc Cogn Affect Neurosci, 9*(9), 1289-1302.

Dai, J., & Busemeyer, J. R. (2014). A probabilistic, dynamic, and attribute-wise model of intertemporal choice. *J Exp Psychol Gen, 143*(4), 1489-1514.

Dai, X., & Fishbach, A. (2013). When waiting to choose increases patience. *Organizational Behavior and Human Decision Processes, 121*(2), 256-266.

Dale, R., Kehoe, C., & Spivey, M. J. (2007). Graded motor responses in the time course of categorizing atypical exemplars. *Mem Cognit, 35*(1), 15-28.

de Jager, C. A., Budge, M. M., & Clarke, R. (2003). Utility of TICS-M for the assessment of cognitive function in older adults. *Int J Geriatr Psychiatry, 18*(4), 318-324.

de Wit, H., Flory, J. D., Acheson, A., McCloskey, M., & Manuck, S. B. (2007). IQ and nonplanning impulsivity are independently associated with delay discounting in middle-aged adults. *Personality and Individual Differences, 42*(1), 111-121.

Drobetz, R., Hanggi, J., Maercker, A., Kaufmann, K., Jancke, L., & Forstmeier, S. (2014). Structural brain correlates of delay of gratification in the elderly. *Behav Neurosci, 128*(2), 134-145.

Dshemuchadse, M., Scherbaum, S., & Goschke, T. (2013). How decisions emerge: action dynamics in intertemporal decision making. *J Exp Psychol Gen, 142*(1), 93-100.

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.

Eppinger, B., Schuck, N. W., Nystrom, L. E., & Cohen, J. D. (2013). Reduced striatal responses to reward prediction errors in older compared with younger adults. *J Neurosci, 33*(24), 9905-9912.

Fassbender, C., Houde, S., Silver-Balbus, S., Ballard, K., Kim, B., Rutledge, K. J., et al. (2014). The Decimal Effect: Behavioral and Neural Bases for a Novel Influence on Intertemporal Choice in Healthy Individuals and in ADHD. *Journal of Cognitive Neuroscience, 26*(11), 2455-2468.

Federal Reserve Board. (2014). *Report on the Economic Well-Being of U.S. Households in 2013*.

Figner, B., Knoch, D., Johnson, E. J., Krosch, A. R., Lisanby, S. H., Fehr, E., et al. (2010). Lateral prefrontal cortex and self-control in intertemporal choice. *Nat Neurosci, 13*(5), 538-539.

Flegal, K. M., Carroll, M. D., Ogden, C. L., & Curtin, L. R. (2010). Prevalence and trends in obesity among us adults, 1999-2008. *JAMA, 303*(3), 235-241.

Forstmeier, S., Drobetz, R., & Maercker, A. (2011). The delay of gratification test for adults: Validating a behavioral measure of self-motivation in a sample of older people. *Motivation and Emotion, 35*(2), 118-134.

Frankfurt, H. (1971). Freedom of the Will and the Concept of a Person. *The Journal of Philosophy, 68*(1).

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature, 40*(2), 351-401.

Freeman, J. B., & Ambady, N. (2009). Motions of the hand expose the partial and parallel activation of stereotypes. *Psychol Sci, 20*(10), 1183-1188.

Freeman, J. B., Ambady, N., Rule, N. O., & Johnson, K. L. (2008). Will a category cue attract you? Motor output reveals dynamic competition across person construal. *J Exp Psychol Gen, 137*(4), 673-690.

Fry, A. F., & Hale, S. (1996). Processing Speed, Working Memory, and Fluid Intelligence: Evidence for a Developmental Cascade. *Psychol Sci, 7*(4), 237-241.

Galvan, A., Hare, T. A., Parra, C. E., Penn, J., Voss, H., Glover, G., et al. (2006). Earlier development of the accumbens relative to orbitofrontal cortex might underlie risk-taking behavior in adolescents. *J Neurosci, 26*(25), 6885-6892.

Gold, J. I., & Shadlen, M. N. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature, 404*(6776), 390-394.

