

Towards a psychological computing approach to digital lifestyle interventions

Citation for published version (APA):

Zhang, C. (2019). *Towards a psychological computing approach to digital lifestyle interventions*. Eindhoven: Technische Universiteit Eindhoven.

Document status and date:

Published: 22/11/2019

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Towards
**A Psychological Computing Approach
to Digital Lifestyle Interventions**

Chao Zhang

The work described in this thesis has been carried out at the Human-Technology Interaction group and Eindhoven University of Technology and at Philips Research, Eindhoven. It is sponsored by the Data Science flagship collaboration between TU/e and Philips.

© Chao Zhang, 2019.

CIP-DATA LIBRARY TECHNISCHE UNIVERSITEIT EINDHOVEN

Zhang, C. (Chao)

Towards a Psychological Computing Approach to Digital Lifestyle Interventions
by Chao Zhang

Eindhoven: Technische Universiteit Eindhoven, 2019. Proefschrift.

A catalogue is available from the Eindhoven University of Technology Library

ISBN: 978-94-638-0579-7

NUR: 770

Keywords: behavior change, lifestyle behavior, self-control, habit formation, computational model, digital health intervention

Cover design by Weihan Chen & Chao Zhang

Towards a Psychological Computing Approach to Digital Lifestyle Interventions

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van
de rector magnificus prof.dr.ir. F.P.T. Baaijens,
voor een commissie aangewezen door het College voor
Promoties, in het openbaar te verdedigen
op vrijdag 22 november 2019 om 11:00 uur

door

Chao Zhang

geboren te Wuhan, Hubei, China

Dit proefschrift is goedgekeurd door de promotoren en de samenstelling van de promotiecommissie is als volgt:

voorzitter:	prof.dr. C.J.H. Midden
1 ^e promotor:	prof.dr. W.A. IJsselsteijn
copromotor(en):	dr.D. Lakens dr. ir. J. Vanschoren
leden:	prof.dr. D.T.D. de Ridder (Universiteit Utrecht) prof.dr. P. Markopoulos prof.dr. W. Hofmann (Ruhr-Universität Bochum) prof.dr. R.M.M. Crutzen (Universiteit Maastricht)
adviseur(s):	dr. A. van Wissen (Philips Research)

Het onderzoek of ontwerp dat in dit proefschrift wordt beschreven is uitgevoerd in overeenstemming met de TU/e Gedragscode Wetenschapsbeoefening.

Contents

1 Introduction

1.1 Introduction	1
1.2 The gap between psychological theories and digital lifestyle interventions ..	4
1.3 The advantages of computational models over traditional theories	6
1.4 A psychological computing approach to digital lifestyle interventions.....	9
1.5 The scope and overview of the thesis	12

2 An Adaptive Decision-Making Framework of Lifestyle Behavior Change

2.1 Introduction.....	15
2.2 Review of individual theories relating to lifestyle behavior change.....	17
2.3 Theory integration: an adaptive decision-making framework.....	23
2.4 Mapping digital intervention techniques to the framework	29
2.5 General discussion	34

3 A Sequential Sampling Model of the Integration of Habits and Goals

3.1 Introduction.....	39
3.2 The conceptual and the computational model	43
3.3 Simulation studies.....	49
3.4 General discussion	58

4 Modeling Habit Development in Dental Behavior in the Real-World

4.1 Introduction.....	63
4.2 Data collection.....	71
4.3 Behavioral results.....	79
4.4 Predicting toothbrushing behavior with theory-based models	92

4.5 General discussion	101
5 Evaluating Mouse-Tracking as a Technique to Reveal Self-Control Processes	
5.1 Introduction.....	109
5.2 Study 1: food-choice experiment with touch-screen and physical-mouse..	118
5.3 Study 2: (conceptual) replications of previous findings.....	129
5.4 Study 3: re-analyses of Sullivan et al. (2015) and simulations	135
5.5 Study 4: food-choice experiment with practice trials and different orientations	146
5.6 General discussion	153
6 Experience Sampling Method to Study the Variations of Self-Control Capacity in Daily Life	
6.1 Introduction.....	159
6.2 Method	162
6.3 Results.....	168
6.4. Discussion	181
7 Epilogue	
7.1 Contributions and insights of the thesis to the psychological computing approach	192
7.2 Challenges and future research.....	194
7.3 Ethical considerations.....	196
7.4 Conclusion	198
Bibliography	199
Summary	225
Acknowledgements.....	227
Curriculum Vitae.....	231

Chapter 1

Introduction

To understand our nature, history, and psychology, we must get inside the heads of our hunter-gatherer ancestors. For nearly the entire history of our species, Sapiens lived as foragers. The past 200 years, during which ever increasing numbers of Sapiens have obtained their daily bread as urban laborers and office workers, and the preceding 10,000 years, during which most Sapiens lived as farmers and herders, are the blink of an eye compared to the tens of thousands of years during which our ancestors hunted and gathered.

Yuval Noah Harari, 2011, p.45.

Foragers mastered not only the surrounding world of animals, plants and objects, but also the internal world of their own bodies and senses. They listened to the slightest movement in the grass to learn whether a snake might be lurking there. They carefully observed the foliage of trees in order to discover fruits, beehives and birds' nests. They moved with a minimum of effort and noise, and knew how to sit, walk and run in the most agile and efficient manner. Varied and constant use of their bodies made them as fit as marathon runners. They had physical dexterity that people today are unable to achieve even after years of practising yoga or t'ai chi.'

Yuval Noah Harari, 2011, p.55-56.

1.1 Introduction

Collective lifestyle, as enabled by technology and culture, largely defines how we live and from what we suffer. Without doubt, the modern lifestyle is a bliss for most people. Compared with our ancestors who hunted and gathered in tribes, we are no longer constantly worried about

¹ This chapter is partly based on Zhang, C., van Wissen, A., Lakens, D., Vanschoren, J., De Ruyter, B., & IJsselsteijn, W. A. (2016). Anticipating habit formation: a psychological computing approach to behavior change support. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct* (pp. 1247-1254). ACM.

getting fatally injured or harmed, or having to starve at the end of a long, industrious, but unfortunate day. Instead, a typical day for many of us is filled with sitting in between walls and concrete that isolate us from any foreseeable danger. Food, often with high-calorie, can be delivered to our fingers at any time. The comfort of modern life, however, comes at some costs. For one thing, perhaps with the exception of our fingers, we exercise our bodies so little that physical exercise becomes a luxury hobby. Moreover, eating has moved far beyond a survival necessity to a kind of entertainment, and our lust for fat and sugar is exploited by the food industry. Finally, the high pace of modern society and a worsened connectedness to nature and kinship contribute to stress, anxiety, and loneliness, which undermine physical health as well. The modern lifestyle is one entangled with chronic diseases and health problems.

According to a report by the World Health Organization (WHO) and the Lancet, chronic diseases account for up to 71% of mortality worldwide (Bennett et al., 2018), and they put enormous burdens on both the affected population and the healthcare system. Research in the last few decades has identified unhealthy lifestyle behaviors as a major risk factor for chronic diseases, including unhealthy diet, physical inactivity, smoking, and alcohol intake as the “big four” (Bauer, Briss, Goodman, & Bowman, 2014; Ezzati & Riboli, 2013). As much as we cannot go back to the old days of hunting in the forests, it seems reasonable for individuals to make efforts to change their unhealthy lifestyle behaviors.

Unfortunately, lifestyle behaviors are notoriously difficult to change. A change in lifestyle requires making not only one healthy decision (e.g., getting a vaccine), but a substantial number of healthy decisions over a prolonged period of time. Without external supports, even motivated individuals find themselves vulnerable to everyday temptations, self-control failures, and bad habits. Ever since the identification of the behavioral risk factors, intervention programs have been designed, tested, and implemented to promote behavior change. Traditionally, interventions target people with diagnosed chronic conditions or at-risk populations, and they are usually centralized and delivered through face-to-face communications. More recently, interests in lifestyle interventions as a solution to everyone’s health problems have really been intensified, thanks to the developments of digital technologies. Digital lifestyle interventions, delivered through computers, websites, and smartphones, are believed by many to be an effective solution for promoting healthy lifestyles at a large scale (e.g., Heron & Smyth, 2010; IJsselsteijn, de Kort, Midden, Eggen, & van den Hoven, 2006; Kaplan & Stone, 2013; Klasnja & Pratt, 2012). Reasons for this optimism about the technology include digital systems’ ubiquity in people’s daily lives, their tremendous ability to collect behavioral data, and the rapid advances in machine intelligence. All these reasons point to a future of personalized and “just-in-time” interventions by smart and mobile applications (see Intille, Kukla, Farzanfar, & Bakr, 2003; Jaimes, Calderon, & Lopez, 2015; Nahum-Shani, Hekler, & Spruijt-Metz, 2015).

1.1 Introduction

So, how far are we in reaching this future? On the positive side, the research and development of digital lifestyle interventions are still booming. In some sense, digital technology has transformed behavior change from a small sub-field of applied psychology to a multidisciplinary field that attracts psychologists, computer scientists, industrial designers, and ethical experts (e.g., see the growing popularity of behavior change research in Figure 1.1). Besides scientific research, the consumer market has already witnessed an abundance of digital lifestyle interventions in many different forms, exemplified primarily by self-tracking wearables (e.g., Fitbit) and mobile health apps (e.g., Google fit), but also by more sophisticated lifestyle coaching apps. For example, the app Habitica uses gamification strategies to support the development and change of health-related habits (Diefenbach & Müssig, 2019).

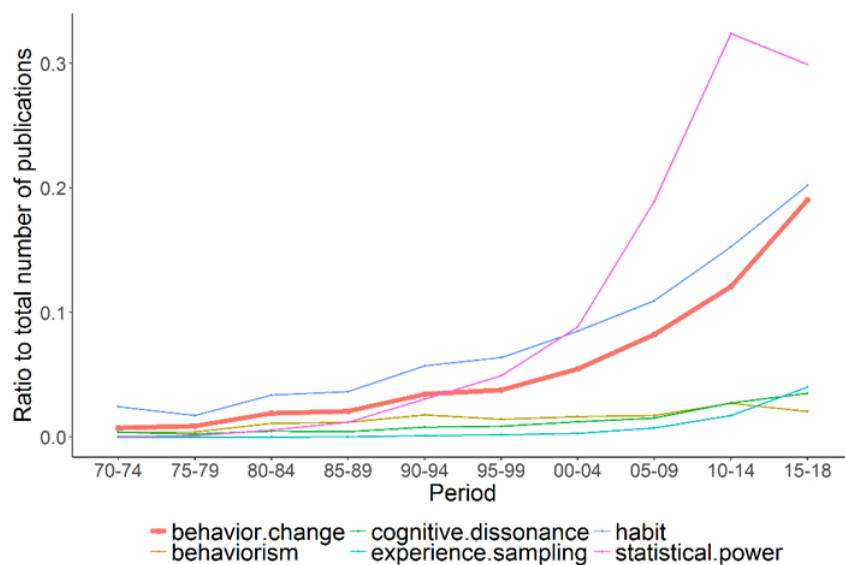


Figure 1.1 Popularity of key words in Google Scholar search results (values represent ratios of numbers of search results to the total number of publications in each period²).

However, when it comes to long-term user engagement and the actual effectiveness of digital lifestyle interventions, the picture is much less encouraging. According to a marketing analysis report (Ledger & McCaffrey, 2014), although one in ten U.S. citizens own a self-tracking device, around one-third lost interests in the products after just 6 months, and over half of them eventually abandoned the devices. This suggests that many users do not perceive the benefits from the devices as promised, such as data-driven self-discovery (see Kersten-van Dijk, Westerink, Beute, & IJsselsteijn, 2017) or increased motivation for change. The lack of effectiveness has been repeatedly reported in the literature. For example, a paper published in the *Journal of the American Medical Association* concluded that the self-tracking devices had negative rather than positive effects on motivating physical activities based on a clinical

² The search was restricted to journals containing the following key words: *psychology, psychological, behavior, behavioral, medicine, medical, neuroscience, health, personality, and statistics*.

trial (Jakicic et al., 2016). Based on a review of systematic reviews, Marcolino and colleagues (2018) also concluded that evidence was still limited regarding the efficacy of mobile health interventions, due to the low quality of reviews and a lack of long-term studies. A typical conclusion can be found in many more papers – short-term positive effects have been reported, but it remains unclear whether the same effects are also to be expected in the long-term, and whether these effects are generalizable to daily situations outside of research trials (Free et al., 2013; Hermsen, Frost, Renes, & Kerkhof, 2016; Kohl, Crutzen, & de Vries, 2013; Mateo, Granado-Font, Ferré-Grau, & Montaña-Carreras, 2015; Nour, Chen, & Allman-Farinelli, 2016).

A central argument in this thesis is that the effectiveness of digital lifestyle interventions depends critically on the progress of psychology and behavioral sciences and their translation to applications. One can consider digital lifestyle interventions as a form of *behavioral technologies* to contrast it with the more traditional *physical technologies*, such as automobiles or computers. For the latter, although human factors certainly play a role, physics and engineering sciences are at the core of their functioning. In contrast, the development information technologies alone cannot realize the promises of digital interventions, unless they are accompanied by knowledge about human behavior. In this thesis, a novel approach, called *psychological computing*, is motivated and explored, with an aim of maximizing the positive mutual influences between psychological theories and digital lifestyle interventions

1.2 The gap between psychological theories and digital lifestyle interventions

The close relationship between psychological theories and digital lifestyle interventions has been recognized by many researchers, and there have been repeated advocates for applying behavior change theories to improve the design of digital intervention systems (Kaplan & Stones, 2013; Michie & West, 2013; Patel, Asch, & Volpp, 2015; Riley et al., 2011; Saranummi et al., 2013). Ideally, theories and interventions should benefit each other. Good theories, when applied properly, are expected to increase the effectiveness of interventions. More specifically, theories may help to identify behavioral determinants as intervention targets, to translate general behavior change techniques to fine-tuned features in applications, and even to predict intervention outcomes. On the other hand, the vast amount of behavioral data collected by digital systems in people's daily lives would potentially contribute to the evaluation of existing behavior change theories (Dunton & Atienza, 2009; Saranummi et al., 2013). Compared with data from more traditional laboratory experiments, digital systems can provide behavioral data with larger and more diverse samples, greater ecological validity, and higher temporal resolution.

Despite these expectations, the synergy between theory development and intervention practice is far from ideal (for a review, see Prestwich, Webb, & Conner, 2015). The role of behavior change theories in digital interventions is not as prominent as one would hope. Several surveys indicate that the application rate of theories in digital intervention trials and commercial health apps ranges between 20% and 50% (Al-Durra, Torio, & Cafazzo, 2015; Conroy, Yang, Maher, 2014; Prestwich et al., 2015; Riley et al., 2011). When theories are applied, only 3 to 5 classical theories dominate in applied settings (Davis, Campbell, Hildon, Hobbs, & Michie, 2015; Webb, Joseph, Yardley, & Michie, 2010). State-of-the-art theoretical models from a broader psychology literature are rarely used by intervention designers. Moreover, the assumed benefits of applying theories have been questioned by an empirical study where no difference in effectiveness was found between theory-informed or theory-free interventions (Prestwich et al., 2014). Finally, as for theory development, data collected by digital systems are usually used for evaluating the effectiveness of specific interventions, but are rarely used to examine predictions derived from theoretical models.

One factor contributing to this “theory-intervention gap” is the lack of theory integration in behavior change research (see Gainforth, West, & Michie, 2015), especially those that are tailored to digital lifestyle interventions. Even for a phenomenon as complex as behavior change, the large number of individual theories pertaining to behavior change – 83 according to a systematic review (Davis et al., 2015) – clearly suggests that theory development in the field is still in the early stages. The number can also be overwhelming for intervention designers who want to have a grasp of the literature, but are not specialized in behavioral science. Perhaps the difficulty to orient in the literature can explain why only a small number of classical theories are popular in applied research (Davis et al., 2015). Many basic theoretical ideas in psychology, despite being highly relevant, are underrepresented in behavior change research, such as decision-making, reinforcement learning, self-control, and habit and goal-directed learning.

The lack of impact of theories on interventions also raises the question whether the level of knowledge about lifestyle behavior change is too limited to be useful. Two specific reasons have been proposed by Riley and colleagues (2011) to argue why traditional behavior change theories are inadequate in the digital age. First, there is a mismatch between traditional theories and digital interventions, in terms of at what temporal scale behaviors are represented. For example, while a healthy-eating app may intervene in its users’ dietary choices on a daily basis, the classic Transtheoretical Model (Prochaska & DiClemente, 1982) only describes stages of behavior change in terms of months. If a theory represents behavior at a coarse temporal scale, processes at finer scales are overlooked and time-intensive digital interventions cannot be informed. Second, most traditional theories are static rather than dynamic. In other words, they provide “snapshot” explanations about what factors determine behavior, but neglect the dynamic interactions between these factors and behavior. For instance, in the

widely applied Theory of Planned Behavior (Ajzen, 1991), no temporal dynamics of the behavioral determinants in the model are specified, nor are any mechanisms to account for the reciprocal influences of behavior on its determinants.

1.3 The advantages of computational models over traditional theories

Theory development in psychology can go through four different phases, as to describe, explain, predict, and control behavior. The problem of lifestyle behavior change would have been solved if theory development is already in the control phase, but most current theories are still in the phases of increasing explanatory power and predictive precision. In the past decades, formalizing theories as computational models have contributed greatly to advance psychological theories in various domains, such as speech perception (McClelland & Rumelhart, 1981), memory (Ratcliff, 1978), and higher-level cognition (Anderson, 1996). However, more applied fields, such as behavior change, are still dominated by descriptive statistical models and verbal explanatory theories.

Some traditional behavior change theories are descriptive, in the sense that they represent correlations or potentially causal relationships between variables, but not psychological processes or cognitive mechanisms underlying a phenomenon. For example, the Theory of Planned Behavior essentially states that the variance in actual behavior is partially related to the variance in three behavioral determinants (attitude, social norm, and perceived behavior control), but it does not offer a process model of how these determinants actually interact to generate behavior. Ajzen (1991) said himself that the Theory of Planned Behavior was developed mainly for behavior prediction rather than explanation. Note that the term *prediction* used by Ajzen, and as often by some social psychologists, means no more than correlation. Precise predictions are usually not possible given the moderate correlations between the variables.

Other traditional theories aspire to explain behavior verbally, such as Social Cognitive Theory (Bandura, 1989), the control theory of self-regulation (Carver & Scheier, 1982), and theories of habit formation in social psychology (e.g., Wood & Neal, 2007). However, they are high-level theories that lack clear operationalizations. For example, the process of habit formation has been described as the build-up of a mental link between a behavior and a context (e.g., Neal, Wood, & Quinn, 2006; Rothman et al., 2015; Wood & Neal, 2007), but this verbal description leaves out many necessary details: Does the rate of habit growth change over time? Or what happens to the mental association if the target behavior is not performed? Answering these questions would force one to formalize the theory into a formal model.

The advantages of computational models over verbal theories have been discussed extensively elsewhere (Farrell & Lewandowsky, 2010; Hintzman, 1991; Marewski & Olsson, 2009; Smaldino, 2017), but we highlight some key points here. First, computational models require

one to precisely define the core mechanism as well as peripheral assumptions in a theoretical system. Compared with verbal theories, which are often based on ambiguous language, computational models can sharpen the theorizations of individual researchers and facilitate more accurate communication among researchers (Farrell & Lewandowsky, 2010). Second, while predictions by verbal theories are mostly directional (e.g., positive or negative correlations, or directional differences between groups), predictions by computational models can be precise and comprehensive, e.g., in the form of point estimates of differences, or distinct quantitative patterns under different conditions. These predictions make the underlying theories more falsifiable, and the comparisons between theories more compelling than what would be possible with null-hypothesis significance testing (Marewski & Olsson, 2009). Third, through simulations, computational models free theorists from their limited human reasoning capacities to understand previously unknown implications of their theories (Smaldino, 2017). This is especially true when one deals with complex systems, such as lifestyle behaviors that are influenced by many internal and external factors.

There is an additional reason why computational models are especially useful for digital lifestyle interventions. Because these models are essentially computer programs or mathematical equations, they can be easily implemented in digital systems. If an implemented model is valid, a digital system is equipped not only with sensing power, but also the ability to reason about the behaviors of its users. There is indeed a growing interest in building computational models of behavior change in order to improve the application of psychological theories to digital interventions (e.g., Hekler et al., 2016; Nilsen & Pavel, 2013; Riley et al., 2011; Spruijt-Metz et al., 2015). Two examples are discussed in detail for illustrative purposes.

Example 1: Computerized Behavior Intervention model

One notable example is the Computerized Behavior Intervention (COMBI) model (Klein, Mogles, & van Wissen, 2011; see Figure 1.2), which has been implemented in real digital intervention systems (Kamphorst, Klein, & van Wissen, 2014). Based on the transtheoretical model of behavior change (Prochaska & DiClemente, 1982), the COMBI model formally represents a behavior change process as going through five stages – precontemplation, contemplation, preparation, action, and maintenance. It then defines the behavioral determinants that would move users from one stage to the next. Thus, a digital system equipped with the COMBI model has an internal representation of a user’s progress and the factors that influence the stage transitions. Through a search algorithm, the system can reason about the “bottlenecks” that hold back an individual user and select intervention techniques that target the restraining determinants (Klein, Mogles, & van Wissen, 2014).