Green, L., Fristoe, N., & Myerson, J. (1994). Temporal discounting and preference reversals in choice between delayed outcomes. *Psychonomic Bulletin & Review, 1*(3), 383-389.

Green, L., & Myerson, J. (2004). A Discounting Framework for Choice With Delayed and Probabilistic Rewards. *Psychol Bull, 130*(5), 769-792.

Green, L., Myerson, J., Lichtman, D., Rosen, S., & Fry, A. (1996). Temporal discounting in choice between delayed rewards: the role of age and income. *Psychol Aging, 11*(1), 79-84.

Green, L., Myerson, J., & Ostaszewski, P. (1999). Discounting of delayed rewards across the life span: age differences in individual discounting functions. *Behav Processes, 46*(1), 89-96.

Gullone, E., & Moore, S. (2000). Adolescent risk-taking and the five-factor model of personality. *J Adolesc, 23*(4), 393-407.

Halfmann, K., Hedgcock, W., & Denburg, N. L. (2013). Age-Related Differences in Discounting Future Gains and Losses. *J Neurosci Psychol Econ, 6*(1), 42-54.

Hare, T., Camerer, C. F., & Rangel, A. (2009). Self-Control in Decision-Making Involves Modulation of the vmPFC Valuation System. *Science, 324*(5927), 646-648.

Hare, T., Hakimi, S., & Rangel, A. (2014). Activity in dlPFC and its effective connectivity to vmPFC are associated with temporal discounting. *Frontiers in Neuroscience, 8*.

Hare, T. A., Camerer, C. F., & Rangel, A. (2009). Self-Control in Decision-Making Involves Modulation of the vmPFC Valuation System. *Science, 324*(5927), 646-648.

Hare, T. A., Malmaud, J., & Rangel, A. (2011). Focusing attention on the health aspects of foods changes value signals in vmPFC and improves dietary choice. *J Neurosci, 31*(30), 11077-11087.

Hare, T. A., Schultz, W., Camerer, C. F., O'Doherty, J. P., & Rangel, A. (2011). Transformation of stimulus value signals into motor commands during simple choice. *Proceedings of the National Academy of Sciences, 108*(44), 18120-18125.

Hare, T. A., Tottenham, N., Galvan, A., Voss, H. U., Glover, G. H., & Casey, B. (2008). Biological substrates of emotional reactivity and regulation in adolescence during an emotional go-nogo task. *Biol Psychiatry, 63*(10), 927-934.

Harris, A., Hare, T., & Rangel, A. (2013). Temporally Dissociable Mechanisms of Self-Control: Early Attentional Filtering Versus Late Value Modulation. *The Journal of Neuroscience, 33*(48), 18917-18931.

Hebb, D. O. (1949). *The Organization of Behavior*. New York: Wiley & Sons.

Hunt, W. A., Barnett, L. W., & Branch, L. G. (1971). Relapse rates in addiction programs. *J Clin Psychol, 27*(4), 455-456.

Hutcherson, C., Plassmann, H., Gross, J. J., & Rangel, A. (2012). Cognitive regulation during decision-making shifts behavioral control between ventromedial and dorsolateral prefrontal value systems *Journal of Neuroscience, in press*.

Ikeda, S., Kang, M.-I., & Ohtake, F. (2010). Hyperbolic discounting, the sign effect, and the body mass index. *Journal of Health Economics, 29*(2), 268-284.

Jenkinson, M., Beckmann, C. F., Behrens, T. E., Woolrich, M. W., & Smith, S. M. (2012). FSL. *NeuroImage, 62*(2), 782-790.

Jocham, G., Hunt, L. T., Near, J., & Behrens, T. E. (2012). A mechanism for value-guided choice based on the excitation-inhibition balance in prefrontal cortex. *Nature Neuroscience, 15*(7), 960-961.

Johnson, E. J., Schulte-Mecklenbeck, M., & Willemsen, M. C. (2008). Process models deserve process data: Comment on Brandstätter, Gigerenzer, and Hertwig (2006).

Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nat Neurosci, 10*(12), 1625-1633.

Kable, J. W., & Glimcher, P. W. (2010). An "As Soon As Possible" Effect in Human Intertemporal Decision Making: Behavioral Evidence and Neural Mechanisms. *J Neurophysiol, 103*(5), 2513-2531.

Kahnt, T., Heinzle, J., Park, S. Q., & Haynes, J. D. (2011). Decoding different roles for vmPFC and dlPFC in multi-attribute decision making. *NeuroImage, 56*(2), 709-715.

Kim, B. E., Seligman, D., & Kable, J. W. (2012). Preference reversals in decision making under risk are accompanied by changes in attention to different attributes. *Frontiers in Neuroscience, 6*.

Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience, 13*(10), 1292-1298.

Krajbich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in Psychology, 3*.

Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences, 108*(33), 13852-13857.

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *The Quarterly Journal of Economics, 112*(2), 443-478.

Liberman, N., & Trope, Y. (2008). The Psychology of Transcending the Here and Now. *Science, 322*(5905), 1201-1205.

Loewenstein, G. (1996). Out of Control: Visceral Influences on Behavior. *Organizational Behavior and Human Decision Processes, 65*(3), 272-292.

Loewenstein, G., & Prelec, D. (1992). Anomalies in Intertemporal Choice: Evidence and an Interpretation. *The Quarterly Journal of Economics, 107*(2), 573-597.

Lohse, G. L., & Johnson, E. J. (1996). A Comparison of Two Process Tracing Methods for Choice Tasks. *Organizational Behavior and Human Decision Processes, 68*(1), 28-43.

Luce, R. D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization*. Oxford: Oxford University Press.

Manly, J. J., Schupf, N., Stern, Y., Brickman, A. M., Tang, M. X., & Mayeux, R. (2011). Telephone-based identification of mild cognitive impairment and dementia in a multicultural cohort. *Arch Neurol, 68*(5), 607-614.

Marketdata Enterprises Inc. (2014). *The U.S. Weight Loss Market: 2014 Status Report & Forecast*.

Mather, M., Canli, T., English, T., Whitfield, S., Wais, P., Ochsner, K., et al. (2004). Amygdala responses to emotionally valenced stimuli in older and younger adults. *Psychol Sci, 15*(4), 259-263.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In J. E. Mazur, J. A. Nevin & H. Rachlin (Eds.), *Quantitative analysis of behavior* (Vol. 5, pp. 55-73). Hillsdale, N.J.: Erlbaum.

McKerchar, T. L., Green, L., Myerson, J., Stephen Pickford, T., Hill, J. C., & Stout, S. C. (2009). A Comparison of Four Models of Delay Discounting in Humans. *Behav Processes, 81*(2), 256-259.

McKinstry, C., Dale, R., & Spivey, M. J. (2008). Action Dynamics Reveal Parallel Competition in Decision Making. *Psychological Science, 19*(1), 22-24.

McNamee, D., Rangel, A., & O'Doherty, J. P. (2013). Category-dependent and category-independent goal-value codes in human ventromedial prefrontal cortex. *Nat Neurosci, 16*(4), 479-485.

Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The Drift Diffusion Model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgment and Decision Making, 5*(6), 437-449.

Mischel, W. (2014). *The Marshmallow Test: Mastering Self-Control*. Boston: Little, Brown and Company.

Mischel, W., Shoda, Y., & Peake, P. K. (1988). The nature of adolescent competencies predicted by preschool delay of gratification. *J Pers Soc Psychol, 54*(4), 687-696.

Mitchell, S. H. (1999). Measures of impulsivity in cigarette smokers and non-smokers. *Psychopharmacology (Berl), 146*(4), 455-464.

Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *The Journal of Neuroscience, 32*(7), 2335-2343.

Muraven, M., & Baumeister, R. F. (2000). Self-regulation and depletion of limited resources: does self-control resemble a muscle? *Psychol Bull, 126*(2), 247-259.

Myerson, J., Green, L., & Warusawitharana, M. (2001). Area under the curve as a measure of discounting. *J Exp Anal Behav, 76*(2), 235-243.

National Institute on Alcohol Abuse and Alcoholism. (2008). *Alcohol Research: A Lifespan Perspective*.

O'Donoghue, T., & Rabin, M. (1999). Doing It Now or Later. *The American Economic Review, 89*(1), 103-124.