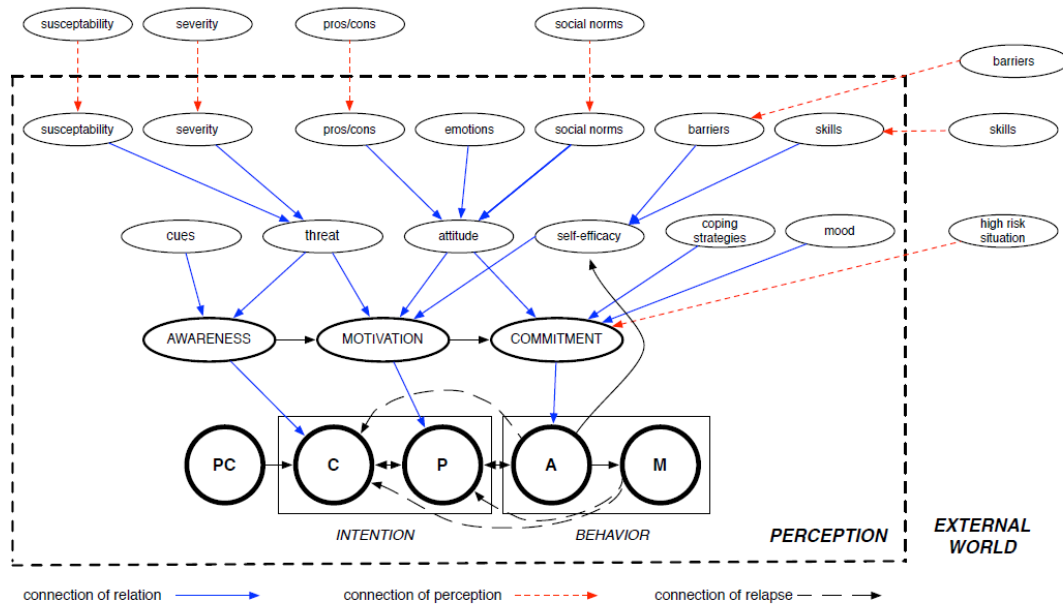


Figure 1.2 A visual representation of the COMBI model (adapted from Figure 1 in Klein et al., 2014).

Example 2: A control system model of Social Cognitive Theory

Another interesting example is a formalization of Social Cognitive Theory (SCT, Bandura, 1989) using a control engineering approach (Riley et al., 2015). Although Social Cognitive Theory encompasses a set of explanatory mechanisms of behavior change verbally (e.g., social learning, and self-efficacy), the evaluations and applications of the theory are largely at the phase of describing the *inter-individual* statistical relationships between variables in the theory. By using a fluid metaphor, variables in the theory are modeled as “inventories” or “reservoirs” with their levels fluctuate over time, depending on inflows from other variables and outflows to other variables (see Figure 1.3). This dynamic-system representation allows the modeling of the relationships between state variables as a set of differential equations. Thus, *intra-individual* dynamics of behavior, self-efficacy, outcome expectancy, among other determinants, and their reciprocal influences can be simulated precisely. According to Riley and colleagues (2015), their dynamic version of Social Cognitive Theory can potentially be used to facilitate intensive and adaptive interventions.

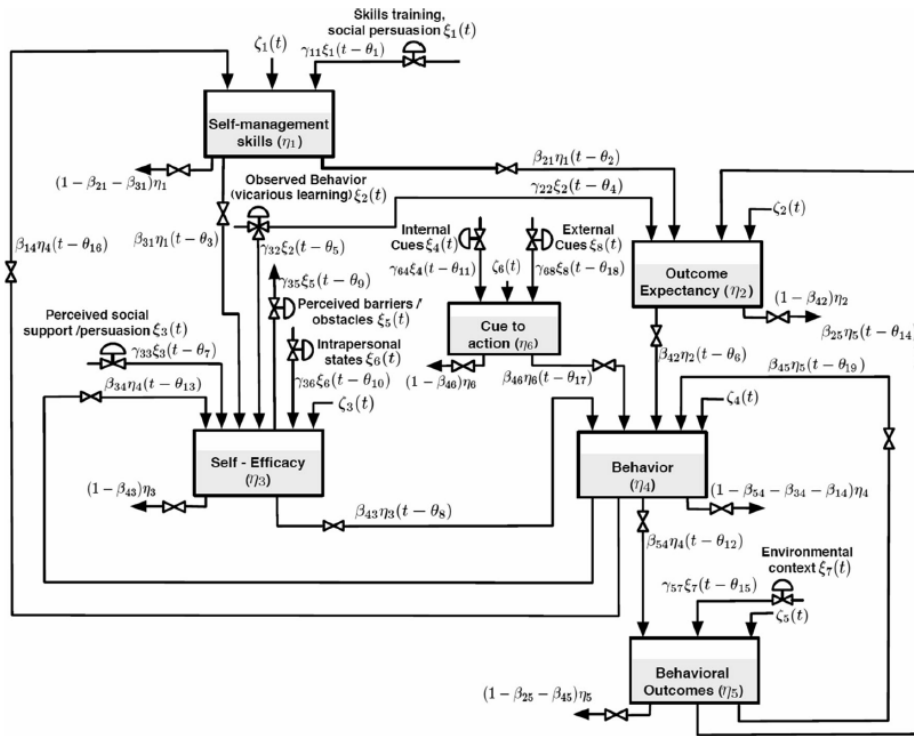


Figure 1.3 A control system model of Social Cognitive Theory (adapted from Figure 2 in Riley et al., 2015).

1.4 A psychological computing approach to digital lifestyle interventions

Motivated by the promise of using computational models to bridge the gap between psychological theories and digital lifestyle interventions, a *psychological computing*¹ approach to digital lifestyle interventions is proposed as a general framework of this thesis. At the core of the approach is a theory-based computational model of human behavior implemented in a digital system. The model assists the system to understand, predict, and change its users’ behaviors.

Figure 1.4 illustrates how the computational model relates to other system components in a digital system that follows the psychological computing approach (see Zhang et al., 2016). The system consists of an input component for gathering necessary data, a processing com-

¹ The choice of the term “psychological computing” can be better understood by comparing it to *affective computing* (Picard, 2000). While in affective computing research, the ultimate goal is to build machines that understand human emotions, in the psychological computing approach the goal is for machines to understand users’ internal cognitive states, such as goals, habits, and beliefs. Understanding emotions could also be relevant if it helps a system to change lifestyle behaviors, and thus we chose to use “psychological” rather than the more narrowed word “cognitive”. The term also emphasizes the use of psychological theories, in addition to data, to infer users’ internal states.

ponent for turning data into useful information based on the underlying model, and an actuation component to choose an intervention based on the information. On the user side, we adopt a cognitive view of behavior that any stimulus from the environment (including the system's intervention) can exert influence on the user's behavior through the mediation of a cognitive system. The input component monitors users' actual behaviors, senses environmental variables, and also collects information about users' cognitive states through self-reports (e.g., experience sampling in Chapter 6) or other indirect measures (e.g., mouse-tracking in Chapter 5). The actuation component selects appropriate intervention techniques to influence users' behaviors through the intended influences on their cognitive systems. Most importantly, the processing component uses a computational model to update cognitive states, predict behaviors, anticipate future change process, and to guide intervention selections (e.g., Chapter 3 & 4). As with the previous two examples, the computational model in this system is based on psychological theories.

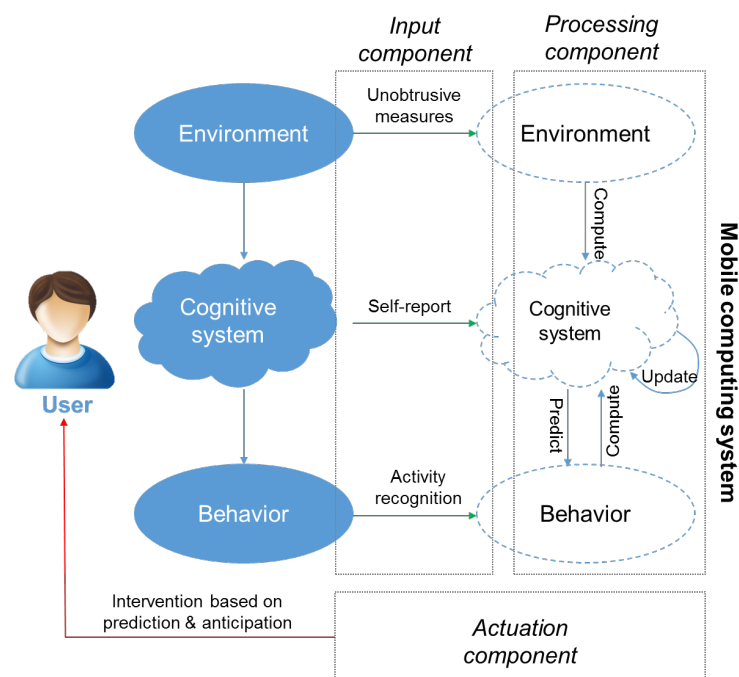


Figure 1.4 A model of a future digital system based on the psychological computing approach.

Moving towards the psychological computing approach requires an iterative process of theory development, model building, and using and evaluating the models in the real-world through digital lifestyle intervention systems (see Figure 1.5). A first step is to search for an overarching theoretical framework that integrates relevant psychological theories and to connect the framework to common behavior change techniques used in digital systems. From there, computational models can be developed and implemented in digital systems to guide intervention practice. In the other direction, data collected by the digital systems in real-world behavior

change processes can facilitate the improvements of the models. Thus, the psychological computing approach is only motivated to enhance intervention effectiveness, but also to maximize the use of time-intensive digital data to improve psychological theories.

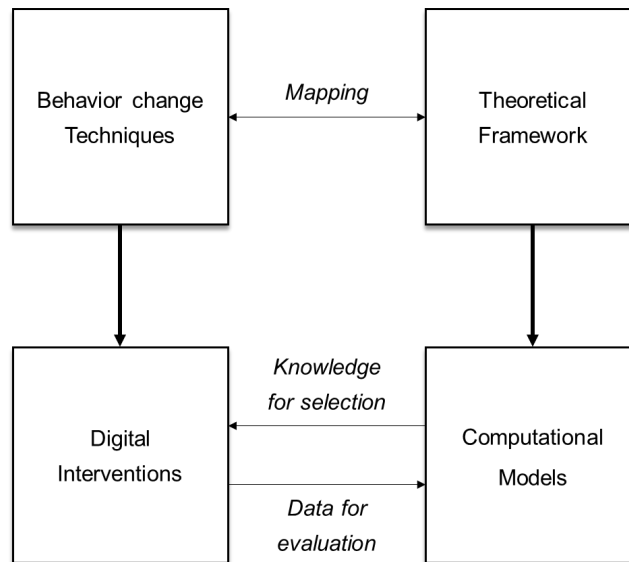


Figure 1.5 An iterative process of moving towards the psychological computing approach.

The psychological computing approach can be regarded as sitting in between a pure theory-driven approach and a pure data-driven approach to digital interventions. In traditional theory-based interventions (see the top half in Figure 1.5), behavior change techniques are derived from verbal theories in psychology and behavioral sciences, and are then implemented in digital systems often as fixed functions for all users. Although empirical studies are often conducted to evaluate the effectiveness of the implemented techniques, the loose connection between the often vague theories and the techniques makes it difficult to evaluate or falsify the underlying theories. In contrast, a pure data-driven approach, as often preferred by computer scientists, disregards psychological theories and uses machine learning models to predict behavior and to adapt interventions to different contexts (for example, see Chapter 3 in op den Akker, 2014). This approach is similar to the processes described in the bottom-half of Figure 1.5, except that the models in the data-driven approach are not derived from psychological theories. In the ideal world where psychological theories are perfect or human data are limitless, either approach alone could work well. However, in the reality where theories are still in developments and cognitive data are difficult to get, connecting the two worlds using theory-based computational models can be an efficient approach that borrows the strengths from both the theory-driven and the data-driven world. The psychological computing approach envisioned is especially applicable to domains where theories can make up for the lack of data for certain behavioral or cognitive processes (see Chapter 4 as an example).

1.5 The scope and overview of the thesis

The current thesis takes a few first steps to bridge the gap between psychological theories and digital lifestyle interventions, and to move towards the psychological computing approach. For this goal, a variety of methodologies are used, including theory review and integration, computational modeling, field intervention studies, and also explorations of novel data collection methods that potentially benefit both science and applications. Four main objectives are:

- To propose a general theoretical framework of lifestyle behavior change that integrates existing psychological theories, which can be mapped to digital intervention techniques
- To develop a computational model of some key processes in framework, and to test if the model can be validated based on empirical findings
- To apply and evaluate (part of) the computational model in a real-world lifestyle intervention problem, which serves as a use case of the psychological computing approach
- To explore methods of measuring cognitive states involved in making daily lifestyle decisions that can potentially be used in digital intervention systems

The work to fulfill the objectives is reported in a series of 6 chapters:

Chapter 2 briefly reviews relevant theories in the psychology literature, and then integrates the diverse theoretical ideas into a unified framework of lifestyle behavior change, called the *adaptive decision-making framework*. Common digital intervention techniques are mapped to the behavioral and cognitive processes in the framework. The chapter also identifies habit formation and self-control as the two key constructs to be examined in later chapters.

Chapter 3 reports a sequential sampling approach of modeling the decision-making and learning processes as described in the adaptive decision framework, focusing on how habits and goals are integrated in determining such decisions. The model is validated through simulation studies to reproduce classic empirical findings, and future extension of the model to daily contexts is discussed.

Chapter 4 evaluates the habit formation part of the computational model in Chapter 3 in two real-world digital intervention studies of dental behavior change. Data collected were used to understand the reciprocal relationship between habit, attitude, and behavior, and to test whether habit strength computed by the computational model could improve behavior prediction.

Chapter 5 & 6 explores two methods of measuring relevant aspects of the self-control in daily lifestyle decisions. Chapter 5 evaluates whether the mouse-tracking technique can be used to measure the cognitive processes that underlie dietary self-control and whether the technique

can be transferred from traditional laboratory settings to touchscreen interfaces that usually come with digital systems. Chapter 6 instead examines the variations of self-control capacity in people's daily lives using an experience sampling method. Data from two field studies contribute to the understanding of how self-control capacity varies inter- and intra-individually and how the variations are related to changes in people's affective states. The results may inform the timing of interventions in future applications.

Finally in Chapter 7, the psychological computing approach is motivated again by considering the challenges of health behavior change in a larger context, connecting it to both an evolutionary and a historical perspective. Contributions of the thesis, future research directions, and ethical issues are discussed in relation to the psychological computing approach.

Chapter 2

An Adaptive Decision-Making Framework of Lifestyle Behavior Change

“There is nothing so practical as a good theory.”

Kurt Lewin, 1951.

2.1 Introduction

In Chapter 1, we have made the case for the importance of basic psychological theories in developing digital lifestyle interventions, and highlighted the causes for the “theory-intervention gap”. It seems that the problem is not the lack of theories; in fact, intervention developers and digital system designers have an abundance of theories they can choose from, for example, from the 83 theories identified by behavior change experts in the book *ABC of Behaviour Change Theories* (Michie, West, Campbell, Brown, & Gainforth, 2014). The problem is that there are too many theories, and a lack of systematic integration. For practitioners without a formal training in psychology, the large number itself might cause difficulties to orientate in the literature and to select theories for particular intervention problems at hand. Without texts that clarify the connections and differences among the individual theories, confusions can arise when more than one theory attempts to explain the same behavioral process but with different terminologies. Such problematic situations are not uncommon since different theories may explain behavior at different levels of analysis (see Crutzen & Peters, 2018), focus on either statistical description or mechanistic explanation, and even take distinct philosophical views on human behavior. To overcome these barriers, perhaps there is nothing so practical as a good theory integration.

⁴ This chapter is partly based on Zhang, C., Lakens, D., & IJsselsteijn, W. A. (2019). Theory Integration for Lifestyle Behavior Change in the Digital Age: An Adaptive Decision-making Framework. <https://doi.org/10.31234/osf.io/fsw8t>.

This is not to say that theory integration has not been done before in the domain of behavior change. In fact, several notable examples can be found, including the i-Change model (de Vries, 2017), PRIME theory (West, 2006), Temporal Self-Regulation Theory (Hall & Fong, 2007), and the COMBI model (Klein, Mogles, & van Wissen, 2011). However, except for the COMBI model, the previous integrations have focused specifically on lifestyle behaviors and on digital interventions. In addition, it can be argued that none of the integrations has sufficiently addressed the two limitations of traditional behavior change theories highlighted in Chapter 1 – the temporal mismatch and the lack of dynamics (see also Riley et al., 2011).

These two limitations are closely related not only to the characteristics of time-intensive digital data (see Dunton & Atienza, 2009; Riley et al., 2011), but also to the characteristics of lifestyle behaviors themselves. Lifestyle behaviors, such as eating, exercising, or toothbrushing, are performed very frequently, usually at least on a daily basis, and on each occasion they are essentially fast daily decisions. This type of decisions (e.g., making one dinner choice) may be relatively inconsequential, but they can form larger behavioral “episodes” (e.g., following a diet), which may affect one’s health significantly. This characteristic of hierarchical organization sets lifestyle behaviors apart from single-time health behaviors or decisions, such as cancer screening or vaccination. The problem of temporal mismatch is rooted in the fact that most traditional behavior change theories were developed mainly to explain single-time decisions. Moreover, unlike single-time decisions, as daily lifestyle decisions are repeated, learning and adaptation through past experience are made possible and they form a big part of the puzzle of lifestyle behavior change and intervention. This requires the inclusion of temporal dynamics into behavior change theories. Overall, the two limitations raised by Riley and colleagues (2011) also reflect the fact that previous theories or theory integrations have not been focused on lifestyle behaviors.

Here we propose a new integrative theoretical framework, called adaptive decision-making, to explicitly account for lifestyle behavior change, addressing directly the two limitations of traditional theories. In doing so, the new framework represents lifestyle behaviors at two temporal levels: a lower level (*action level*) that matches the daily individual decisions and the time-intensive interventions realized by digital systems, and a higher level (*reflection level*) that matches the episodes of repeated decisions (see Figure 2.1). In addition, both decision-making processes (how behaviors are determined or decisions are made) and learning processes (how earlier behaviors or decisions influence later ones through cognitive determinants) at each level will be included in the framework. More broadly, the framework is intended to bridge the above-mentioned theory-intervention gap.

In the rest of the chapter, we first review important theoretical ideas relevant to lifestyle behavior change from a broad psychology literature. To facilitate theory integration, the individual theories are compared in terms of their temporal scales and their emphasis on either learning or decision-making. Next, the adaptive decision-making framework is introduced by

2.1 Introduction

integrating the relevant but disparate theoretical ideas into a two-level representation of lifestyle behavior change presented above (Figure 2.1). Afterwards, we relate the framework to intervention practice by mapping common behavior change techniques used in digital systems onto the behavioral processes in the framework. The chapter is concluded with a general discussion on the added value of the framework to behavior change theorists and digital intervention designers.

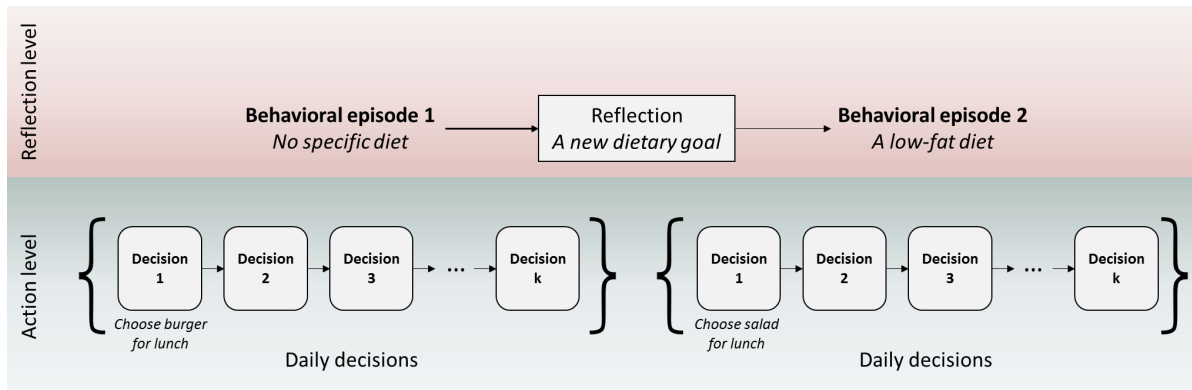


Figure 2.1 A two-level representation of lifestyle behavior (change).

2.2 Review of individual theories relating to lifestyle behavior change

There are two distinct and complementary traditions of explaining human behavior – a learning tradition and a decision-making tradition (see Michie, West, & Spring, 2013). The learning tradition, as its name suggests, focuses on the time-course of learning a behavior – in particular, the interdependence among behavioral occasions in a sequence rather than the exact determinants of each occasion. In contrast, researchers in the decision-making tradition care more about what factors determine a behavior on specific occasions, and what information is processed at such moments, but much less on how repeated decisions are related. As an example, to understand and change someone’s obsession with fast food, learning researchers would study how the rewarding experience of eating fast food (e.g., taste and energy) leads to stronger tendencies to repeat the behavior in the future, and how this reinforcement loop can be broken. To tackle the same problem, decision-making researchers would instead search for the determinants of either a particular decision of choosing fast food (e.g., hunger, time-pressure, etc.) or the general tendency of having this lifestyle (e.g., innate preference for fat, social environments, etc.).

As both learning and decision-making aspects are crucial for developing a dynamic framework, the review below is organized based on the theories’ roots in either tradition. After the review, a brief discussion is provided on the temporal scales implied in the theories – whether they focus primarily on explaining individual daily decisions (action level) or episodic behavioral processes (reflection level).

2.2.1 Theories in the learning tradition

Reinforcement learning theory

Reinforcement learning, or learning by outcomes, is a fundamental form of learning discovered in the early age of modern psychology (for a review, see Postman, 1947) and is still influential in today's behavioral and brain sciences (Yin & Knowlton, 2006). The success of reinforcement learning theory is also reflected in the development of reinforcement learning algorithms in the field of artificial intelligence to solve practical problems, where the development was greatly inspired by the theory in psychology (Sutton & Barto, 1998). Humans, animals, and also artificial agents are theorized to adapt their behaviors through their interactions with the changing environments in order to achieve their goals (e.g., finding food, avoiding predators, or to schedule digital interventions). Basically, if a behavior results in goodness to an organism, the frequency of performing the same behavior increases; conversely, if a bad outcome follows, the behavior will be performed less often in the future. This was summarized by Thorndike (1932) as the *law of effect*.

Reinforcement theory becomes more complex when one also considers the *law of exercise* (Thorndike, 1932). The above-mentioned response-outcome (R-O) learning, or goal-directed learning, is accompanied by stimulus-response (S-R) learning, also known as a process of habit formation (see Yin & Knowlton, 2006). The distinction between goal-directed learning and habit learning has been demonstrated in instrumental learning experiments where animals or humans are trained to acquire reward-generating responses (e.g., press a lever to receive food): when a response is overly trained, it persists to be triggered by the corresponding stimulus even when the reward becomes goal-irrelevant (e.g., when a rodent is satiated) (e.g., Adams, 1982; Dickinson, 1985; Yin, Knowlton, & Balleine, 2004). The recent resurgence of interest in habit formation in social and health psychology also follows the theory to define habits as mental associations between behaviors and environmental cues (Wood & Neal, 2007; Wood & Runger, 2016). When a behavior becomes strongly habitual, goal-related determinants of behavior, such as attitude and intention, cease to influence behavior (Gardner, 2015).

Control theory of self-regulation

The classical reinforcement learning theory focuses on the role of external immediate rewards in controlling behavior, but neglects the role of distal behavioral outcomes that may be cognitively represented. Following criticisms on this bias (e.g., Kanfer & Karoly, 1972), the control theory of self-regulation assumes that people can mentally represent distal outcomes of goals, and the regulation of behavior is generally towards reducing the discrepancies between the goals and their current status (Carver & Scheier, 1982; Powers, 1973). When a behavior leads to a reduced discrepancy, the reduction itself becomes a reinforcer to motivate the behavior, just like external rewards. This discrepancy-reduction mechanism is analogous to

feedback control systems in engineering, where discrepancy between perceived status and a reference value is constantly monitored to maintain homeostasis.

Control theory also represents goals and self-regulation hierarchically. Carver and Scheier (1982) proposed a 9-level hierarchical control system in which a behavior output from a higher level serves as the goal reference to the next level below. For lifestyle behaviors, it is sufficient to consider three levels⁵ – long-term goals (e.g., improving health), short-term goals (e.g., walking ten-thousand steps a day), and actions (e.g., taking a specific walk). Taking actions lead to fulfillment of short-term goals, which in turn bring one closer to the adherence to long-term objectives. Self-regulation operates most frequently at the action level (i.e., making daily decisions), but people's attention can be shifted to higher or lower levels. Downward shifting happens when lower-level motor-control, which is normally highly automated, becomes temporarily impeded during action executions (e.g., when learning a new motor skill or when a dysfunctional action needs to be inhibited; see Norman & Shallice, 1986). Upward shifting can be understood as self-reflective moments when a person reconsiders the attainability of a higher-level goal, which is more difficult to predict (but see Psarra, 2016).

Social Cognitive Theory

Social Cognitive Theory proposed by Albert Bandura is one of the most cited and applied theories in behavior change research (Davis et al., 2015; Webb et al., 2010). The theory encompasses three key concepts – *social learning* (Bandura, 1971), *self-efficacy* (Bandura, 1982), and *proactive control* (Bandura, 1989). First, based on extensive research on children's learning behaviors (e.g., Bandura & McDonald, 1963), *Social Learning Theory* posits that behaviors or attitudes are not only acquired through direct reinforcement, but also by observing the behaviors and their corresponding consequences to others (Bandura, 1971). For many health-related behaviors, the long-term health consequences are often learned by observing other people's behavioral outcomes. Second, based on organizational decision-making research (e.g., Bandura & Wood, 1989), it was found that the subjective belief in one's ability to perform a behavior was closely related to actual performance. Relating to the control theory above, this self-efficacy belief can be understood as a cognitive mechanism that simulates a series of future actions (e.g., dinner choices every day) in an extended episode of goal-pursuit (e.g., adherence to a diet). If the mentally simulated actions fail to bring sufficient progress, a person may decide to abandon the goal-pursuit altogether. Third, Bandura was among the earliest scholars to discuss proactive control – a discrepancy-production process, in which a person sets higher goals to further motivate behavior (Bandura, 1989). It thus complements the discrepancy-reduction mechanism at the core of control theory. Relatedly,

⁵ These levels were termed *system concepts, principles, and programs* in Carver and Scheier (1982).

the idea that goals are susceptible to changes also allows the possibility to adjust an unattainable goal downwards to reduce its discrepancy to the current status. Altogether, the three concepts contribute to extend reinforcement learning and control theory by incorporating flexibility to complex human behaviors.

2.2.2 Theories in the decision-making tradition

Classical expected utility model

Across behavioral sciences (e.g., psychology and economics), many mathematical models have been developed to describe how people make choices given a fixed set of alternatives (options, e.g., fries and salad) and attributes (e.g., healthiness and tastiness). A fundamental theoretical idea behind many models is the expected utility theory, which assumes that people integrate multiple attributes – the choice alternatives’ values to satisfy various personal goals – into a unidimensional construct called subjective or expected utility, and then choose the alternative with the highest utility (Oppenheimer & Kelso, 2015). Formally, expected utility is computed as $EV = \sum_{j=1}^J \sum_{n=1}^{N_j} V(x_{jn}) \times P(x_{jn})$, where $V(x_{jn})$ is the subjective value function for the n^{th} possible value of attribute j , and $P(x_{jn})$ is the probabilistic belief that attribute j would take that value (Savage, 1954; von Neumann & Morgenstern, 1947). The equation implies that the expected utility of one choice alternative increases when choosing the alternative is likely to produce certain outcomes (large $P(x_{jn})$) and when the outcomes are highly valuable (large $V(x_{jn})$). For example, whether people choose salad over fries depends on both their beliefs about its benefits for health their valuations on good health. The theory does not imply that people always consciously follow the equation to compute utilities, but rather reflects key neural mechanisms that underline decision-making (see Busemeyer & Townsend, 1993). In reality, conscious and deliberative computations are more common for single-time important decisions (e.g., comparing different health insurance policies), but rarely for fast daily lifestyle decisions.

Sequential sampling models

Empirical data from choice experiments have repeatedly shown that people are less rationale than the classical choice models following the expected utility theory would suggest (for a review, see Bhatia, 2013). People are prone to be influenced or biased by information that is seemingly irrelevant, for example, the addition of an inferior choice option (e.g., Tversky, 1972) or framing of losses versus gains (Tversky & Kahneman, 1981). To account for the anomalies, a sequential sampling approach has been developed to model the cognitive process of decision-making dynamically, such as the *multialternative decision field theory* (Roe, Busemeyer, & Townsend, 2001) and the *associative accumulation model* (Bhatia, 2013). The new models share the idea that preferences for different choice alternatives are accumulated over time (e.g., a few seconds) and a choice is finally committed when its preference signal

exceeds a decision threshold the earliest. At each timestep, the preference signals of choice alternatives fluctuate according to a round of utility comparison based on one (e.g., Bhatia, 2013) or multiple attributes (as in drift diffusion models, e.g., Ratcliff & Rouder, 2001). The stochastic property of sequential sampling models enable them to explain the sensitivity of choices to subtle changes in choice sets and to predict decision time (Busemeyer & Townsend, 1993). Finally, sequential sampling models suggest a mechanism for habitual choices – repeatedly choosing an alternative may shift its starting position of preference accumulation towards a decision threshold at the baseline (Roe et al., 2001).

Reasoned action approach

Influenced by the expected utility theory (Fishbein & Ajzen, 1975) but with a strong focus on application, the reasoned action approach (see Noar & Head, 2014) has produced some of the most applied theories in behavior change research, such as the Theory of Planned Behavior (Ajzen, 1991; Ajzen & Madden, 1986) and the Health Belief Model (Janz & Becker, 1984). For a decision-making perspective, this approach categorizes attributes in possible choice situations into a smaller set of behavioral determinants that are generalizable to a wide range of behaviors and are measurable by self-report. For example, in the Theory of Planned Behavior, regardless of the specific alternatives and attributes considered, factors affecting choices are categorized into three determinants, namely attitude, social norm, and perceived behavioral control (Ajzen & Madden, 1986). When a specific behavior is considered (e.g., dinner choice), attitude towards a choice alternative is further determined by many attributes (see Ajzen, 1991), such as taste, nutrition, and price, while social norm is influenced by the perceived social consequences of choosing an alternative (e.g., presenting oneself to be environmental friendly). Perceived behavioral control, similar to self-efficacy, measures one's confidence in maintaining certain choices in the future. Ajzen (1991) explicitly considered his model as a model for behavioral prediction, rather than for explaining what processes underlie overt behaviors or decisions or how such processes can be intervened.

It is also worth noting that the reasoned action approach makes a strong assumption on the intentionality of behavior (see Karoly, 1993). For example, behavioral intention is a prerequisite to actual behavior in the Theory of Planned Behavior (Ajzen & Madden, 1986). Thus, this approach considers behaviors as “planned” or “intended”, and as results from careful deliberations on the pros and cons of certain behaviors. Such a theoretical position suggests that the reasoned action approach was developed to mainly deal with single-time decisions or the planning of behavioral episodes, rather than the “small” daily decisions.

Dual-processing models

A recurrent idea in psychology is that humans possess two distinct modes or systems for processing information and making decisions. Despite that different dual-system models use different terminologies (for a review, see Evans, 2008), it is widely accepted that one system is

fast, impulsive, and largely automatic, and the other system is slow, reflective, and deliberate (Kahneman, 2003).

The Reflective-Impulsive Model (RIM, Strack & Deutsch, 2004) is a representative of this approach, and it has been explicitly applied to health-related behaviors (Hofmann, Friese, & Strack, 2009; Hofmann, Friese, & Wiers, 2008). The reflective system hosts various higher order mental operations that rely on controlled processes and symbolic representations, including deliberate judgments, planning for goal pursuit, and the inhibition of prepotent responses. In contrast, the impulsive system operates fast on associative clusters in long-term memory that group stimuli, affective states and behavioral responses together. At the moment of a specific decision, whether self-control succeeds or not depends on the relative ability of the processes in the two systems to activate the corresponding behavioral schemas. Several boundary conditions have been proposed to moderate the relative strength of the two systems (Hofmann, Friese, & Wiers, 2008). For example, a person is believed to behave more impulsively, when the behavior is highly habitual, when their cognitive load is high, and when their mood is positive.

2.2.3 Temporal scales used in the above theories

Figure 2.2 summarizes learning and decision-making theories based on their temporal scales, or their levels of behavior representation. A similar distinction was made by Karoly (1993), where theories at the action level are called *online* theory while theories at the reflection level are called *offline* theories.

	Learning	Decision-making
Reflection level	<ul style="list-style-type: none"> • Social Cognitive Theory • <u>Control theory of self-regulation</u> 	<ul style="list-style-type: none"> • <u>Expected utility theory</u> • <u>Sequential sampling models</u> • Reasoned action approach
Action level	<ul style="list-style-type: none"> • Reinforcement learning theory • <u>Control theory of self-regulation</u> 	<ul style="list-style-type: none"> • <u>Expected utility theory</u> • <u>Sequential sampling models</u> • Dual-processing models

Figure 2.2 Categorization of reviewed theories based on their theoretical traditions and temporal scales (theories apply to both temporal scales are underlined).

In the learning tradition, the reinforcement learning theory clearly represents behavior at the action-level, since the outcome of each specific action/decision is modeled to have concrete impacts on the frequency of repeating the same action in the future. Reinforcement learning experiments also involve repeated trials within a relatively short period of time (e.g., a few hours). The control theory of self-regulation, due to its hierarchical structure, covers both behavioral processes at the reflection level and the action level. Social Cognitive Theory and particularly its processes of self-efficacy and proactive control apply mainly to behaviors at the reflection level. Although the two processes may have counterparts at a lower level as in the control theory, Bandura's focus was clearly on voluntary and deliberative human behaviors (Bandura, 1989).

In the decision-making tradition, mathematical models as part of the expected utility theory and the sequential sampling approach can be equally applied to decisions at both temporal scales, as long as decisions with clearly defined choice sets are considered. As discussed above, theories in the reasoned action approach deal mainly with decisions at the reflection level because of its assumption on intentionality. In contrast, as dual-process models are intended to account for "small" daily decisions, for which both reflective and impulsive processes play a role.

2.3 Theory integration: an adaptive decision-making framework

To reiterate, our goal of theory integration is to develop a unified framework that includes and connects all relevant decision-making and learning processes at both the action-level and reflection-level of lifestyle behavior change. Most processes in the framework came directly from the theories reviewed, but efforts were made to unify terminologies from different theories to form a coherent system, and to tailor the theoretical system to lifestyle behaviors. Taking dietary behavior as a primary example, the framework should explain not only how daily meal choices are made and how each decision outcome influences future choices, but also how a goal of adhering to a specific diet is made and how such goals are evaluated. The following sections introduce the adaptive decision-making framework in four parts: action-level decision-making, action-level learning, reflection-level decision-making, and reflection-level adaptation.

2.3.1 Action-level decision-making: daily meal choices

Daily lifestyle decisions, such as daily meal choices, can be modeled as a two-step process – first *option generation* and then *option evaluation* (see Figure 2.3). Thus, the framework assumes that when choosing a meal, different meal options have to be generated or recalled by a decision-maker first, before evaluations of a few options can be made to inform a final choice (Kamphorst & Kalis, 2015; Tobias, 2009). The notion of option generation has not

been examined in any of the decision-making theories reviewed, probably because those theories are based on laboratory choice experiments, where options are simply provided by the experimenters. For lifestyle behaviors in daily environments, how choice alternatives are generated is an important question.

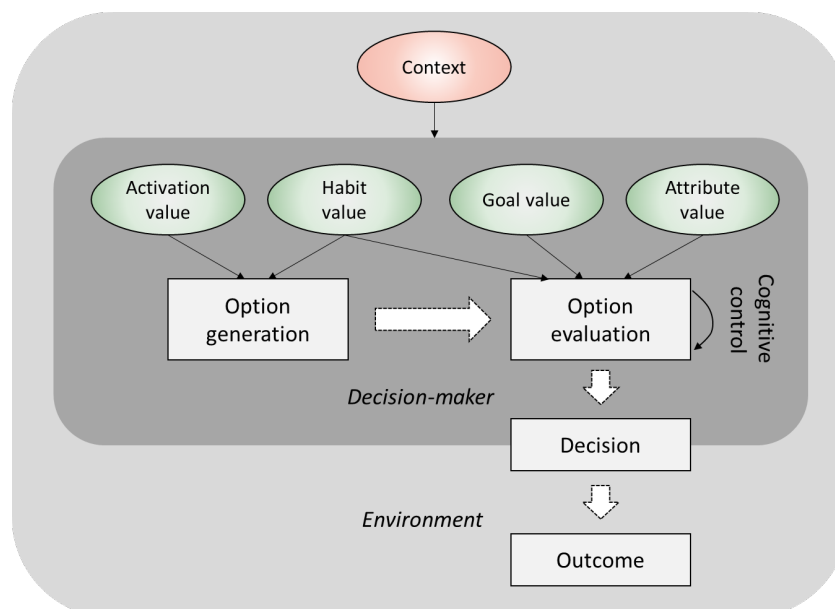


Figure 2.3 A two-step model of daily lifestyle decision-making.

In general, behavioral options can be generated through three different means. First, if an option is habitual, it will be activated when the associated cues are encountered, such as location and time (e.g., lunch at the office) or a complex combination of contextual cues (e.g., a busy Wednesday evening). This follows directly from the construct of habit in reinforcement learning theory. Second, options may be remembered at the right moments because people intentionally try to maintain them in their prospective memory (i.e., not to forget to do something in the future, McDaniel & Einstein, 2000; Tobias, 2009). This usually happens when there is a salient goal guiding daily decisions, such as a goal of adhering to a low-carb diet. People may also intentionally associate important options with external cues, so that encountering the cues is likely to trigger the options (see Gollwitzer, 1999). Third, options can be triggered by direct external suggestions at the decision moments, for example, a coaching message from a mobile health app to recommend health foods (Kamphorst & Kalis, 2015). Through these means, behavioral options that are sufficiently activated (e.g., by passing an activation threshold) would be later evaluated.

Option evaluation can be modeled as a process of comparing several options and then choosing the one with the highest goal-satisfying value. The exact computation of utilities can follow either classical expected utility models or the more dynamic sequential sampling models, but for the framework it is sufficient to identify three main cognitive determinants of the evaluation process. First, when multiple personal goals are relevant for a daily decision, these

goals can be regarded as more or less importance by a decision-maker, thus entailing higher or lower *goal values*. For example, between the goals of keeping health and enjoying delicious food, a person who regards the former goal as more valuable would be more likely to choose food options for meals that satisfy the health goal.

Second, for each personal goal, a behavioral option has its *perceived attribute value* relating to that goal, which determines the total utility of the option. These attribute values are subjective beliefs held by people about the causal relationships or contingencies between choosing certain behavioral options and the realizations of personal goals. While goal values are relatively stable within-person, attribute values are more context-dependent and are prone to changes through learning and experience. For example, the perceived tastiness of a particular meal option may depend on a person's momentary appetite and it may change over time through repeated tasting of the food.

There is a particular challenge for healthy decisions, such as healthy food choices, because usually two distinct types of attributes are considered – an immediate hedonic aspect such as tastiness, and a long-term consideration of health consequences. This challenge is essentially a problem of self-control from a decision-making perspective (see Berkman, Hutcherson, Livingston, & Inzlicht, 2017). According to the idea of temporal discounting in decision-making theories (Frederick, Loewenstein, & O'donoghue, 2002; Green & Myerson, 2004), because any rewards from potential health improvements are delayed in time when compared with the immediate hedonic aspects, the value of the attribute healthiness is discounted before it is integrated in option evaluation (Chapman, 1996; Story, Vlaev, Seymour, Darzi, & Dolan, 2014). Another reason why health aspects are often weighted less than hedonic aspects in actual decisions is that the former are more abstract concepts, so they might be more difficult or take longer to be processed (Maier, Raja Beharelle, Polanía, Ruff, Hare, 2018; Sullivan & Huettel, 2018; Sullivan, Hutcherson, Harris, & Rangel, 2015). Finally, from a dual-processing perspective (e.g., Hofmann, Friese, & Wiers, 2008), dietary self-control may sometimes succeed because people can voluntarily exert top-down cognitive control on the option evaluation process, especially if a momentary preference for a meal option conflicts strongly with a diet goal. It has been shown experimentally that cognitive control may either modulate the valuation process to be more in favor of healthiness rather than tastiness (Hare, Camerer, & Rangel, 2009), or filter people's attention away from hedonic attributes in the early stage of option evaluation (Harris, Hare, & Rangel., 2013). Of course, top-down control depends on many contextual variables to be effective, such as motivation (Inzlicht & Schmeichel, 2012), mental fatigue (van der Linden, Frese, & Meijman, 2003), stress level (Chajut & Algom, 2003), and daily affective states (Zhang, Smolders, Lakens, & IJsselsteijn, 2018).

Third, *habit values* or habit strengths, which represent the history of choosing certain behavioral options, may influence the evaluation of options. As discussed in the theory review, learning experiments have shown convincingly that even when two options are provided to

decision-makers, habitual options have more chances to be chosen than non-habitual options (e.g., Dickinson, 1985; Yin & Knowlton, 2006). In sequential sampling models, the influence of habits on the dynamic process of option evaluation can be understood as positively biasing the baseline preferences for the habitual options (Roe et al., 2001). As an intuitive example, if someone often chose fast food in the past, fast food is by default more favorable than other options when no additional deliberations are made.

2.3.2 Action-level learning: developments of eating habits

Action-level learning processes can be added to the framework by integrating the ideas of goal-directed learning and habit learning from the reinforcement learning theory to the two-step decision-making model proposed above (see Figure 2.4). First, feedback from decision outcomes to perceived attribute values represents goal-directed learning. For example, when a new canteen is built at a workplace, employees may have initial but very uncertain beliefs about the tastes and calories of different lunch options, but after a few weeks of trying it out, they gradually form more certain and accurate perceptions about the qualities of the food. Computationally, the updates of perceived attribute values can be done through model-based and model-free reinforcement learning algorithms (e.g., temporal difference learning, Sutton & Barto, 1998), or Bayesian belief update (Russo, van Roy, Kazerouni, Osband, & Wen, 2018). For health-related attributes, because concrete decision outcomes are infrequent, it is less clear how direct learning from experience works, if it's possible after all (but see Gershman & Daw, 2017). People's beliefs about the health consequences of different food are more susceptible to social learning and education.

Second, there is direct feedback from decisions themselves to habit values, as in a process of habit formation or habit learning. Although daily lunch decisions in a new canteen are driven primarily by goal-related attribute values, through repeated decisions, mental associations between frequently chosen food options and the environment cues (e.g., the physical setting of the canteen, lunch time) are gradually strengthened. These associations, as habit values, then influence future decisions through both the option generation and option evaluation processes discussed above. The exact mechanism of habit learning is beyond the scope of this chapter, but it has been modeled using algorithms inspired by Hebbian learning (Hebb, 1949) in the literature (e.g., Klein, Mogles, Treur, & van Wissen, 2011; Miller, Shenhav, & Ludvig, 2019; Psarra, 2016; Tobias, 2009).

Besides the two main learning processes, there is also a direct link from decisions to the activation values of options, which has been discussed much less in the learning literature. When a decision is made and the corresponding behavior is executed, the behavior execution increases the activation level of the behavioral option in memory, even though such an increase has been shown empirically to be very small (Tobias, 2009). As discussed in the section about

2.3 Theory integration: an adaptive decision-making framework

option generation, the dynamics of activation values are primarily memory processes and are mostly affected by physical and social stimuli in the environment.

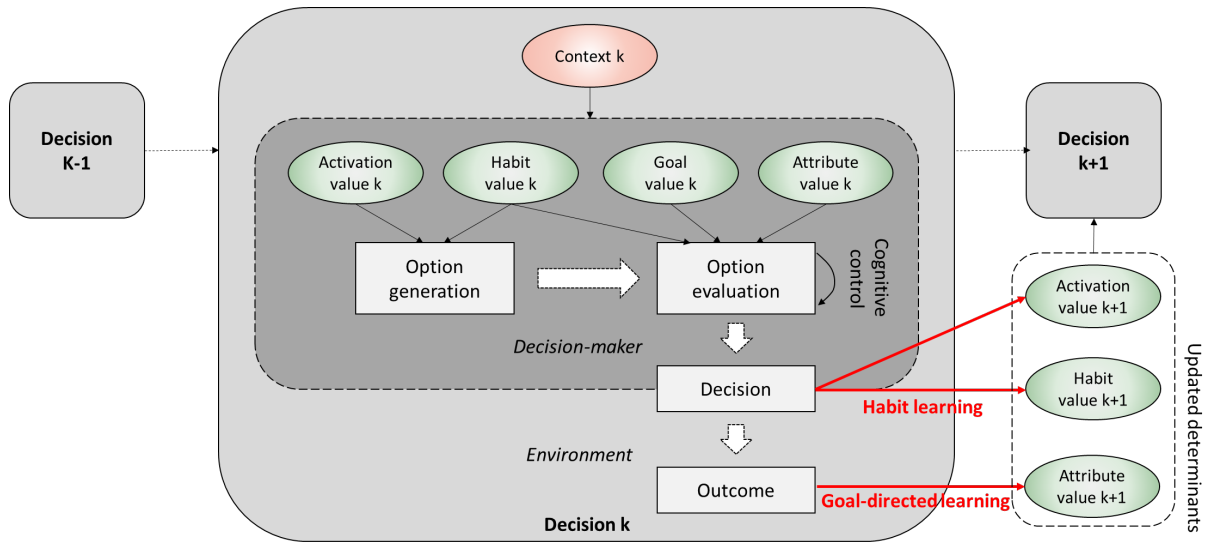


Figure 2.4 Action-level learning processes added to the decision-making model.