Paeratakul, S., York-Crowe, E. E., Williamson, D. A., Ryan, D. H., & Bray, G. A. (2002). Americans on Diet. *Journal of the American Dietetic Association, 102*(9), 1247-1251.

Park, S. Q., Kahnt, T., Rieskamp, J., & Heekeren, H. R. (2011). Neurobiology of value integration: when value impacts valuation. *The Journal of Neuroscience, 31*(25), 9307-9314.

Pawitan, Y. (2001). *In All Likelihood: Statistical Modelling and Inference Using Likelihood*. New York: Oxford University Press.

Peters, J., & Buchel, C. (2010). Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-mediotemporal interactions. *Neuron, 66*(1), 138-148.

Peters, J., & Buchel, C. (2011). The neural mechanisms of inter-temporal decision-making: understanding variability. *Trends Cogn Sci, 15*(5), 227-239.

Petersen, A., Crockett, L., Richards, M., & Boxer, A. (1988). A self-report measure of pubertal status: Reliability, validity, and initial norms. *Journal of Youth and Adolescence, 17*(2), 117-133.

Pine, A., Seymour, B., Roiser, J. P., Bossaerts, P., Friston, K. J., Curran, H. V., et al. (2009). Encoding of marginal utility across time in the human brain. *J Neurosci, 29*(30), 9575-9581.

Presser, S., & Stinson, L. (1998). Data Collection Mode and Social Desirability Bias in Self-Reported Religious Attendance. *American Sociological Review, 63*(1), 137-145.

Rangel, A. (2013). Regulation of dietary-choice by the decision-making circuitry. *Nature Neuroscience*.

Rangel, A., & Clithero, J. (2013). The computation of stimulus values in simple choice. In P. W. Glimcher & E. Fehr (Eds.), *Neuroeconomics: Decision Making and the Brain* (Second ed.). London: Academic Press.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychol Rev, 85*(2), 59-108.

Ratcliff, R. (1980). A note on modeling accumulation of information when the rate of accumulation changes over time. *Journal of Mathematical Psychology, 21*(2), 178-184.

Ratcliff, R., Cherian, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *J Neurophysiol, 90*(3), 1392-1407.

Ratcliff, R., & McKoon, G. (1982). Speed and accuracy in the processing of false statements about semantic information. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8*(1), 16-36.

Ratcliff, R., & McKoon, G. (1989). Similarity information versus relational information: differences in the time course of retrieval. *Cogn Psychol, 21*(2), 139-155.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput, 20*(4), 873-922.

Ratcliff, R., & Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science, 9*(5), 347-356.

Ratcliff, R., & Smith, P. L. (2004). A Comparison of Sequential Sampling Models for Two-Choice Reaction Time. *Psychol Rev, 111*(2), 333-367.

Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychon Bull Rev, 9*(3), 438-481.

Read, D., & Read, N. L. (2004). Time discounting over the lifespan. *Organizational Behavior and Human Decision Processes, 94*(1), 22-32.

Reutskaja, E., Nagel, R., Camerer, C. F., & Rangel, A. (2011). Search Dynamics in Consumer Choice under Time Pressure: An Eye-Tracking Study. *American Economic Review, 101*(2), 900-926.

Reynolds, B. (2006). A review of delay-discounting research with humans: relations to drug use and gambling. *Behav Pharmacol, 17*(8), 651-667.

Rodriguez, C. A., Turner, B. M., & McClure, S. M. (2014). Intertemporal choice as discounted value accumulation. *PLoS ONE, 9*(2), e90138.

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionst model of decision making. *Psychol Rev, 108*(2), 370-392.

Rolls, E. T., Grabenhorst, F., & Deco, G. (2010). Choice, difficulty, and confidence in the brain. *NeuroImage, 53*(2), 694-706.

Samanez-Larkin, G. R., & Knutson, B. (2015). Decision making in the ageing brain: changes in affective and motivational circuits. *Nat Rev Neurosci, advance online publication*.

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. *J Neurosci, 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. *Front Neurosci, 5*.