2.3.3 Reflection-level decision-making: dietary goal-setting

The action-level decision-making and learning processes cover a substantial part of what people do in their daily lives. However, the framework as it is now gives the impression that lifestyle behaviors are mindless, and without much thinking or purpose. In fact, people do periodically have moments when they reflect on their health status, contemplate about possible improvements, and make action plans. According to the control theory (Carver & Scheier, 1982), what connects people's abstract long-term goals (i.e., what they want or strive towards in their lives) and daily "small" decisions are the more concrete short-term goals they set. Short-term goal-setting can be understood also as a process of decision-making, albeit at the reflection level rather than the action level. The decisions made are commitments to goals that guide future daily decisions, rather than overt behaviors that trigger motor programs.

Thus, the two-step model of daily lifestyle decisions generally applies to the setting of short-term goals. In selecting a dietary plan, for example, people first search for diet candidates that serve their long-term goals in their memory as well as external sources (e.g., dietary books or mobile e-coaching systems). A few candidates are then evaluated based on relevant attributes, such as taste, ease of preparation, and expenses. At the reflection-level, often these attributes can indeed be categorized into a few determinants, such as attitude, social norm, and perceived behavior control (Ajzen, 1991). There are also a few more distinct features for the decision-making process concerning short-term goals. First, because goal-settings occurs at a much lower frequency than daily decisions, habits are less likely to be formed and to influence the decision-making process. More specifically, the search for goal candidates tends to be more thorough, and the longer evaluation time also reduces any habitual bias. Second,

because goal-setting aims for an extensive period of time, and possibly requires a more abstract mental construal (see Trope & Liberman, 2010), it is largely detached from direct sensory information and visceral attributes, such as effort and tastiness. Thus, the self-control problem is less prevalent for goal-setting than for daily lifestyle decisions. Third, self-efficacy plays an important role for goal-setting. People may carefully consider the feasibilities of different diet goals by mentally simulating a series of daily dietary choices in the future.

Motivating functions of short-term goals

When short-term goals are generated, they can influence daily lifestyle decisions through both option generation and option evaluation (see Figure 2.5). First, setting up a short-term goal can increase the activation values of desirable behavioral options through a process termed *planning*. Planning can be done through two mechanisms discussed earlier: an effortful prospective memory process (e.g., rehearsing eating salads; see McDaniel & Einstein, 2000), or an *implementation intention* process, i.e., mentally associating a behavioral option with certain environmental cues (e.g., eating an apple as snack when watching TV; see Gollwitzer, 1999). Second, compared with long-term goals (e.g., to improve health), short-term goals are more concrete, so complying with these goals brings immediate satisfaction (Locke & Latham, 2002). The goal-compliance satisfaction functions as an additional attribute that compete with other hedonic attributes (e.g., tastiness) during the option evaluation process.

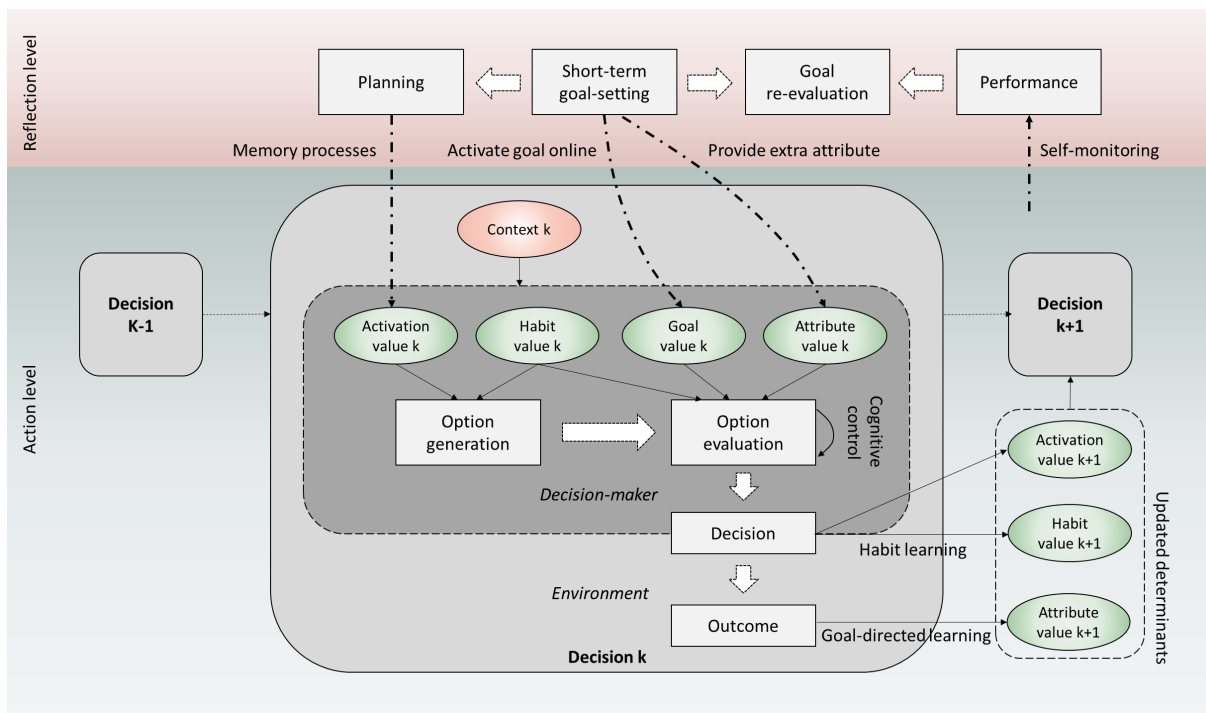


Figure 2.5 A full representation of the adaptive decision framework (reflection-level processes and interactions between the two levels added to the previous decision-making model).

2.3.4 Reflection-level adaptation: self-monitoring and re-evaluation of dietary goals

Short-term goals have to be re-evaluated in reflection moments periodically, in order to change the goals that are no longer adaptive (e.g., too difficult, too easy, or irrelevant). Such reflection-level adaptation processes are well-described by the control theory (Carver & Scheier, 1982) and the Social Cognitive Theory (Bandura, 1989). A goal re-evaluation first requires inputs from the action-level processes through *self-monitoring*. Through repeated daily dietary choices, not only are dietary habits formed, but also actual choices and their outcomes are stored in episodic memory and later retrieved and integrated into a mental representation of overall performance over a past time period (see Figure 2.5). Next, discrepancy between a goal reference (e.g., a dietary plan) and a performance representation is computed and used to inform the reflection level adaptation. Depending on the size of the discrepancy and other personal and contextual factors, a few outcomes are possible. First, a people may further employ the motivating functions of dietary goals to promote healthy daily food choices, in order to reduce goal-performance discrepancies. However, when a discrepancy is deemed to be too large, people may instead lower the goal standard (e.g., be less strict on calorie intake) or abandon the goal altogether (e.g., give up a diet). Finally, when performance matches or even exceeds a current goal standard, they may proactively adjust the goal standard upwards to further improve health (see Bandura, 1989).

2.4 Mapping digital intervention techniques to the framework

We refer to *digital intervention techniques* as behavior change techniques that target lifestyle behaviors and are implemented in digital systems (e.g., web, mobile, or wearable systems). Behavior change techniques, in turn, are generally defined as the active ingredients of interventions that can influence behaviors in desirable ways (Abraham & Michie, 2008; Michie, Richardson et al., 2013). In this section, in order to demonstrate the relevance of our framework to digital lifestyle interventions, we map common digital intervention techniques to the adaptive decision-making framework. Specifically, these techniques are categorized and interpreted according to the behavioral processes and cognitive variables in the framework being targeted. Strengths and limitations of some techniques are discussed based on the implications of the framework.

2.4.1 Digital intervention techniques targeting action-level decision-making

Because digital systems are prevalent in people's daily lives, they are well-positioned to influence people's daily lifestyle decisions in the decision moments. The ability to target action-level decisions is indeed considered by many as a promising direction for digital lifestyle interventions, as reflected in research on *ecological momentary interventions* (Heron & Smyth, 2010) and *just-in-time adaptive interventions* (JITAI; Hekler et al., 2016; Intille et al., 2003; Jaimes et al., 2015; Nahum-Shani et al., 2015; Riley et al., 2015). According to our framework,

there are many different ways that digital intervention systems can influence online decision-making processes, depending on whether the techniques target option generation or option evaluation, and which cognitive variables are targeted (e.g., activation, attribute, or goal value). Four main categories can be distinguished.

Option-based techniques

Option-based techniques make certain desirable behavioral options salient, but leave the option evaluation process completely to the users themselves. When a desirable behavior is obvious but may not be constantly salient to users, digital systems can simply prompt users to actively make decisions to engage in a desirable behavior, for example, to take exercise breaks when overly sedentary behaviors are detected by the system (Pina, Ramirez, & Griswold, 2012; Thomas & Bond, 2015). In some other cases, it might be possible to provide users with new options that are better than the ones known by the users (Kamphorst & Kalis, 2014). Finding such “attractive” options relies on a system’s sensor network and smart algorithms, which potentially make it more knowledgeable than its users in a given behavioral domain and/or context. For example, to promote physical activities, Guo (2016) developed a system which recommends new commuting routes to users “in-situ”, based on automatic detection of users’ habitual existing routes and Google map data.

Attribute-based techniques

Attribute-based techniques aim at changing users’ beliefs about the attribute values of options by providing health-related knowledge or facts. They are referred to as *providing information about behavior-health link* or *providing information about consequences*, in the taxonomy of behavior change techniques (Abraham & Michie, 2008). Given the common assumption that humans are rational decision-makers, providing information about attribute values has been considered as a logical approach to behavior change in traditional health education campaigns. However, our framework implies that attribute value is only one of many factors that influence option evaluation and providing information alone may not always change people’s beliefs about attribute values. Nonetheless, information about attribute values can be provided to justify the recommendations of behavioral options whenever appropriate in digital systems (e.g., calorie information for different meal choices).

Goal-based techniques

Because goal values modulate attribute values in option evaluation, activating health-related goals in the decision moments provides yet another type of intervention techniques. When implemented in digital systems, this technique links the suggestions of concrete behavioral options with the reminder of the associated short-term or long-term goals. For example, when a mobile application prompts a user to take a lunch walk, the user’s goal of walking 10,000 steps a day (and the achieved steps) can be presented along with the option of taking a lunch walk.

Structure-based techniques

Structure-based techniques differ from the previous types because they do not change the availability of options nor do they alter the existing payoffs (i.e., attribute values). Because they are less susceptible to user reactance and require less processing efforts from users than other techniques, structure-based techniques have attracted strong research interests (e.g., Adams, Costa, Jung, & Choudhury, 2015; Schneider, Weinmann, & vom Brocke, 2018), usually under the header of nudging or choice architecture (Johnson et al., 2012; Thaler & Sunstein, 2008). As examples, *default* and *context effects* (e.g., compromise effect) are two widely used structure-based techniques. Lee and colleagues (Lee, Kiesler, & Forlizzi, 2011) adopted the default technique to promote healthy snacking in an online environment by making healthier options the default choices. Zhang and colleagues built on the compromise effect to promote physical exercise at work: Intensive exercise options were added in an app called BeActive! to make the moderate exercise options to be perceived as compromise options and thus as more attractive (Zhang, Starczewski, Lakens, & IJsselsteijn, 2018).

In general, digital interventions for daily lifestyle decisions face a challenge that people are often triggered by cues in the environments to make these decisions spontaneously, rather than through the mediation of digital systems. In other application domains, such as e-commerce, people are accustomed to shopping online, so e-commerce systems do not need to worry about missing intervention opportunities. To intervene lifestyle decisions at the critical moments, interfaces between the information in the digital systems and people's spontaneous behaviors in the physical world might need to be created. Current approaches include predicting users' spontaneous decision-making moments using sensor network (e.g., predicting "about-to-eat" moments, Rahman, Czerwinski, Gilad-Bachrach, & Johns, 2016), and initiating decisions when interventions are predicted by the system to be most valuable (e.g., predicting stressful moments, Jaimes, Llofriu, & Raij, 2014). This challenge will continue to stimulate new intelligent digital solutions and at the same time debates on the associated ethical implications (e.g., Zuboff, 2019).

2.4.2 Digital intervention techniques targeting action-level learning

Digital intervention techniques targeting action-level learning processes operate in between rather than at decision moments. The goal is to support either goal-directed learning (e.g., to change perceived behavior-health links) or the formation of healthy habits. If these techniques are effective, the cognitive variables that influence decision-making will be in a health-promoting state, so that users are expected to maintain the learned healthy behaviors without being continuously intervened by digital systems at a daily basis.

A main challenge for lifestyle behavior change is learning the causal relationships between one's behaviors and health consequences, as these consequences are usually delayed. As discussed, researchers have speculated on the role of episodic memory in tracking internal and

external events to support this type of learning (Gershman & Daw, 2017). In this regard, the self-tracking function of many digital systems can support learning by externalizing user' memory systems (Kersten-van Dijk et al., 2017). Behavioral and contextual data can be objectively recorded and can be reviewed later by users when consequential health events take place. Because self-tracking studies mostly focus on the effectiveness of the technology as a whole, evidence regarding its specific role in supporting goal-directed learning is lacking (Kersten-van Dijk et al., 2017). Some interview data indicated that users of self-tracking systems believed that they acquired knowledge about behavior-health links through the self-tracking technology (e.g., Choe, Lee, Lee, Pratt, & Kientz, 2014; Li, Dey, & Forlizzi, 2011).

Instead of directly supporting the learning of health consequences, a different and popular approach is to provide extra rewards that may reinforce the desirable behaviors. In such gamification systems, most common extra rewards are virtual rewards, such as badges, trophies, or complements (Cugelman, 2013; Deterding, Dixon, Khaled, & Nacke, 2011; Nacke & Deterding, 2017). These virtual rewards are expected to steer users to healthy behavioral options by competing with the inherent hedonic values of many unhealthy behaviors.

Despite its popularity, the effectiveness of virtual rewards in lifestyle behavior change is questionable, as empirical studies found no positive effects in several health domains, such as physical activity (Zuckerman & Gal-Oz, 2014) and sexual protection behavior (DeSmet, Shegog, van Ryckeghem, Crombez, & de Bourdeaudhuij, 2015), while users in a study perceived such virtual rewards implemented in an exercise-promoting application as “not motivating” or even “unnecessary” (Munson & Consolvo, 2012). Our framework implies that the problem with virtual rewards is not in the learning of the contingencies – given sufficient learning, users would know what behaviors are rewarded – but in the corresponding goal values of these rewards: the goal values of virtual rewards are often low, when compared with other hedonic attributes, such as tastiness and reduced effort. Future research on gamification should focus on making the virtual rewards more goal-relevant, for example, by embedding them as a game mechanic that users care about (Berkovsky, Coombe, Freyne, Bhandari, & Baghaei, 2010), or by making the rewards socially meaningful (Shahrestani et al., 2017).

Another technique in this category is habit formation support, usually by reminding users about a new and desirable behavioral option. This is especially valuable at the beginning of habit formation when new options are not always remembered by users themselves. Unlike the technique of suggesting options at decision moments (see last section), reminders that support habit formation are sent offline and according to time-based schedule (e.g., once every morning). They do not persuade users to act immediately, but to increase the activation values of certain options so they are more likely to be generated when decision moments arrive. Reminders have been widely used and have been shown to be effective in domains where forgetting is the main obstacle for behavior change (e.g., Armstrong et al., 2009; Thakkar et

al., 2016). More research is warranted to understand its value in changing more complex lifestyle behaviors, when activation value is one of many influential cognitive variables.

2.4.3 Digital intervention techniques targeting reflection-level decision-making

Setting up a short-term goal, as a reflection-level decision-making process, is often the starting point of behavior change, either by people themselves or supported by digital systems. Without external interventions, goal-setting is only triggered under special conditions, for example, when someone has learned new health-related knowledge (e.g., become aware of the risk of smoking) or has experienced a sudden change of their health status (e.g., being diagnosed with diabetes). Thus, a straightforward intervention technique is to proactively prompt users to set up new goals to improve their lifestyles. In many digital systems, following a goal-setting prompt, a user can choose a goal and then record it in the system, which allows the system to remind the user of the goal when needed.

As goal-setting is a decision-making process, most techniques discussed in the section of targeting action-level decision-making also apply to goal-setting, including option, attribute, and structure-based techniques. As a particularly promising direction, digital systems may utilize their data-gathering power and artificial intelligence to recommend novel and attractive options for short-term goals (Kamphorst & Kalis, 2015). To address the subtlety and complexity of goal-setting in the health domain, the systems need to personalize options based on users' abilities (e.g., Radha, Willemsen, Boerhof, & IJsselsteijn, 2016) and also on their unique life experiences (see Rutjes, Willemsen, & IJsselsteijn, 2019). In the future, the difficult task of setting up challenging, motivating, yet realistic goals may indeed be transferred from people to intelligent intervention systems.

After the step of goal-setting, digital systems can go further to support the planning phase that connects short-term goals to daily decisions in the future. A simple technique is to prompt users to make concrete plans in the system, for example, by adding activities to a calendar. Data provided by the users allow digital systems to check adherence and send reminders when necessary. In addition to this time-based planning technique, digital systems may also encourage users to use the event-based planning technique of implementation intention discussed earlier (Gollwitzer, 1999). Implementation intention has been shown to be effective in the health domain (e.g., Adriaanse, Vinkers, de Ridder, Hox, & de Wit, 2011; Luszczynska, Sobczyk, & Abraham, 2007), and it has also been implemented in digital interventions where no human instructions are required (e.g., Pinder, Vermeulen, Wicaksono, Beale, & Hendley, 2016; Stawarz, 2017). A recent system even uses sensor data to automatically generate "if-then" rules that were adapted to the living contexts of individual users (Dogangün, Schwarz, Kloppenborg, & Le, 2017).

2.4.4 Digital intervention techniques targeting reflection-level adaptation

In this category, providing behavior feedback to users to support self-monitoring is the most used behavior change in digital systems (Conroy et al., 2014; Lehto & Oinas-Kukkonen, 2011; Mercer, Li, Giangregorio, Burns, & Grindrod, 2016; Payne, Lister, West, & Bernhardt, 2015; Zhao, Freeman, & Li., 2016). Technically, with the development of increasingly powerful sensors, digital systems are able to track lifestyle behaviors and related variables more accurately and in greater details than people's own memories. Moreover, these systems can transform the rich raw data into numerical or visual information (e.g., weekly summary of step count) to facilitate better comparison with the short-term goal references (e.g., Kersten-van Dijk et al., 2017).

Although self-monitoring as a general behavior change technique has been identified as effective (Michie, Abraham, Whittington, McAteer, & Gupta, 2009), the evaluation of this technique in digital systems have yield mixed results (e.g., Hermsen et al., 2016; Zhao et al., 2016) and is impeded by the lack of high-quality studies and a lack of focus on self-monitoring per se (Kersten van Dijk et al., 2017). At least it is evident that the abundance of self-tracking devices has not solved the problem of lifestyle behavior change. From an evolutionary perspective, since people's natural self-monitoring function has existed long time before the emergence of digital systems and quantitative data, it is not guaranteed that the technology-enhanced information can lead to better functioning. A recent study indicates that some self-tracking users may have an exaggerated focus on numeric feedback as the replacement of bodily experience as feedback, potentially leading to negative consequences such as rumination (van Dijk, Westerink, Beute, & IJsselsteijn, 2015). The bottom line is that even if technology-enhanced self-monitoring is beneficial to some extent, our framework implies that it is only one step in reflection-level adaptation. Future research should investigate how digital systems can also support the reflective processes that immediately follow self-monitoring, including the comparison between goal references and monitored performance, and the adjustments of goals and behaviors.

2.5 General discussion

Understanding and changing lifestyle behaviors in the digital age require a theoretical perspective that combines decision-making and learning, and a representation of behavior at the level of both daily decisions and episodic reflections. These two requirements have guided our review of individual theories and their integration, and the outcome is temporally fine-grained, dynamic, and process-oriented theoretical framework of lifestyle behavior change. Through a mapping exercise, we also linked common digital intervention techniques to the behavioral processes and cognitive constructs in the framework.

2.5.1 Theoretical contributions and comparisons with previous integration works

A primary objective of developing the adaptive decision-making framework was to address the mismatch between theory and digital intervention in terms of their temporal granularity (Riley et al., 2011). This was done by considering lifestyle behaviors at two different time scales – one that represents the individual daily decisions or actions, and another that groups the repeated daily decisions into a larger episode and incorporates self-regulatory processes (e.g., goal-setting, self-monitoring). This two-level representation contrasts our framework with previous integration attempts that were based on the stage model of change (Prochaska & Di-Clemente, 1982), such as the COMBI model (Klein, Mogles, & van Wissen, 2014) and the i-Change Model (de Vries, 2017). While the COMBI and i-Change Model postulate a more general process of behavior change (e.g., through contemplation, preparation, action, and maintenance), our framework zooms in to explain how repeated daily actions, with the help of reflection-level regulatory processes, lead to maintenance. Thus, our framework complements earlier work and contributes uniquely to the research on digital interventions, for which time-intensive intervention is a main strength.

There are earlier frameworks that are more similar to the adaptive decision-making framework when it comes to behavior representation. Both PRIME theory (West, 2006) and Temporal Self-Regulation Theory (Hall & Fong, 2007) model behavior change as a continuous process rather than a series of discrete stages. However, our framework is the first to explicitly distinguish the two distinct levels of lifestyle behaviors and the different time scales involved. Although a two-level representation is arguably a simplification, it helps to elaborate the behavioral and cognitive processes at the two different levels, and more importantly, allows researchers to include interactions between the levels – for example, the top-down motivating impacts of short-term goals on daily decisions, and the bottom-up process of self-monitoring to facilitate re-evaluation of goals.

A second limitation in the current literature we addressed is the lack of dynamic processes in traditional behavior change theories (Riley et al., 2011). By integrating theories from both the learning and the decision-making tradition, the resulted adaptive decision-making framework depicts a dynamic bidirectional relationship between behaviors and cognitive constructs that influence the behaviors. Our framework thus complements previous frameworks that focused exclusively on learning processes, such as the framework of evolutionary learning processes (Crutzen & Peters, 2018) and Action Change Theory (Vlaev & Dolan, 2015). More broadly, we believe that the need to capture the complexities of lifestyle behaviors for designing better digital interventions provides a strong and timely motivation to integrate decision-making and learning theories in basic psychological research (see Hastie, 2001). In this respect, the adaptive decision-making framework not only adds value to theoretical thinking in behavior change, but also to psychology in general.

Furthermore, our work may stimulate some reconsideration of the popular dual-processing models, in which health behaviors are assumed to be driven by two distinct forces, one reflective and one impulsive (Hofmann, Friese, & Wiers, 2008). In our view, such dichotomous categorization of diverse processes and constructs might be too coarse for a full understanding of the dynamic lifestyle behavior change process. The adaptive decision-making framework instead suggests several dualities. First, there is a contrast of long-term health benefits versus immediate hedonic rewards in option evaluation, where effortful cognitive control is required to battle one's impulses. Second, the goal-directed evaluation based on attributes (both long-term and short-term) competes with the influences from habits. This has been discussed extensively in the learning literature as the dual action control by goals and habits (e.g., Dolan & Dyan, 2013). Third, the faster processes at the action level can certainly be contrasted with the more thorough processes at the reflection level, but one should also realize that they operate on very different time scales.

Finally, our framework focuses strongly on behavioral processes and cognitive mechanisms. Connections between theoretical constructs in the framework are meant to represent causal mechanisms rather than statistical relationships as in the COMBI model or the i-Change Model. These processes or mechanisms are described at a level of specificity that they can be transformed to computational models by introducing additional assumptions and formal algorithms. For example, key cognitive variables are defined for option evaluation, but the exact computational process of these variables is left open to different assumptions (e.g., Roe et al., 2001; Usher & McClelland, 2001). Similarly, although the framework acknowledges the joint influence of habit and goal-directed control, the exact arbitration between the two is subjected to different computational accounts (e.g., Daw, Niv, Dayan, 2005; Keramati, Dezfouli, & Piray et al., 2011; Miller et al., 2019). The process-oriented nature of our framework makes it an ideal scaffold to develop new dynamic computational models envisioned by many researchers (e.g., Hekler et al., 2016; Nilsen & Pavel, 2013; Riley et al., 2011; Spruijt-Metz et al., 2015).

2.5.2 Added value to the synergy between theory and digital intervention

The adaptive decision-making framework was developed with the aim to bridge the gap between behavior change theories and digital intervention applications. As a modest first step, the framework provides a good summary of theoretical ideas in psychology to applied behavior change research and it can be used as a reference if practitioners want to read more about specific theories and computational models. The wide coverage of our framework may help to broaden the theoretical knowledge of applied researchers, encouraging them to experiment with more cutting-edge theoretical propositions rather than restricting themselves to a few classical theories.

Moreover, the framework's emphasis on behavioral processes and their corresponding digital intervention techniques should contribute to the identification, implementation, and evaluation of intervention techniques. First, the framework makes a clearer distinction between people's behavioral processes and the techniques that may influence the processes, when compared with some existing taxonomies. For example, *habit formation* has been considered as a behavior change technique (Abraham & Michie, 2008), but it is essentially a behavioral process that also operates without interventions and is driven by multiple lower-level processes. It is more informative and actionable for system designers, if they are told to look at the specific processes underlying habit formation and how they can be changed, rather than to simply implement a technique called habit formation. Second, by mapping digital intervention techniques to behavioral processes in our framework, it should become clear that often a technically well-defined function can target multiple distinct behavior processes. For example, self-tracking may increase users' knowledge about behavior-health links, but may also support self-monitoring (Kersten-Van Dijk et al., 2017). We argue that evaluation research (e.g., review and meta-analysis) should focus more on the effects of intervention techniques have on individual processes rather than the effectiveness of broadly defined categories of technologies (e.g., "feedback system", Hermsen et al., 2016), in order to gain a better understanding of how and why certain intervention techniques work. Third, when combining multiple intervention techniques to a single digital system, our framework can inform designers about whether the techniques target complimentary processes/constructs or the same process/construct. In the latter case, the combination of techniques as a package may not necessarily be more effective than its components, and more careful analysis is needed. For example, as implementation intention and "just-in-time" reminder both increase the activation values of desirable options, it is questionable whether combining them would yield better results (cf. Pinder et al., 2016).

The adaptive decision-making framework may also have a positive impact of digital applications on basic science. With a clear mapping between digital intervention techniques and the basic processes/constructs in the framework, theorists would be more aware of the digital systems that happen to target the processes/constructs of their interests. Such systems can then be utilized to collect time-intensive and ecologically valid behavioral data to test basic theories or computational models.

2.5.3 Limitation of the framework

Several limitations of the adaptive decision-making framework can be noted. First, the framework is restricted to the explanation and modeling of individual lifestyle behaviors, so social processes and interactions between individuals are not considered. Second, as the framework focuses on the universal behavioral and cognitive processes across people and across behavior domains, it does not uncover the goal values and attribute values for a specific behavior domain and a specific individual. Third, although the framework includes the most important

processes and constructs from a theoretical perspective, it does not address which processes are easier to be changed by digital systems. For example, higher-level goal values (e.g., seizing the moment or caring for the future) may be very influential on people's behaviors, but they are relatively difficult to change given one's social and developmental circumstances (see Hall & Fong, 2007).

2.5.4 Conclusion

Despite the limitations, we developed the framework in the hope that the adaptive decision-making framework will benefit behavior change theorists, digital system designers, and most importantly, facilitate a better communication between the two communities. A stronger synergy will help to bring the future where digital systems become ubiquitous tools to support healthy living. In the meantime, a wider adoption of more effective digital interventions will offer ample opportunities for building and testing new theories of human behavior.

For the remaining chapters in this thesis, the adaptive decision-making framework has also highlighted to particular challenges for changing lifestyle behaviors – the challenge of breaking bad habits and forming good habits through repeated daily decisions, and the challenge of enabling sufficient self-control in everyday environments full of temptations. Chapter 3 and 3 directly develop and evaluate a computational of habit formation based on the framework. Chapter 5 and 6 address the topic of self-control.

Chapter 3

A Sequential Sampling Model of the Integration of Habits and Goals

“Echo, Stop!”

Me shouting at my Amazon Echo Dot (March 28th, 2019).

3.1 Introduction

It was an early morning before I went to High-Tech Campus Eindhoven to give a talk about the content of this very chapter at the annual PhD DeVo event of Philips Research. For me as an evening person, it was quite rare that I was already having breakfast in my living room at around 7:15. Suddenly, music started playing from my bedroom – it was Michael Jackson’s Billie Jean, the third in a series of alarms set in my smart speaker, an Amazon Echo dot, to wake me up every day. “Echo, stop!”, I immediately shouted, but very soon I realized something was quite off. Why would I stop my favorite music from playing, since I was already eating my breakfast and the music could certainly cheer me up? At the same time, I somehow felt that muting it was the most appropriate response.

This personal experience has become my favorite example of how habits and goals can be in conflict. There are many everyday situations where people repeat behaviors that worked for them in past but compromise their current best interests. At a road junction, a driver may quickly turn to the route that they usually take for years, even though they are aware of the ongoing constructions on that road. In a supermarket, a consumer may pick up a familiar product at a familiar location without much thoughts about whether the product is needed on that day. In this chapter, we show through simulation studies that this phenomenon can be explained as an integration of habits and goals in a dynamic decision-making process.

⁶ This chapter is based on Zhang, C., van Wissen, A., Dotsch, R., Lakens, D., & IJsselsteijn, W. A. (submitted). A sequential sampling approach to the integration of habits and goals.

3.1.1 The devaluation paradigm and the arbitration models

The habit-goal conflict has been most convincingly demonstrated in instrumental learning experiments using the so-called devaluation paradigm (e.g., Adams, 1982; Dickinson, 1985; Tricomi, Balleine, & O'Doherty, 2009). In such experiments, human or animal subjects are trained to acquire certain behavioral responses (e.g., pressing a lever or a key) in order to obtain rewards that satisfy their goals (e.g., food or virtual money). A robust finding is that when the subjects are overtrained, their behaviors in extinction tests (i.e., where the behaviors are no longer rewarded) become insensitive to whether the corresponding goals are inhibited or not (e.g., by poisoning the food or devaluing the virtual money). This *devaluation effect* has been established as the empirical basis for the well-accepted theoretical proposition that two distinct learning systems – a habit system and a goal system – jointly control the selection of behaviors (Dolan & Dayan, 2013; Thorndike, 1932; Yin & Knowlton, 2006).

The exact mechanism of how habits and goals interact to control behavior has recently been modeled as an *arbitration process* based on the properties of the two learning systems (e.g., Daw et al., 2005; Keramati et al., 2011; Miller et al., 2019). In their seminal work, Daw et al. (2015) proposed an uncertainty-based arbitration process. Following reinforcement learning theory (Sutton & Barto, 1998), habit and goal-directed learning are mapped to model-free and model-based reinforcement learning respectively. As model-based reinforcement learning is faster but computationally more complex, the goal system is the more reliable than the habit system at the early stage of instrumental training, but it eventually ends up with slightly higher uncertainty at the asymptotes, due an additional computational noise (Daw et al., 2005). Since a model-free algorithm does not respond immediately to environmental changes, when the habit system takes control after extensive training, the behavior becomes insensitive to goal devaluation. Lee, Shimojo, and O'Doherty (2014) provided some initial evidence that human brain may employ such an arbitration mechanism.

Keramati et al. (2011) devised a different arbitrator that looks not only at uncertainty (or what they called value of information), but also the cost of computation. The goal-directed system is assumed to be always more informative, but switching to it to get more information comes with the cost that the additional time used for model-based computation may have been used otherwise for receiving reward. Thus, as the habit system gradually reduces uncertainty through extensive training, it wins the arbitration as the advantage of switching to the goal-directed system is suppressed by its cost. Following a value-free view of habit, Miller et al. (2019) suggested an arbitration based on the variance of action values or habit strengths learned in the two systems. The logic is that if one system is more effective at distinguishing different response options, it should be preferred or weighted more. Critically, extensive training increases the variance in the habit system, but decreases the variance in the goal system. Both models can reproduce the classical devaluation effect (Adams, 1982; Dickinson, 1985), but also findings in other paradigms, including reversal learning (Pessiglione et al.,

3.2 The conceptual and the computational model

2005), learning with concurrent schedule (Kosaki & Dickinson, 2010), and learning with different reinforcer scheduling (e.g., Dickinson, Nicholas, & Adams, 1983).

All arbitration models can be summarized in a single conceptual scheme (Figure 3.1a). In two distinct systems, action values of different behavioral responses are learned, representing how these responses compare to each other on some relevant dimensions. Because action values learned in the two systems may disagree with each other, an arbitration process is needed to decide which system controls behavior, based on the relative strength of the two systems. After arbitration, the behavioral response with the highest action value learned in the dominated system is selected. Because the habit system lags behind the goal system in reaching its maximum performance but is ultimately more efficient, so the control of behavior switches from the goal-directed system to the habit system in the later stage of learning (Figure 3.1b).

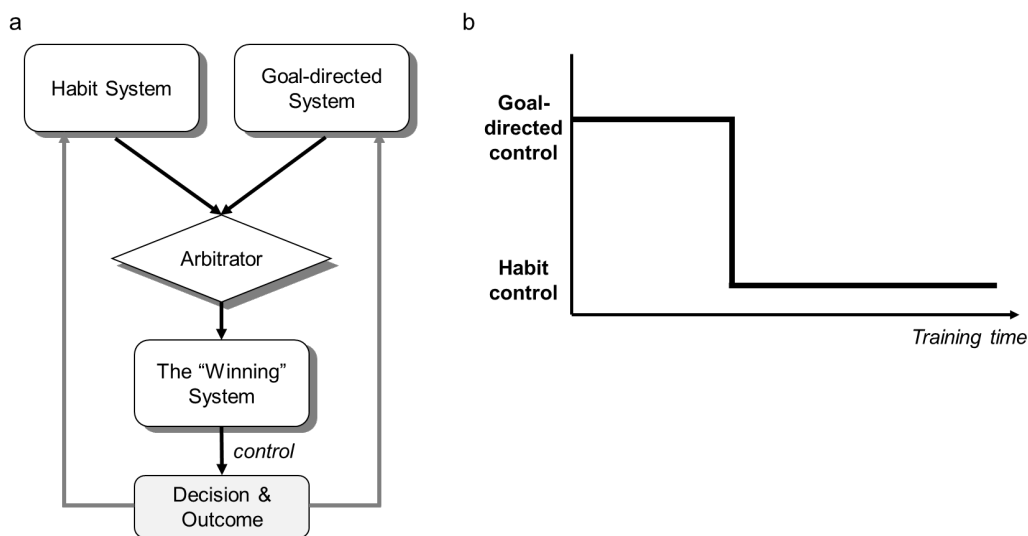


Figure 3.1 (a) A common scheme for arbitration models; (b) Predicted by arbitration models, control of response selection switches from the goal-directed system to the habit system after a certain amount of training.

3.1.2 Motivation for a sequential sampling approach

Although arbitration models can qualitatively reproduce some important empirical findings in the instrumental learning literature (Daw et al., 2005; Keramati et al., 2011; Miller et al., 2019), their general approach can be questioned on several grounds. First, while the two separate learning systems and their neurological substrates are well-established (Yin & Knowlton, 2006), the existence of an additional arbitrator remains a critical assumption, awaiting more neurophysiological evidence (but see Lee, Shimojo, & O’Doherty, 2014). Moreover, after arbitration, the response selection process is the same for either of the learning systems, regardless of which one is in control. This contradicts with the seemingly qualitative differences in how habits and goals influence behaviors – habitual responses are often conceptualized as

impulses triggered by environmental cues (see Wood & Neal, 2007), which are sometimes overruled by goals. Finally, the arbitration models are not well-equipped to account for an important aspect of learning – the temporal change of decision time. The Daw et al. (2005) and Miller et al. (2019)'s models predict identical decision times for responses controlled by habits and goals, while Keramati et al. (2011)'s model produces unrealistic sudden switches between very fast (habitual) and very slow (goal-directed) responses.

The above issues are all related to the fact that response selection is oversimplified in those models grounded in learning theories. Taking decision-making as an alternative theoretical perspective may help to solve these issues. Specifically, we might consider response selection in instrumental learning as value-based decision-making, and it is plausible that computational models of decision-making also apply to habit-goal conflicts. Sequential sampling models have been used to model choice patterns and decision times in many problems in psychological science (for a review, see Oppenheimer & Kelso, 2015), including memory retrieval (e.g., Ratcliff, 1978), perceptual decision-making (e.g., Ratcliff & Rouder, 1998), and value-based decision-making (e.g., Roe et al., 2001). The sequential sampling approach generally assumes that decision-makers accumulate preferences for different choice alternatives in multiple steps by sampling goal-related attributes of the alternatives, before the preference for one alternative eventually exceeds a decision threshold. If habits and goals can be mapped to distinct parameters in a sequential sampling model, they may be integrated dynamically to produce habit-goal conflicts and other habit-related phenomena, without resorting to an additional arbitrator.

Two important and distinct determinants of any sequential sampling process are the *starting position* of preference accumulation (baseline preference) and the *drift rate* at each step of preference accumulation (Forstmann, Ratcliff, & Wagenmakers, 2016). In their seminal paper on multialternative decision field theory, Roe et al. (2001) discussed a possible mapping of habits and goals to starting position and drift rate respectively, but the idea was not examined any further in value-based decision-making research. A stronger rationale for this mapping comes from research results in the neighboring field of perceptual decision-making, where a typical task requires choosing the correct movement direction of groups of dots. It was found that while drift rate related to stimulus ambiguity in the current trial, starting position related instead to past choices (Bode et al., 2012; Mulder, Wagenmakers, Ratcliff, Boekel, & Forstmann, 2012). If a similar distinction between past and current information applies to value-based decision-making, then habits and goal-related values may play the same roles as past choices and current perceptual evidence respectively. Furthermore, Akaiishi and colleagues (Akaiishi, Umeda, Nagase, & Sakai, 2014) found that the way past choices influence current choice in the perceptual domain is mathematically equivalent to a form of Hebbian learning (Hebb, 1949), which has been theorized previously to also underlie habit learning (Klein, Mogles, Treur, & van Wissen, 2011; Miller et al., 2019).

3.2 The conceptual and the computational model

In the next sections, we formally introduce a sequential sampling model in which habits and goals are integrated dynamically, and show through simulations that the model can reproduce choice patterns and predict gradual changes of decision times in three instrumental learning tasks – classic devaluation paradigm, devaluation paradigm with a concurrent schedule, and reversal learning. The simulations studies provide a good test of the validity of our model, as the reproduced findings are considered as classic demonstrations of habit-goal conflicts and were previously addressed by the arbitration models.

3.2 The conceptual and the computational model

In order to model the cognitive processes in instrumental learning, we first define the structure of a typical instrumental learning task and the relevant variables involved in the task. We focus on an example of animal experiments where rodents learn to press a lever to obtain food (Figure 3.2a), but the same task definition also applies to human instrumental learning tasks. In a constrained environment (e.g., a feeding cage), a learning agent is assumed to have a fixed number of goals. For example, a rodent may strive to obtain food, water, mating opportunities, and also to rest. At one moment, these goals can differ in their importance to the agent, numerically represented by *goal values*. To satisfy its goals, the agent needs to engage in certain behaviors, and it can be assumed that given the constrained environment, only a limited number of behavioral responses are available. In a typical experiment, it can be simplified that rodents only have two possible responses – to press a lever or to rest. For each goal and each behavioral response, *attribute values* can be specified to represent the likelihood of achieving the goal by executing the behavior (e.g., pressing lever scores high on attribute *food*, resting scores high on attribute *leisure*). Note that among all the goal-related attributes, some can be called unattainable attributes as no behavioral response in the constrained environment can satisfy those goals (e.g., *mating* is an unattainable attribute given the absence of other rodents in the cage).

As in classic expected utility theory (e.g., Savage, 1954; von Neumann & Morgenstern, 1947), the overall value of a response depends on both its attribute values and the corresponding goal values. In addition, each behavioral response also has a different *habit value*, depending on its frequency of being selected in the past. Overall, the task of the learning agent is to search for the behavioral response that can maximize the satisfaction of its various goals through repeated decisions. This representation is similar to the multi-armed bandit task in the reinforcement learning literature, where an agent learns the pay-offs of multiple slot machines through repeated trials (Sutton & Barto, 1998; for a similar representation of instrumental learning, see Fontanesi, Gluth, Spektor, & Rieskamp, 2019).

Conceptually, the learning task consists of a sequence of interconnected decision-making (response-selection) and learning processes (Figure 3.2b). At each iteration, the current goal

values, attribute values, and habit values are integrated in a sequential sampling process to produce the current decision (e.g., press the lever) and the associated outcome (e.g., food delivered). Two distinct learning processes then follow. In goal-directed learning, perceived outcomes are used to update the agent's beliefs about the attribute values of the behavioral responses (*law of effect*, Thorndike, 1932). With enough repetition, perceived attribute values should converge to the true likelihoods of satisfying different goals by executing different behaviors. In a separate habit learning process, the habit values of behavioral responses are updated based simply on whether the responses are selected at this iteration (*law of exercise*, Thorndike, 1932). Cognitively, habit values reflect the mental associations between behaviors and environmental cues strengthened through their repeated co-occurrences⁷ (Klein, Mogles, Treur, & van Wissen, 2011; Wood & Neal, 2007; Wood & R nger, 2016). The updated attribute values and habit values are used for the subsequent decisions.

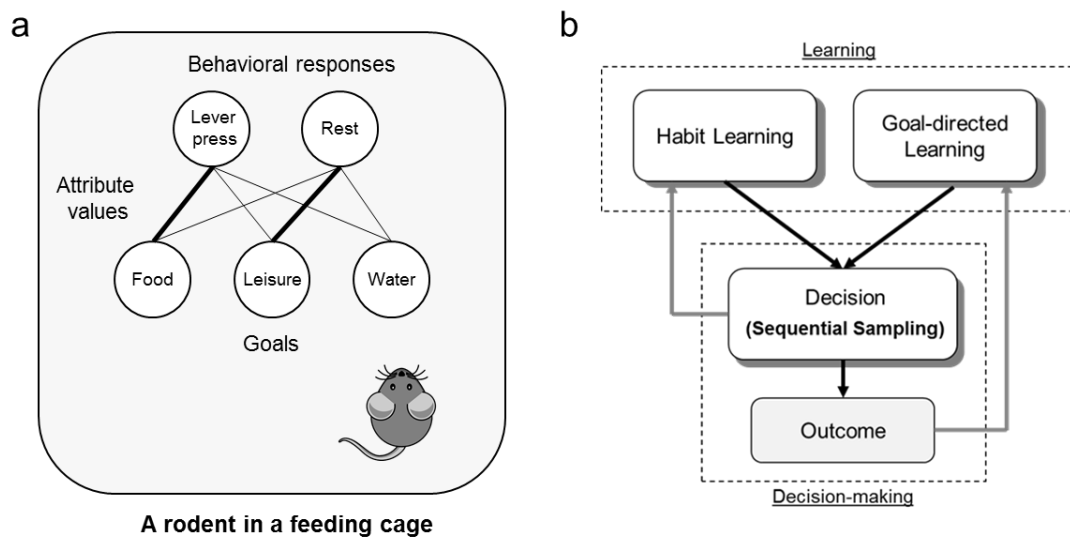


Figure 3.2 (a) A representation of behavioral responses, goals, and goal-related attribute values in a typical instrumental learning experiment with rodents; (b) A representation of the task as a repeated alternations between decision-making and learning.

⁷ There is an ongoing debate on whether habit learning depends on the decisions alone (value-free, see e.g., Miller et al., 2019; Miller, Ludvig, Pezzulo, & Shenhav, 2018; Pauli, Cockburn, Pool, P rez, & O'Doherty, 2018) or also on decision outcomes (e.g., as model-free reinforcement learning, see Daw et al., 2005; Keramati et al., 2011). Because our main objective is to propose a new model of habit-goal integration in response selection, we take the value-free view of habit learning for its simplicity and its similarity to the updating rule of prior choice's effect in perceptual decision-making (Akaishi et al., 2014). In theory, our sequential sampling approach should remain effective even if the alternative view of habit learning is taken.