Sangrock, J. (2002). *A Topical Approach to Life-Span Development*. New York: McGraw-Hill.

Scherbaum, S., Dshemuchadse, M., Fischer, R., & Goschke, T. (2010). How decisions evolve: the temporal dynamics of action selection. *Cognition, 115*(3), 407-416.

Scherbaum, S., Dshemuchadse, M., & Goschke, T. (2012). Building a bridge into the future: Dynamic connectionist modeling as an integrative tool for research on intertemporal choice. *Frontiers in Psychology, 3*.

Scherbaum, S., Dshemuchadse, M., Leiberg, S., & Goschke, T. (2013). Harder than Expected: Increased Conflict in Clearly Disadvantageous Delayed Choices in a Computer Game. *PLoS ONE, 8*(11), e79310.

Schwarz, N., & Hippler, H. (1995). Subsequent questions may influence answers to preceding questions in mail surveys. *Public Opinion Quarterly, 59*(1), 93-97.

Shamosh, N. A., & Gray, J. R. (2008). Delay discounting and intelligence: A meta-analysis. *Intelligence, 36*(4), 289-305.

Song, J.-H., & Nakayama, K. (2008). Target selection in visual search as revealed by movement trajectories. *Vision Research, 48*(7), 853-861.

Sowell, E. R., Peterson, B. S., Thompson, P. M., Welcome, S. E., Henkenius, A. L., & Toga, A. W. (2003). Mapping cortical change across the human life span. *Nat Neurosci, 6*(3), 309-315.

Spear, L. (2009). *The Behavioral Neuroscience of Adolescence*. New York: W. W. Norton & Company.

Spielberger, C. D., Gorssuch, R. L., Lushene, P. R., Vagg, P. R., & Jacobs, G. A. (1983). *Manual for the State-Trait Anxiety Inventory.*: Consulting Psychologists Press, Inc.

Spivey, M. J., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proc Natl Acad Sci U S A, 102*(29), 10393-10398.

Steinbeis, N., Haushofer, J., Fehr, E., & Singer, T. (2014). Development of Behavioral Control and Associated vmPFC–DLPFC Connectivity Explains Children's Increased Resistance to Temptation in Intertemporal Choice. *Cerebral Cortex*.

Steinberg, L., Graham, S., O'Brien, L., Woolard, J., Cauffman, E., & Banich, M. (2009). Age differences in future orientation and delay discounting. *Child Dev, 80*(1), 28-44.

Sullivan, N., Hutcherson, C., Harris, A., & Rangel, A. (2015). Dietary self-control is related to the speed with which attributes of healthfulness and tastiness are processed. *Psychol Sci, 26*(2), 122-134.

Suter, R. S., & Hertwig, R. (2011). Time and moral judgment. *Cognition, 119*(3), 454-458.

Thaler, R. H., & Benartzi, S. (2004). Save more tomorrow™: Using behavioral economics to increase employee saving. *Journal of political Economy, 112*(S1), S164-S187.

Thaler, R. H., & Sunstein, C. R. (2003). Libertarian paternalism. *American Economic Review*, 175-179.

Thompson, E. R. (2007). Development and Validation of an Internationally Reliable Short-Form of the Positive and Negative Affect Schedule (PANAS). *Journal of Cross-Cultural Psychology, 38*(2), 227-242.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychol Rev, 108*(3), 550-592.

Wang, Y.-L., Lin, Y.-C., Huang, C.-L., & Yeh, K.-H. (2012). Benefitting from a different perspective: The effect of a complementary matching of psychological distance and habitual perspective on emotion regulation. *Asian Journal of Social Psychology, 15*(3), 198-207.

Weber, E. U., Johnson, E. J., Milch, K. F., Chang, H., Brodscholl, J. C., & Goldstein, D. G. (2007). Asymmetric discounting in intertemporal choice: a query-theory account. *Psychol Sci, 18*(6), 516-523.

Wechsler, D. (2011). *Wechsler Abbreviated Scale of Intelligence, Second Edition*. San Antonio, Tx: Psychological Corporation.

Zhang, Y., Brady, M., & Smith, S. (2001). Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans Med Imaging, 20*(1), 45-57.