3.2.1 Modeling response selection as a sequential sampling process

For modeling response selection in the instrumental learning task, we adopted the general framework of the multialternative decision field theory (MDFT; Roe et al., 2001), but other sequential sampling models of value-based decision-making should also work in principle (e.g., Trueblood, Brown, & Heathcote, 2014; Usher & McClelland, 2001). Figure 3.3 provides a visual summary of the model, showing how the outcome and time course of a response selection (between lever press and rest) are determined in a sequential sampling process as influenced by four variables – *starting positions*, *sampling probabilities*, *drift rates* at each time step, and a *decision threshold*⁸.

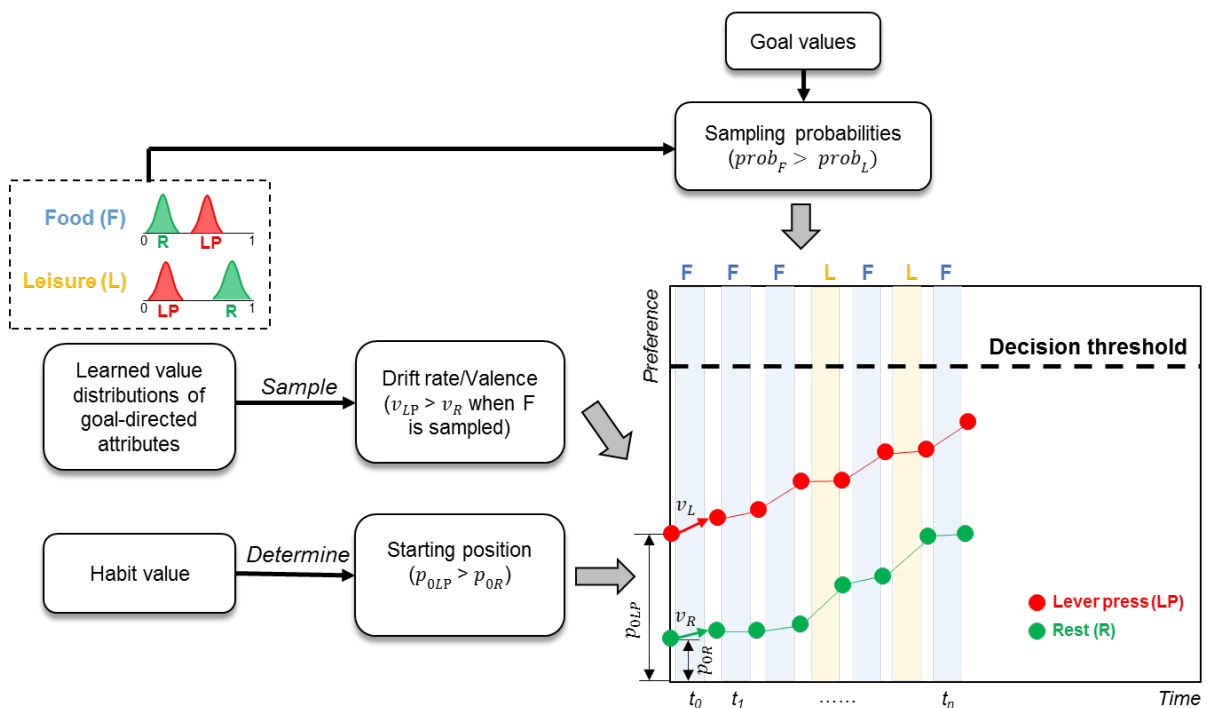


Figure 3.3 A detailed representation of a sequential sampling process and its inputs.

At the start of a sequential sampling process, starting positions represent a decision-maker’s baseline preference towards a set of behavioral responses. The model proposes that habitual responses (i.e., responses that have been chosen more frequently in the past) are by default more favorable than the less habitual ones, represented by higher starting positions⁹ (Roe et al., 2001). The starting positions or the preferences at t_0 for all responses equal to their habit

⁸ For the applications in this chapter, we assume that decision-making processes are terminated by an internal representation of decision threshold. However, decision-making processes can also be forced to terminate at an arbitrary time t , and the response with the highest preference at that time is chosen.

⁹ A different cognitive mechanism with the same consequence is that the decision-maker start to accumulate preferences for habitual responses earlier than other responses (for a similar idea, see Psarra, 2016).

strengths (\mathbf{H}) scaled by a scalar parameter θ :

$$\mathbf{P}_{(0)} = \theta \mathbf{H} \quad (1)$$

From their starting positions, a learning agent’s preferences for different responses drift over time, and at each time step, the drifts depend on which goal-related attribute is sampled and how each response scores on the sampled attribute. For example, if attribute *food* is sampled, the preference for the response lever-press will increase greatly because lever-press scores high on attribute *food*. A key assumption of MDFT is that at one time step, the decision-maker only sample one attribute (e.g., at t_1 *food* is sampled while at t_2 *leisure* is sampled). In the original MDFT, the sampling probabilities are equal for all attributes (i.e., sampling randomly). Instead, our model proposes that sampling probabilities of attributes are determined by two variables – the *goal values* of the attributes and the *attainability of attributes*, which measure the importance and relevance of the attributes respectively in the current task. If, for example, obtaining food is more important than conserving energy for rodents, *food* will be sampled more than *leisure*. Also, if one attribute is more attainable in the current task (contained more in the responses) than another attribute (e.g., some behavioral responses result in *food*, but none results in *mating*), it will be more likely to be sampled. Mathematically, a softmax function is used to calculate sampling probability (Pr_j), with the multiplications of goal value (G_j) and the attainability of attributes (A_j) as inputs and τ as a scaling parameter,

$$Pr_j = \frac{e^{\tau G_j A_j}}{\sum_{k=1}^K e^{\tau G_k A_k}} \quad (2)$$

where the attainability of each attribute is the sum of all responses’ scores on that attribute (X_{ij}), $A_j = \sum_{i=1}^N X_{ij}$. Attribute values are often given externally in choice experiments, but in our learning task they are derived from learned probability distributions for each attribute. For the calculation of A_j , the model assumes that the expected mean reward values (EMRs) of the distributions are used. Later, attribute values sampled at each time step, the lowercase M_{ij} , are instead randomly sampled from the distributions.

Two implications of Equation 2 are worth noting. First, the unattainable attributes will have very low though non-zero sampling probabilities. As there can be many unattainable attributes in a controlled instrumental learning task, the sum probability of sampling any unattainable attribute can be non-trivial (e.g., 5%), and it is similar to the probability of sampling noise, which is usually arbitrarily defined in other sequential sampling models (e.g., Roe et al., 2001). Second, goal values for different attributes are assumed to be stable (fixed) in short time frames for each decision-maker, but can be substantially changed through experimental procedures such as goal devaluation (e.g., Adams, 1982; Dickinson, 1985). Consequently, if a food is devalued, its sampling probability also decreases towards zero.

3.2 The conceptual and the computational model

The rest of the model follows MDFT closely. When an attribute is selected based on sampling probabilities at time t , the momentary drift rates of behavioral responses (or *valences* as in Roe et al., 2011) are their attribute values on the sampled attribute, as in the matrix form¹⁰:

$$\mathbf{V}_{(t)} = \mathbf{M}_{(t)} \mathbf{W}_{(t)} \quad (3)$$

where $\mathbf{V}_{(t)}$ is an N -dimensional valence vector representing the drift rates of different behavioral responses at different time steps. $\mathbf{W}_{(t)}$ is a J -dimensional vector of attribute weights, in which the sampled attribute is weighted 1 and all others are weighted 0. Lastly, $\mathbf{M}_{(t)}$ is an N -by- J matrix containing all attribute values for all responses. Unlike the original MDFT, where $\mathbf{M}_{(t)}$ is fixed for all t , $\mathbf{M}_{(t)}$ elements are randomly sampled according to the underlying probability distribution learned for each response-attribute pair at each time step.

Next, preferences $\mathbf{P}_{(t)}$ at time t are determined by the preferences at the previous time step ($\mathbf{P}_{(t-1)}$) and the current drift rates $\mathbf{V}_{(t)}$. Between two successive time steps, there is a decay or leakage of each preference itself, and there are influences from the preferences of competing responses, often in the form of lateral inhibition. Both processes are summarized in an N -by- N matrix \mathbf{S} , in which elements on the main diagonal are equal to a self-decay parameter (S_{self}) and all other elements are equal to a lateral inhibition parameter ($S_{lateral}$). Thus, preferences are calculated in the matrix form:

$$\mathbf{P}_{(t)} = \mathbf{S} \mathbf{P}_{(t-1)} + \mathbf{V}_{(t)} \quad (4)$$

When a behavioral response's preference exceeds the decision threshold, a decision is made and the behavior is executed by the learning agent. Reward to be received relating to each attribute or goal is calculated by reward probabilities pre-defined by the learning task (e.g., the reinforcement schedule of a learning experiment). Before making the next decision, habit values and goal-related attribute value distributions are updated.

3.2.2 Modeling habit and goal-directed learning

We assume that habits are value-free, meaning that their updates depend only on the decisions (and the associated behavioral executions) themselves but not on the consequences brought by the decisions. Specifically, the model for habit learning uses the same Hebbian learning equation as in Miller et al. (2019), but is also conceptually compatible with other

¹⁰ In the original MDFT, valence is computed as $\mathbf{V}_{(t)} = \mathbf{C} \mathbf{M}_{(t)} \mathbf{W}_{(t)}$, where \mathbf{C} is an N -by- N contrast matrix with all the elements on the main diagonal equal to 1 and all other elements equal to $-1/(N-1)$. In this way, valences measures relative advantages or disadvantages of responses, rather than absolute attribute values. We tested this version as well, but it would severely attenuate the impact of habit strength. This strong contrasting mechanism is not essential to the sequential sampling approach, and it is not used in other models (e.g., Bhatia, 2013; Trueblood et al., 2014; Usher & McClelland, 2004).

equations (Klein, Mogles, Treur, & van Wissen, 2011; Psarra, 2016; Tobias, 2009):

$$\mathbf{H}_{(T)} = \mathbf{H}_{(T-1)} + \alpha_H(\mathbf{A}_H - \mathbf{H}_{(T-1)}) \quad (5)$$

where learning rate α_H controls how much habit values (\mathbf{H}) change from one time point to the next¹¹, and \mathbf{A}_H is a scaling parameter which limits the upper-bound of habit values. The equation implies that with repeated behaviors, habit values increase fast at the beginning and then their growth slow down until the values reach their asymptotes. This pattern is consistent with empirical data on the dynamics of self-reported habit strength (Lally, van Jaarsveld, Potts, & Wardle, 2010).

Compared with habit learning, the algorithm for goal-directed learning can be very complex. Previous models have implemented model-based reinforcement learning algorithms (Daw et al., 2005; Keramati et al., 2011; Miller et al., 2019). Since we simplified our task representation to single-state repeated decision-making or multi-armed bandit problem (in contrast to Markov decision process), goal-directed learning may be modeled with a simple algorithm of Bayesian belief update – combining prior distributions (beliefs about attribute values before a decision) and data (i.e., rewards) to obtain posterior distributions (beliefs after a decision). Assuming that the reward generation processes in learning experiments are Bernoulli processes, beta distributions can be used for priors and posteriors. Formally, the updating rule is expressed as:

$$(\alpha_{ij}, \beta_{ij}) \leftarrow \begin{cases} ((1 - \gamma)\alpha_{ij} + \gamma\bar{\alpha}, (1 - \gamma)\beta_{ij} + \gamma\bar{\beta}), & D_{(T)} \neq i \\ ((1 - \gamma)\alpha_{ij} + \gamma\bar{\alpha} + R_{j(T)}, (1 - \gamma)\beta_{ij} + \gamma\bar{\beta} + 1 - R_{j(T)}), & D_{(T)} = i \end{cases} \quad (6)$$

where the alpha and beta parameters defining the beta distribution of response i on attribute j are only updated by reward $R_{j(T)}$, if decision at T ($D_{(T)}$) is to choose response i . To account for the nonstationary environments in typical experimental setups (e.g., reward functions can be suddenly changed by the experimenter), parameter γ is used to inject uncertainty in the distributions. In other words, belief distributions always regress to a default distribution defined by $\bar{\alpha}$ and $\bar{\beta}$ (a uniform beta distributions with both equaling 1), ensuring fast reactions of learning agents to changes in the environment. Note that sampling values from the beta distributions defined in equation 6 resembles the widely-used Thompson sampling approach to solve practical Bandit problems (see Russo, van Roy, Kazerouni, Osband, & Wen, 2015), although in our model the distributions are sampled many times (i.e., sequential sampling) rather than only once within each single decision.

¹¹ The uppercase T in the equation denotes time point or decision point (e.g., trial number in experiments), which is different from the time step t in the sequential sampling of each decision.

3.3 Simulation studies

We conducted three simulation studies to validate our model. First, the model was used to reproduce the classic devaluation effect (e.g., Adams, 1982; Dickinson, 1985), and sensitivity analyses were performed to see if the effect was robust against changes of parameter values. Next, we extended the model to a devaluation paradigm with two competing response options (Kosaki & Dickinson, 2010). Finally, following previous works (Keramati et al., 2011; Miller et al., 2019), the model was used to reproduce the findings in a reversal learning task (Pessiglione et al., 2005), focusing on how it can produce gradual changes in decision time. For all model parameters (i.e., excluding task-specific parameters), the same values were used for all three studies, as shown in Table 3.1.

Table 3.1 Parameter values used in all three studies.

	Parameter	Explanation	Value
Decision-making (Sequential sampling)	θ	Scaling parameter for transforming habit strengths to starting positions. The exact value is arbitrary, but it should scale the largest habit strength possible (close to 1) to the decision threshold (e.g., 1).	1
	τ	Scaling parameter for the softmax function used in Equation 2. The larger the value, the more dominant the largest input is in calculating the outputs. The value is arbitrary, but depends on the scale used for goal values, e.g., [0, 1]).	10
	S_{self}	Memory parameter that measures on the information loss ($1 - S_{self}$) in preference accumulation (e.g., 0.94 used in Roe, et al., 2001).	0.99
	$S_{lateral}$	Lateral inhibition parameter that measures the competition among choice options (e.g., -0.001 and -0.025 used in Roe et al., 2001).	-0.03
	DT	Decision threshold for sequential sampling. The exact value is arbitrary, as it depends on the scales used for attribute values (e.g., [0, 1] in our studies).	1
	$maxStep$	The maximum time step allowed in a sequential sampling process if no option's preference exceeds decision threshold.	100
	$N_{unattain}$	Number of unattainable attributes.	10

(To be continued)

	Parameter	Explanation	Value
<i>Habit learning</i>	α_H	Learning rate in the Hebbian equation for habit learning. The larger its value, the faster habit strengths update. Miller et al. (2019) used much smaller values (e.g., 0.001), and indeed many more training trials were required to reach full habit strengths (e.g., 6000).	α_H
	A_H	Scaling parameter determining the upper bound of habit strength (usually 1, Miller et al., 2019).	1
<i>Goal-directed learning</i>	γ	Uncertainty parameter that determines the rate of uncertainty injected in the Bayesian belief updates. The larger its value, the faster a learner discounts “old” information, or “forgets” faster (e.g., 0.01 used in Russo et al., 2015).	0.1
	$\bar{\alpha}$	Alpha parameter of the convergence distribution in the absence of observations (uniform beta distribution was used, see Russo et al., 2015).	1
	$\bar{\beta}$	Beta parameter of the convergence distribution in the absence of observations.	1

3.3.1 Study 1: Classic devaluation effect

The classic devaluation effect shows that learning agents become insensitive to goal devaluation after extensive training, but remain sensitive after moderate training. The effect has been repeatedly replicated for both animals and humans (e.g., Adams, 1982; Dickinson, 1985; Killcross & Coutureau, 2003; Liljeholm, Dunne, & O'Doherty, 2015; Tricomi et al., 2009; Yin et al., 2004; Yin, Knowlton, & Balleine, 2005), and is considered a seminal finding for differentiating habits from goal-directed behaviors. In a typical animal devaluation experiment, rodents learn to press a lever to obtain food pellets through either moderate or extensive pairing of the response and the food. After training, half of the rodents are subjected to a devaluation procedure, where the food becomes undesirable because of either a satiation procedure or food-aversive conditioning (indicated as the “devalued” or “paired” group). The other half undertakes a similar procedure but with a different food not used in training (indicated as the “non-devalued” or “control” group). Finally, in the extinction test, no food pellets are delivered no matter how frequently the rodents press the lever. The devaluation effect manifests as an interaction effect. After moderate training, rodents in the devalued group press the lever less often than their peers in the control group. For rodents that receive extensive training, their lever-pressing responses seem to become insensitive to goal devaluation – both the devalued and the control group press the lever with equal frequency.

3.3 Simulation studies

In the simulated experiment, learning agents were trained to press the lever for either 40 or 240 trials (as in Keramati et al., 2011), in which they were assumed to have a higher goal value for obtaining food ($G_{food} = 0.8$) than for having some rest ($G_{leisure} = 0.4$). Pressing the lever would lead to food 60% of the time¹², but never any leisure. Relaxing (no lever-pressing), on the other hand, always led to leisure but no food. Besides food and leisure, the agents were assumed to have 10 other important goals ($G_{unattain} = 0.8$), but these goals were unattainable by either of the two responses. Devaluation was implemented as the diminishing of G_{food} to 0 for half of the agents. In the 100 extinction trials, the probability of obtaining food by lever-pressing was reduced from 0.6 to 0. Five-hundred simulations of homogenous agents were run.

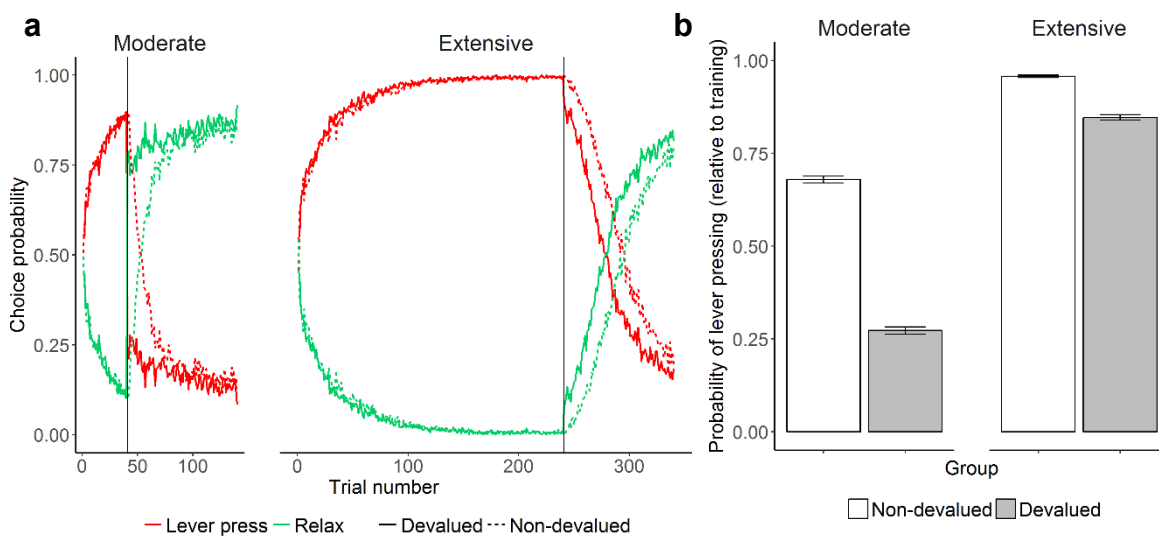


Figure 3.4 Simulated behavioral results for a classic devaluation experiment. (a) Change of choice probability over time; (b) Aggregated lever-pressing rates in the first 20 trials after devaluation relative to the level at the end of training.

Figure 3.4 shows simulated choice probabilities over time and aggregated response rates. Our model produced a main effect of training (higher lever-pressing rates after extensive training), a main effect of devaluation (lower lever-pressing rates when G_{food} is devalued), and most importantly a clear *training duration* by *devaluation* interaction effect. As can be seen in Figure 3.4a, the lever-pressing rates in the two groups decreased almost in parallel after extensive training, while after moderate training the lever-pressing rate of the devalued group declined sharply as compared to the non-devalued group.

Our model also predicted that decision times decreased gradually over the course of training, but increased abruptly after devaluation, before eventually decreasing again (see Figure 3.5a).

¹² The exact reward probability for food was not decisive for reproducing the devaluation effect, as long as it was high enough so that the full acquisition of the lever-pressing response was achieved.

Note that an increase of decision times after devaluation was observed in all conditions, regardless of whether strong habits were formed or not (cf. Keramati et al., 2011). This is a novel prediction of the model that can be tested in future research.

The effect-generating mechanisms of the model are illuminated in the temporal changes of the underlying cognitive variables in the model, especially their values at the transition from training to extinction (point of devaluation for the devalued group) (see Figure 3.5b - 3.5f). First, as expected, the habit values for the two groups after extensive training were very close to 1, while the habit values for the two groups after moderate training were just below 0.75 (Figure 3.5b). Second, there was a sudden change in sampling probabilities for the devalued group – these agents stopped to sample attribute *food* because of the goal devaluation, but instead started to sample the unattainable attributes most of the time (Figure 3.5c, left). In contrast, agents in the control group continued to sample *food* frequently before they gradually unlearned the association between lever-pressing and food in the extinction phase (Figure 3.5c, right). Thus, when looking at the expected mean reward values (EMR) for attribute *food* and the unattainable attributes (Figure 3.5d & 3.5f), it was clear that the response lever-pressing was at disadvantage in the devalued group compared to the control group. The lever-pressing rate of the devalued group dropped significantly faster (Figure 3.4a, left), unless the high habit values for the agents after extensive training functioned as a counteracting mechanism.

3.3 Simulation studies

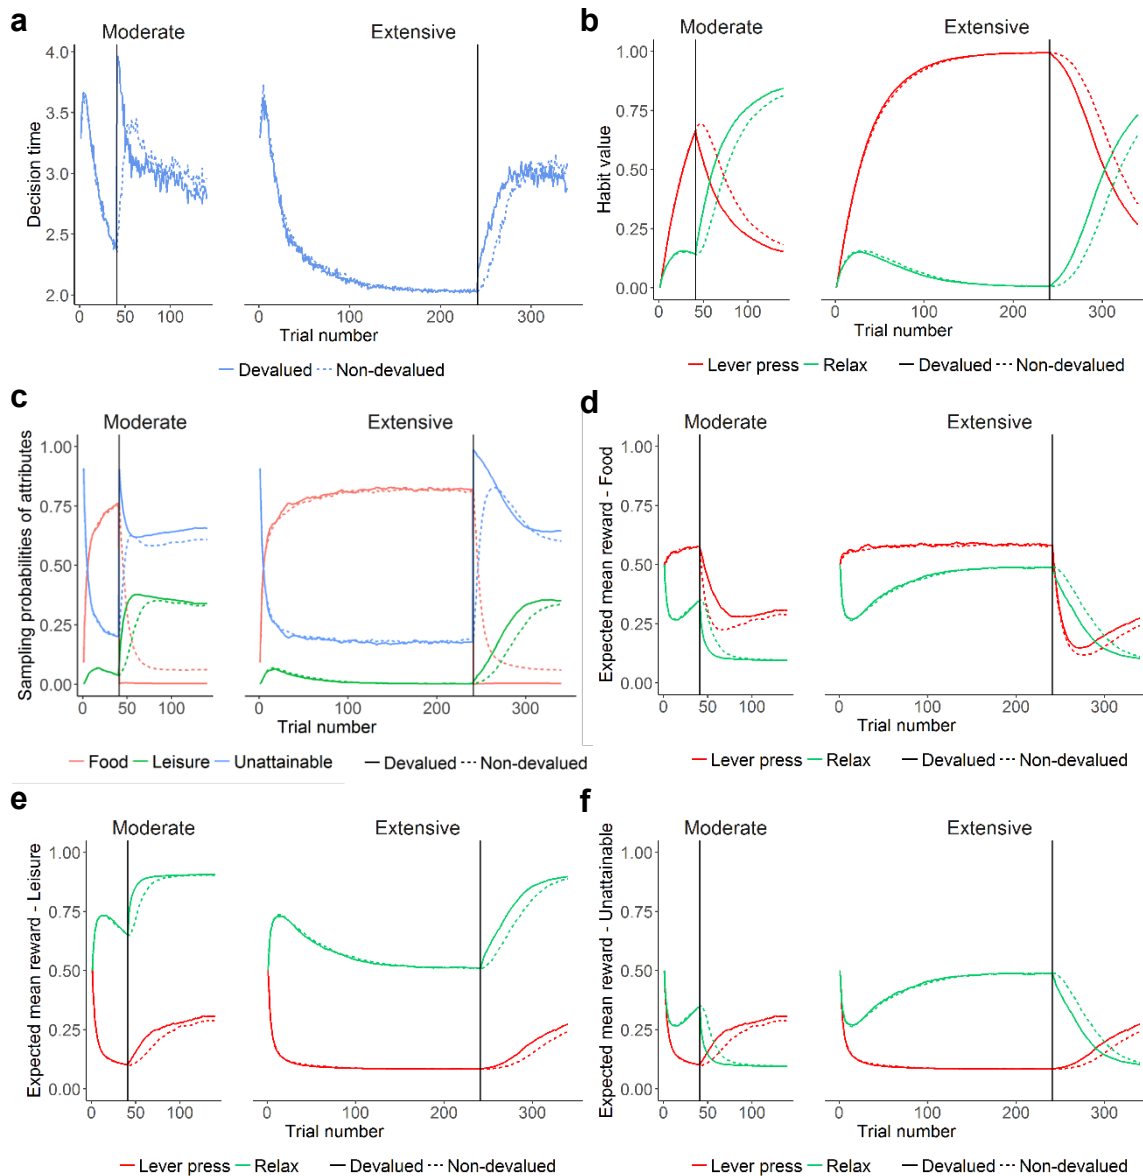


Figure 3.5 Temporal changes of decision time and underlying cognitive variables in the simulated devaluation experiment. (a) Decision time; (b) Habit value; (c) Sampling probability of attributes; (d) EMR of attribute *food*'s distributions; (e) EMR of attribute *leisure*'s distributions; (f) EMR of unattainable attributes' distributions.

Sensitivity analyses showed that the model was reasonably robust in reproducing the devaluation effect against changes in parameter values (Figure 3.6). First, as expected, lever-pressing rates in the devalued and control group only became comparable when the training was more than approximately 170 trials (Figure 3.6a). This result basically reaffirmed the devaluation effect that insensitivity to goal devaluation only happens when the response is over-trained. Second, a very high value for the memory parameter (S_{self}) was needed to reproduce the devaluation effect (Figure 3.6b), consistent with the small memory leakages implemented in sequential sampling models in the literature (e.g., Roe et al., 2001). Third, variation of the

lateral inhibition parameter ($S_{lateral}$) in the range of -0.3 and 0 did not change simulation results to any extent (Figure 3.6c), and the relative low values used were consistent with the literature (e.g., -0.001 and -0.025 used in Roe et al., 2001). Since theoretically lateral inhibition has an effect of reinforcing the responses with default high preferences (due to strong habits), a very large $S_{lateral}$ would result in an unrealistic pattern of no decay of lever-pressing rate in the extinction phase.

Fourth, the curves for habit learning rate confirmed that some habit formation was needed to reproduce the devaluation effect, but if habits were made to form too fast (e.g., $\alpha_H > 0.15$), responses would become insensitive to goal devaluation even after moderate training (Figure 3.6d). Fifth, results of the gamma parameter suggested that a small uncertainty injection was needed to reproduce the devaluation effect (Figure 3.6e), as the parameter positively related to the value distributions of the unattainable attributes that were mostly sampled for the devaluated groups. If there was little uncertainty (e.g., $\gamma < 0.03$), the resultant low value distributions would lead to drift rates that were too small to push the baseline preference of lever-pressing to the decision threshold even after extensive training. In contrast, if a lot of uncertainty was injected (e.g., $\gamma < 0.2$), very large drift rates would be sampled from the value distributions of unattainable attributes and they would push baseline preferences of lever-pressing after both moderate and extensive training to the decision threshold. Finally, the number of unattainable attributes did not seem to have any substantial impact on the generation of the devaluation effect (Figure 3.6f).

3.3 Simulation studies

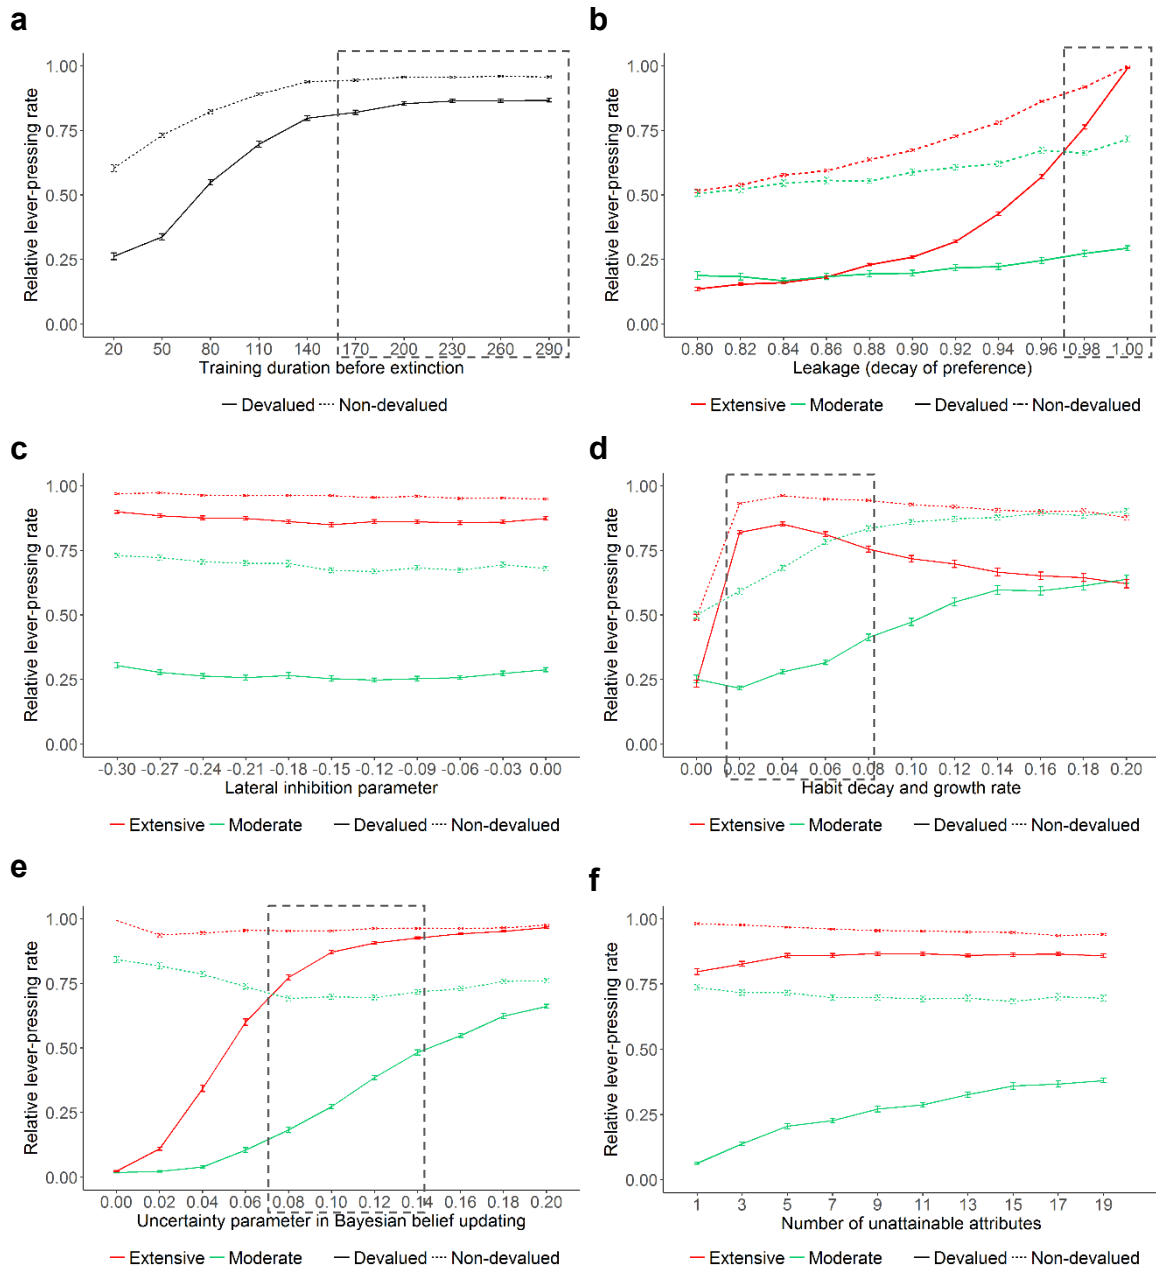


Figure 3.6 Sensitivity of the devaluation effect to different parameter values. (a) Training duration; (b) Leakage parameter (S_{self}); (c) Lateral inhibition parameter ($S_{lateral}$); (d) Habit learning rate (α_H); (e) Uncertainty parameter in Bayesian belief updating (γ); (f) Number of unattainable attributes ($N_{unattin}$). The dashed squares indicate the effect-producing ranges.

3.3.2 Study 2: Devaluation paradigm with a concurrent schedule

We extended our simulation to devaluation experiments with a concurrent schedule. In Kosaki and Dickinson (2010), instead of training one response-outcome pair, rodents were trained to learn two instrumental responses with two types of food concurrently. With this schedule, even if extensive training was used, rodents remained sensitive as to which food

was devalued. Thus, we simulated 500 homogeneous agents only in extensive training to see if the model would produce a clear difference between responses to the devalued and non-devalued food. Other setups were similar to the previous scenario, except that two food attributes (with goal values $G_{food_A} = G_{food_B} = 0.8$) and two lever-pressing responses were used. Each food was again reinforced to the correct response 60% of the time.

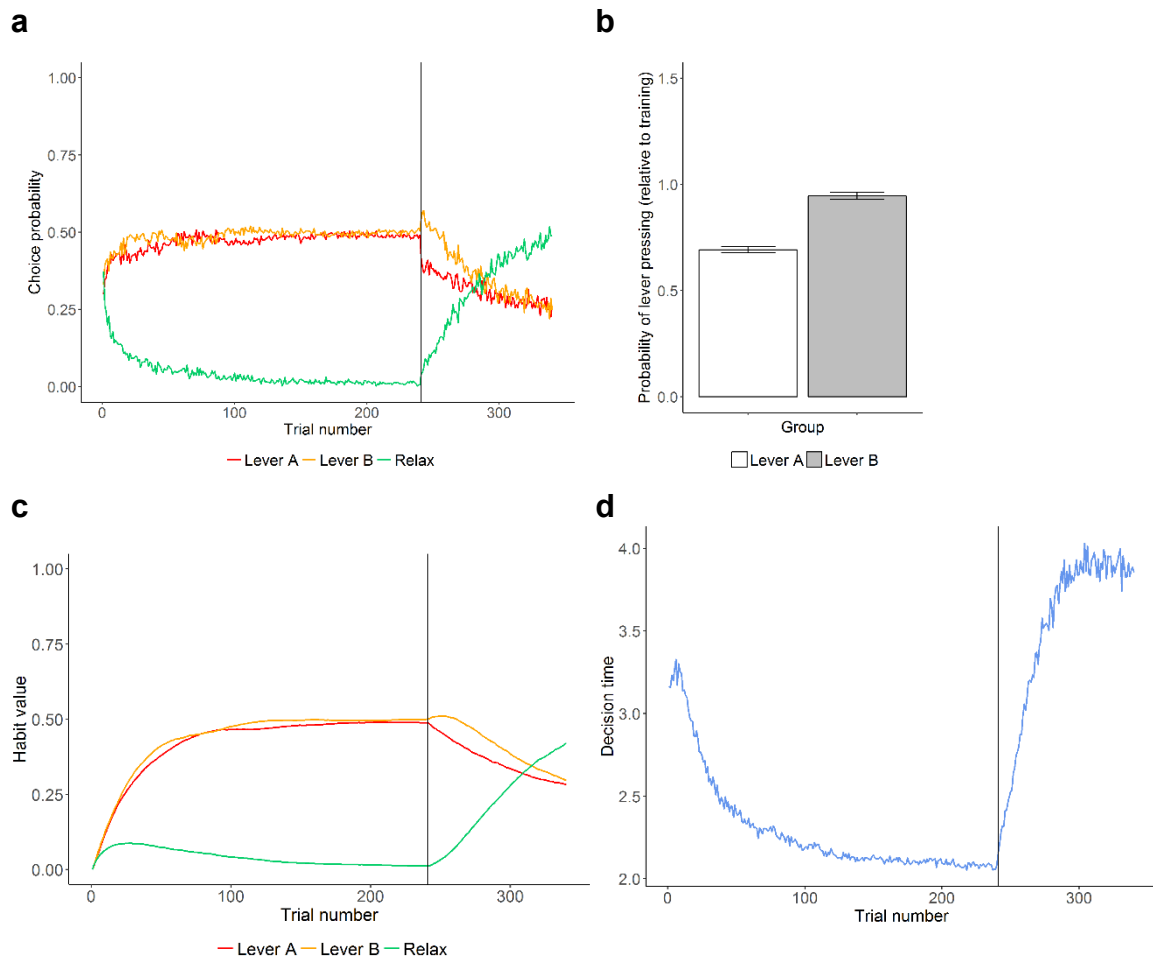


Figure 3.7 Simulated results for devaluation paradigm with concurrent schedule. (a) Change of choice probability over time; (b) Aggregated response rates (relative to the end of training) after devaluation (first 20 trials used); (c) Habit value; (d) Decision time.

As in Figure 3.7a and 3.7b, results were consistent with the empirical finding: at the point of devaluation, choice probability decreased sharply for the devalued response (lever-press A), while it increased for the non-devalued one (lever-press B). Unlike the classic devaluation experiments, even after extensive training, habit strengths for both responses were only moderate (around 0.5, see Figure 3.7c) because of the competition, so the shift in starting positions could not compensate for the disadvantages of the devalued response in terms of sampled attribute values. The model also predicted decision time to decrease gradually during

training, and to increase greatly in the extinction phase, eventually becoming slower than the decision time at the start of training.

3.3.3 Study 3: Reversal learning

Reversal learning refers to learning tasks where payoffs of behavioral responses are occasionally reversed during the task. For example, in Pessiglione et al. (2005), following two stimuli with equal appearance probability, human participants learned in three phases to either press a button (go response) or withdraw from pressing a button (no-go (NG) response) in order to earn as many points as they could. In the training phase, the go-response earns points for one stimulus, while the NG-response earns points for the other. In the reversal phase, the reward-generating stimulus-response mapping was reversed. In the final extinction phase¹³, the NG-response earns points for both stimuli. The basic finding was that people needed time to gradually learn the changes in the underlying reward probabilities and decision time fluctuated in time: responses became faster when a reward-structure was learned and slower when the structure was reversed.

We used the same task structure as in Pessiglione et al. (2005). Learning agents were assumed to primarily focus on accumulating points ($G_{point} = 0.8$) and to a lesser degree to conserve energy (or to obtain leisure, $G_{leisure} = 0.1$). Probabilities of obtaining points were either 0 or 1 for the responses depending on the phases (training, reversal, or extinction), while probabilities of obtaining leisure were all set to 1, since the button-pressing responses do not consume much energy for humans (so attribute *leisure* should have negligible influence on decision). The numbers of trials in the three phases were set to 150, 200, and 150 (as in Keramati et al., 2011). Five-hundred simulations with homogenous agents were run to obtain the results.

Result of choice probability in Figure 3.8a confirmed that the simulated agents could learn to adapt to changes in reward structure, and indeed the changes of response patterns were gradual rather than immediate. Thus, our model produced similar results as with previous models (Keramati et al., 2011; Miller et al., 2019). It should be noted that the habit system or a non-zero α_H is not essential for producing the basic pattern. Even without habit formation ($\alpha_H = 0$), the changes in response pattern cannot be completely abrupt, as it takes time to update beliefs about reward probabilities (see Figure 3.8b). However, it was clear that with habit, the changes were much slower (in over 100 trials instead of only 30 trials).

Unlike Keramati et al. (2011), our model predicted gradual rather than sudden changes of decision time (measured as the sampling steps taken to make decisions, see Figure 3.8c). Also

¹³ To avoid confusion, it is important to note that extinction phase in Pessiglione et al. (2005) does not mean “no reward” as in other animal learning experiments, but only implies that the active go-response (button-pressing) is unlearned.

consistent with the empirical results (Pessiglione et al., 2005), decision time after the extinction phase increased about 1/2 less than after the extinction phase, because in the extinction phase reversal only applied to one stimuli.

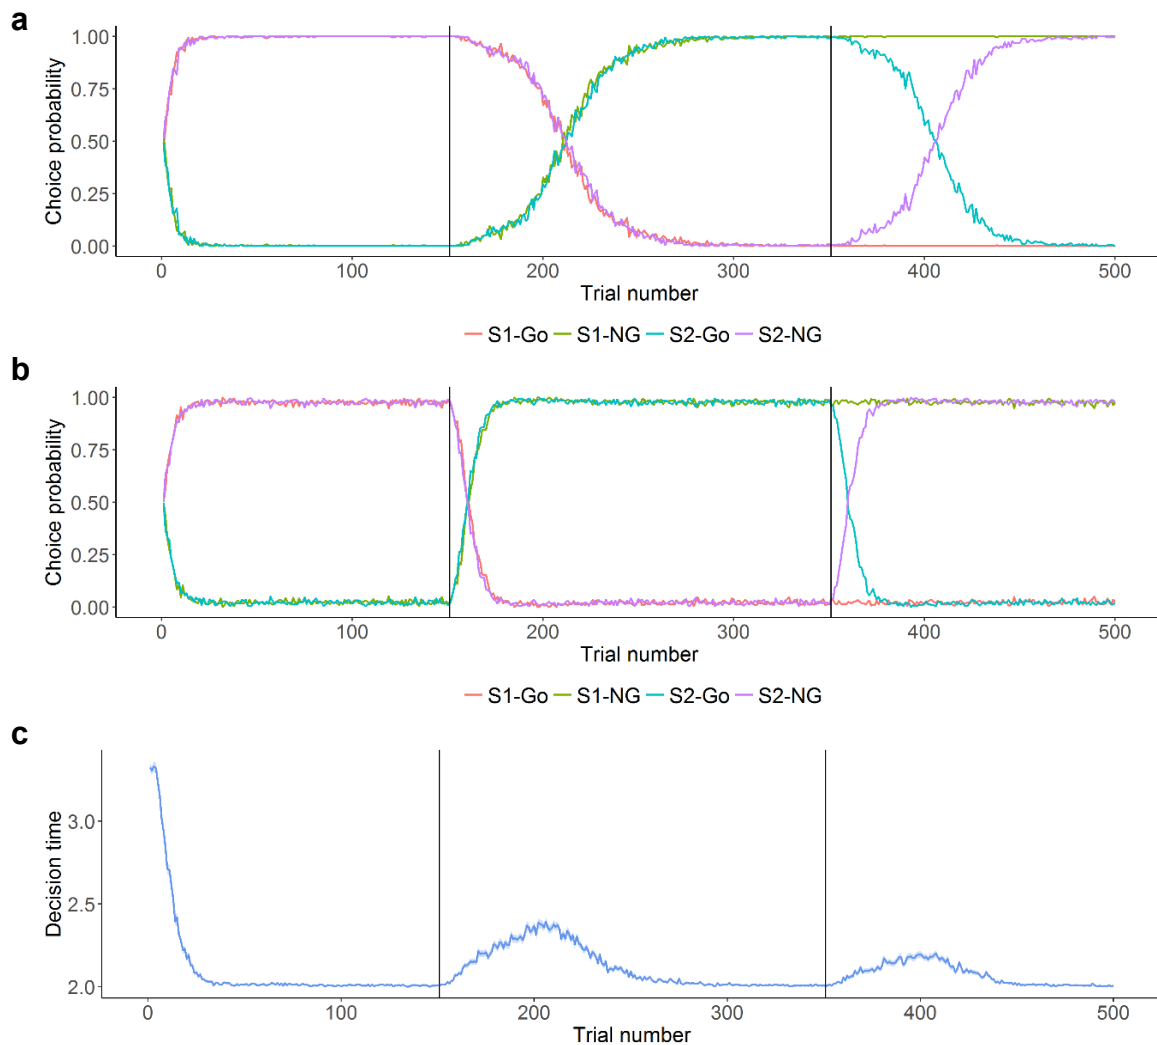


Figure 3.8 Simulation results of reversal learning. (a) Choice probabilities with $\alpha_H = 0.04$; (b) Choice probability with $\alpha_H = 0$; (c) Decision time with $\alpha_H = 0.04$.

3.4 General discussion

In this chapter, we have built a sequential sampling model to explain why strong habits prevent people from making decisions that satisfy their current goals. Simulation results have shown that our model can reproduce empirical results from three instrumental learning paradigms: classic devaluation, devaluation with a concurrent schedule, and reversal learning. This was achieved by a rather straightforward implementation of the multialternative decision field theory, with only two additional theoretical assumptions: (1) Starting positions of preference accumulation are determined by the habit values of behavioral responses; (2) At-

tribute sampling probabilities are based on the importance and task-relevance of the corresponding goals. The sensitivity analysis and the fact that the same parameters were used in all three studies speak to the strength of our central theoretical propositions.

3.4.1 Theoretical contributions

Contrasting previous models (Daw et al., 2005; Keramati et al., 2011; Miller et al., 2019), our work raises the possibility that instead of competing with each through a centralized arbitration, habits and goals may be integrated dynamically to produce behavioral responses. Sometimes, habits and goals are congruent and they jointly push responses in the same direction (e.g., during the learning phase of a devaluation experiment). In other cases, habit-goal conflicts emerge from the same process, when the goal-related attribute values become incongruent with the habit values obtained from prior behavior repetitions, for example, after goal devaluation or reward structure reversal. It remains entirely possible that habit strengths and goal-related attribute values are learned in distinct neural systems (Yin & Knowlton, 2006), but at the decision moments both value signals are integrated into a single decision-making circuit. This hypothesis should be evaluated in future neurophysiological research, preferably by incorporating existing insights about the neural underpinning of learning (e.g., Dolan & Dayan, 2013; O’Doherty, Cockburn, & Pauli, 2016; Yin & Knowlton, 2006) and of decision-making (e.g., Kable & Glimcher, 2009; Rangel, Camerer, & Montague, 2008; Shadlen & Shohamy, 2016; Summerfield & Tsetsos, 2012).

Our model shared two theoretical stances with Miller et al. (2019). First, both models separate goal values from goal-related attribute values, even though goal values as static decision weights in their model rather than the dynamic precursors of attribute sampling probabilities as in ours. This separation implies a double disassociation that devaluation only depletes goal values, while extinction test only affects goal-related attribute values. In contrast, other models implement both devaluation and extinction as changes to reward probabilities or directly to state-action values (Daw et al., 2005; Keramati et al., 2011). We believe that a separation is theoretically favorable, as it has been made in other theoretical frameworks (e.g., as *outcome value* and *outcome contingency* in learning theories, and as *decision weight* and *attribute value* in decision-making models), and there is evidence that they have distinct neural substrates (Kable & Glimcher, 2009; Rangel et al., 2008). Second, our work adds to Miller et al. (2019) that for explaining classic findings in instrumental learning, a value-free view of habit (Miller et al., 2018; Pauli et al., 2018) is at least as effective as the previous value-based view of habit (Dolan & Dayan, 2013). Our work cannot directly evaluate the verisimilitudes of the two views, but the assumption of mapping habit values to starting positions in sequential sampling models is more consistent with Hebbian learning algorithms (value-free) than with model-free reinforcement learning algorithms (value-based) of habit learning (see Akaiishi et al., 2014).

Our model also has implications for the role of uncertainty and speed-accuracy trade-offs in instrumental learning. Some conceptualizations of uncertainty and speed-accuracy trade-offs have been made in earlier models (Daw et al., 2005; Keramati et al., 2011), but uncertainty was computed as a higher-order mathematical property, such as variance of distributions. Rather, uncertainty is realized in our model as the sampling of values from distributions and as the stochastic process of preference accumulation. In addition, speed-accuracy trade-offs are naturally incorporated in any sequential sampling model (e.g., Ratcliff & Rouder, 1998), as more accumulation steps reduce uncertainty but lead to longer decision times.

Furthermore, the idea of sampling values from distributions for decision-making coincides with the Thompson sampling approach of solving repeated decision problems (Bandit problems), which usually achieves optimal balance between exploration and exploitation (Russo et al., 2015). Thompson sampling can be seen as a special case of sequential sampling with only one step. In this sense, sequential sampling with more than one step would favor exploitation more than exploration, depending also on the decision threshold. By shifting starting positions closer to threshold, strong habits further enhance exploitation. In contrast, unattainable attributes in our model provide a mechanism against over-exploitation, since the under-explored responses tend to have higher mean expected values for those attributes (see Figure 3.5f). In the events of sudden environmental changes (e.g., devaluation of primary goals), this mechanism counteracts habits to promote exploration. Future research should examine the role of habits in the exploration-exploitation dilemma and in reverse the role of the dilemma in instrumental learning.

3.4.2 A note on a critical pattern in the devaluation effect

A strong notion of insensitivity to devaluation would require a completely equal response rates between the devalued and non-devalued groups after extensive training. This critical pattern has indeed often be found in the literature (e.g., Adams, 1982; Dickinson et al., 1983; Killcross & Coutureau, 2003; Yin et al., 2004). As discussed earlier, the central mechanism of habits shifting starting positions should always produce a strong interaction effect, but not necessarily the exactly same response rates for the two groups of extensive training (see Figure 3.3b). The difference between the response rates of the two groups depend on the temporal dynamics of the expected mean rewards of all attributes, which may be altered by changes of various parameter values. For example, Figure 3.5e shows that larger uncertainty parameter γ may produce smaller differences between the two groups.

It should be noted that the critical pattern is potentially also a challenge for previous computational models (Daw et al., 2005; Keramati et al., 2011; Miller et al., 2019). Actually, none of the three models mentioned here addressed this problem directly, because they mainly compared response rates before and after devaluation, but not the relative response rates in the devalued and control groups after devaluation as usually reported in the empirical studies.

Whether or not these models can produce the critical pattern is beyond the scope of this chapter.

3.4.3 Model predictions and future work

Besides recreating key empirical results, the simulation studies produced full temporal dynamics of behaviors and underlying cognitive variables. As not all variables are observable and many specific local patterns may be subjected to parameter values, we only discussed several obvious model predictions. First, as an important strength of the sequential sampling approach, our model predicts gradual changes of decision time throughout the course of learning, rather than abrupt changes (see Keramati et al., 2011). The basic property of gradual change can be evaluated in trial-based instrumental learning experiments, from which decision times for all trials are plotted. Second, a recurrent pattern of decision time was that after devaluation it increased greatly regardless of whether extensive training lead to strong habits. This is because although strong habits and high starting positions can make decision fast, losing the positive drifts from sampling *food* (because of either devaluation or extinction) results in longer decisions. In contrast, Keramati et al. (2011)'s model predicts no such increase immediately after devaluation because habit system is assumed control actions completely. Thus, examining change of decision time following devaluation can provide a strong inference (see Platt, 1964) for comparing the two models. Third, for the study with concurrent schedule, the model predicts a brief increase in choice probability of the non-devalued response (see Figure 3.6a). This is consistent with Keramati et al. (2015)'s model and with the decision-making phenomenon that decrease of choice share in one option usually lead to increase of shares in all other options. Future studies are needed to further test this prediction.

Our model should be extended to other empirical scenarios in the future. One important example is the influence of reinforcement scheduling on the basic devaluation effect. It is well-known that variable-interval schedule promotes habit formation and thus insensitivity to devaluation than variable-ratio schedule (e.g., DeRusso et al., 2010; Dickinson, 1985; Dickinson et al., 1983), and this effect has been simulated by Miller et al. (2019). Adapting the sequential sampling approach to this scenario would require a continuous-time transformation of our model. Another interesting effect but in the domain of choice reaction time is the Hick's law (Hick, 1952) simulated by Keramati and colleagues with their model (2011). Because Hick's law has been accounted by a different sequential sampling model (Usher, Olami, & McClelland, 2002), our model should in principle explain it as well. Finally, as discussed by Keramati et al. (2011), novel scenarios using concurrent schedule might provide strong tests for different models, for example, two responses with one identical outcome, and two responses with two outcomes that sufficiently differ in their values. Our model would predict sensitivity to devaluation in the former scenario, but insensitivity in the latter scenario.

3.4.4 Conclusion

Our modeling work provide a preliminary demonstration of the sequential sampling approach's ability to explain the integration of habits and goals. It sheds new lights on the old problem of habit-goal conflicts and encourages a more uniform approach to learning theories and decision-making theories. The answer to the question of why people chose habitual but inferior options may indeed be explained by the dynamic interactions between the two forces. More broadly, our work extends an emerging research line of applying sequential sampling models to human reinforcement learning (Fontanesi et al., 2019; Frank et al., 2015; Pedersen, Frank, & Biele, 2017), and encourages a more unified approach to learning and decision-making theories in psychological science.

In terms of the adaptive decision-making framework, our sequential sampling model can be considered as a model of the action-level processes. Because sequential sampling models are extensively used for modeling human value-based decision-making, our approach helps to connect basic instrumental learning research to human habits in real-life context (see Marien, Custers, & Aarts, 2019). Of course, to extend the model's scope to people's daily lifestyle decisions, at least the process of option generation needs to be added. The next chapter explores this possibility in detail and also tests the habit formation part of the model in a real-world intervention application.

Chapter 4

Modeling Habit Development in Dental Behavior in the Real-World

4.1 Introduction

Research on habits, both empirical and computational, has traditionally focused on instrumental learning tasks in controlled laboratory settings (see Chapter 3). This line of research has generated robust findings on how animals and humans behave in those environments, and advanced the knowledge about how habit plays a role in people's pursuits of goals (e.g., Dickinson, 1985; Yin & Knowlton, 2006; Dolan & Dayan, 2013). However, for the purpose of understanding lifestyle behavior change and informing digital interventions, the knowledge learned from laboratory studies has to be generalized to habit developments in people's daily environments. For two reasons, it is unlikely that knowledge learned in controlled environments can be directly applied to the real-world.

First, the temporal span of habit formation in the lab is relatively short: trials are typically separated by just a few seconds and the whole learning tasks take a few hours or at most several days. In contrast, most lifestyle behaviors are repeated daily, and strong habits may take weeks or even months to form. The much longer gaps between actions imply that different processes may underlie learning in the two very different environments.

Second, whereas behavioral options are provided in laboratory experiments, in daily lives, people need to recall good options from their memories or to look for new options from external sources. As in the adaptive decision-making framework (Chapter 2), strong habits may facilitate the option generation process. As habits are cue-behavior associations, encountering certain environmental cues may trigger the associated behavioral options (Tobias, 2009; Wood & Neal, 2007). Since this cognitive mechanism is not used in laboratory tasks, it has to be understood by studying habit formation in the real-world.

In order to bring habit research closer to real-world problems and applications, in this chapter we investigate how people form a new dental habit in their daily lives. We chose to study dental behavior because of its relative simple cue-behavior associations (e.g., always performed before going to sleep) and its regularity (e.g., twice a day). Before turning to our own

investigation, we review existing behavioral and computational literature on habit formation in the real-world to motivate our research questions.

4.1.1 Literature review

Measuring habits in the real world

To study habits in the real-world, the construct of habit needs to be measured. This implies an assumption that habits are real entities at a cognitive or a neurological level and their strength can be represented by numbers. In the instrumental learning literature, habit strength has never been measured independently, but is inferred from the behavioral pattern that behavioral responses become insensitive to goal devaluation after extensive training (see Chapter 3). Ever since the discovery of the classic devaluation effect (Dickinson, 1985), it is believed by learning theorists and social psychologists alike that at a cognitive level, habits are mental associations between environmental cues and behaviors, and the associations can vary in their strengths (Gardner, 2015; Wood & Neal, 2007). With this assumption, several different approaches have been used by social and applied psychologists to measure habit strength (for a review, see Gardner, 2015).

The earliest method is to measure habit strength based solely on self-reported behavior frequency. This approach follows directly from the notion that habit formation requires behavior repetition, as in Thorndike's law of exercise (Thorndike, 1932; also see Hull, 1943). In the 70s and 80s, when habit was largely neglected in behavioral theories, this method helped to bring some attention to the construct of habit by showing that past behavior frequency predicted future behavior in addition to attitude and behavioral intention (e.g., Bagozzi, 1981; Landis, Triandis, & Adamopoulos, 1978; Triandis, 1977). However, using past behavior as a proxy for habit strength has been criticized by many theorists (e.g., Beck & Ajzen, 1991; Mittal, 1988; Ajzen, 2002), because such a measure conceptually confounds the habit construct with the behavior it is supposed to explain. A slightly improved method recognizes the fact that habit formation requires not just any behavior repetition, but repetition under a stable environment. Specifically, the method includes an additional question about context stability (e.g., "*when you do X, how often is cue Y present?*"), and multiply the score of context stability and behavior frequency to form a measure of habit strength (e.g., Ji & Wood, 2007; Ouellette & Wood, 1998).

A more defining characteristic of habits omitted in the behavior-based measures is the feeling of automaticity or uncontrollability when a behavior becomes habitual. This characteristic is also implied in the most corroborated finding in the habit literature – the devaluation effect that habitual behaviors do not respond immediately to goal changes (Dickinson, 1985). More recent approaches of measuring habits were all developed to address the automaticity character. Verplanken and colleagues constructed a measure based on fast responses to behavior-triggering scenarios (Aarts, Verplanken, & van Knippenberg, 1997; Verplanken, Aarts, & van

Knippenberg, 1994, 1997). For example, to measure habit strengths of using different travel modes, participants would be asked to respond to a series of typical travel scenarios (e.g., going to the beach with some friends) by reporting the first mode coming to their minds given a set of options (e.g., cycling, driving, walking, etc.) (Verplanken et al., 1994). The concept of automaticity is captured under the assumption that habitual options are recalled more easily and are thus considered first (see Psarra, 2016). A major drawback of this approach is that the administration of the measure requires a structured interview setting, which is usually not feasible in field studies.

Self-report measures that focus the automaticity character of habits have also been developed and used extensively. In several earlier studies, single-item measures of automaticity were added alongside behavior frequency items – for example, “*when I got into my car, I was not even aware that I put on my seatbelt*” (Mittal, 1988), and the extent to which participants felt they were using certain transportation modes by “*force of habit*” (Aarts & Dijksterhuis, 2000). A more comprehensive and structured scale, called the *Self-Report Habit Index* (SRHI), was developed by Verplanken & Orbell (2003), which consists of 12 items that tap onto the automaticity, behavior frequency, and self-identity aspect of habits. The automaticity sub-scale of the SRHI uses 4 items that describe the feeling of automaticity with different expressions (e.g., “Behavior X is something I do *without thinking*”, “Behavior X is something I start doing *before I realize I'm doing it*”), and requires participants to rate the applicability of the descriptions when they perform a target behavior. Some researchers have suggested to use the sub-scale of automaticity only (i.e., “Self-Report Behavioral Automaticity Index”, or SRBAI), as it provides a more parsimonious measure without sacrificing the convergent and predictive validity of the full SRHI (Gardner, 2012; Gardner, Abraham, Lally, & de Bruijn, 2012).

Using the self-report measures, scientists are starting to know a fair bit about the dynamics of habits in people’s daily lives. For example, by monitoring how people developed self-selected new habits for 12 weeks (e.g., drinking water during lunch, going to bed early), it was shown that habit changes followed asymptotic curves (i.e., they grew fast in the beginning and then the growth decelerated) and on average 66 days were required to reach the asymptotes (Lally et al., 2010). Another finding from the same study was that habits decayed only slightly if the target behaviors were omitted for one day. Similar asymptotic curves for habit strengths have been found in a habit-formation intervention study on dietary behavior change, based on self-assessments at 6 time points in about 3 months (Gardner, Sheals, Wardle, & McGowan, 2014). Finally, by tracking workers’ habit strengths for different travel modes before and after an office relocation, it was shown that the growth of new habits and the decay of old ones occurred in parallel, and even when the old modes were rarely used, weak habit strengths usually remained (Walker et al., 2015).

The effects of habit strength and attitude on behavior

A particular interest in social and applied psychology is the influence of habit strength on behavior, in addition to the more goal-related constructs such as attitude and behavioral intention. The original interest in habit in social psychology came from a discontent with the over-emphasis on rational or goal-related constructs in behavioral theories (e.g., the Theory of Planned Behavior) and the recurrent finding that past behavior predicted future behavior on top of the goal-related constructs (Bagozzi, 1981; Landis et al., 1978; Triandis, 1977). Triandis (1977) was the first to formally propose that the probability of performing a behavior is partially a weighted function of habit and behavioral intention, and the weights for the two distinct determinants are supposed to inhibit each other, thus implying a negative habit by intention interaction effect. Such a proposal is certainly consistent with findings from the instrumental learning literature on the interplay between habit and goal-directed control (Dickinson, 1985; Yin & Knowlton, 2006), and with the more general theorizing of dual-system accounts of behaviors in psychology (e.g., Kahneman, 2003; Strack & Deutsch, 2004).

However, the early findings based on behavior frequency cannot provide strong support for a distinct role of habit, because statistically any variance in behavior not explained by attitude and intention may be autocorrelated in time. Thus, by measuring automaticity using the scenario-based method mentioned above, Verplanken et al. (1994) provided the first clear evidence in the context of travel mode choice that: (1) Habit strength accounts for variance in behavior in addition to attitude; (2) When habit is strong, the positive effect of attitude on behavior is attenuated (i.e., a negative moderation effect). Since the invention of SRHI and SRBAI, many more studies in various behavioral domains have demonstrated the independent influence of habit strength on behavior and a negative interaction effect between habit and intention/attitude (e.g., fruit consumption, de Bruijn et al., 2007; physical activity, van Bree et al., 2015; sun protection use, Allom, Mullan, & Sebastian, 2013; for a review, see Gardner, 2015).

Despite the relatively robust demonstration of the effect (for some negative results, see Murtagh et al., 2012; Norman, 2011; Gardner et al., 2012), there are two limitations in the previous research. First, in almost all previous studies, behavior frequency as the main dependent variable was self-reported (with only one exception in 23 studies reviewed by Gardner, 2015). Besides the commonly recognized limitations of self-report measures, such as memory biases and social desirability, if both behavior and behavioral determinants are self-reported, common-method bias may inflate the true correlations between these variables (see Podsakoff, MacKenzie, Lee, & Podsakoff, 2003). In fact, the only study that used an objective measure of behavior (walking habits) did not find the significant interaction effect (Murtagh et al., 2012). For these reasons, measuring behavior objectively would consolidate the earlier findings, and provide a more accurate estimation of effect sizes. Second, most studies employed

either a cross-sectional or prospective design, where both behavior and behavioral determinants were measured only once (for one exception reviewed in Gardner, 2015, see Conroy, Maher, Elavsky, Hyde, & Doerksen, 2013). These designs cannot be used to distinguish between-person and within-person effects. This implies that the findings reviewed above do not necessarily answer questions about within-person processes (see Borsboom, Mellenbergh, & van Heerden, 2003; Molenaar, 2004), for example, whether someone would perform a behavior more often when a stronger habit is formed. Answering this type of questions requires time-intensive longitudinal designs (see Dunton & Atienza, 2009), such as daily diary or experience sampling method.

The effects of behavior repetition on habit formation

The relationship between habit and behavior is not unidirectional but reciprocal. As reviewed in Chapter 2, the change of habit strength depends on the repetition of behavior. However, only until recently, empirical data have been collected to test this relationship in daily environments (de Bruijn, Gardner, van Osch, & Sniehotta, 2014; Fleig et al., 2003; Gardner & Lally, 2013; Verplanken & Melkevik, 2008; Wiedemann, Gardner, Knoll, & Burkert, 2013). Among these studies, the study on exercising habit (Verplanken & Melkevik, 2008) had the best study design in our opinion, because a prospective design was used in which habit strengths were measured at two time points (T1 and T2) and exercising behavior in-between the two points (a 1-month gap) was self-reported. By controlling for baseline habit strength at T1, the correlation between behavior frequency during the month and the habit strength afterwards (T2) offered an estimate of the effect of behavior repetition on the change of habit strength. In contrast, other studies suffered from one or more methodological issues, with either a lack of control for baseline habit strength (de Bruijn et al., 2014; Gardner & Lally, 2013; Wiedemann et al., 2013), or a mismatch between the temporal spans of the behavior measures and the timing of habit measures (de Bruijn et al., 2014; Fleig et al., 2003; Wiedemann et al., 2013).

There are other limitations to all the previous studies. As with the research on the influence of habit on behavior, the above studies only used self-reported behavior as the predictor, and their prospective design did not allow for a separate estimate of within-person effects of behavior on habit strength over time. In addition, none of those studies were done in the context of habit formation or in behavior change intervention trials, so new data are required to test if trained behavior repetitions would indeed lead to strong habits.

Computational modeling of habit formation and option generation

Although Thorndike's law of exercise was proposed nearly a century ago (Thorndike, 1932), it was not until recently when researchers started to describe the dynamics of habit formation computationally. In Chapter 3, we have used the model of habit learning by Miller et al. (2019)

in our sequential sampling approach to habit-goal integration. Here, we provide a more detailed review of a few recent computational models of habit formation (Klein, Mogles, Treur, & van Wissen, 2011; Miller et al., 2019; Psarra, 2016; Tobias, 2009). All models are inspired by the theoretical principle of Hebbian learning (Hebb, 1949), which states that neurons fire at the same time tend to form stronger connections among themselves. For habit formation, this principle implies a theory at the cognitive level: when nodes representing a behavior and an environmental cue are often activated simultaneously, a stronger link is formed between the two nodes. Figure 4.1 shows the exact mathematical equations used in these models, and their simulation results in a simple scenario of habit development. To show how habit strength increases and decays in these models, the scenario assumes that a new behavior is consistently performed in the first 60 time steps before it is abandoned completely. Using reasonable parameter values for each of these models, the general temporal patterns share great similarities and are also consistent with the empirically observed asymptotic curves in self-reported habit formation (Lally et al., 2010).

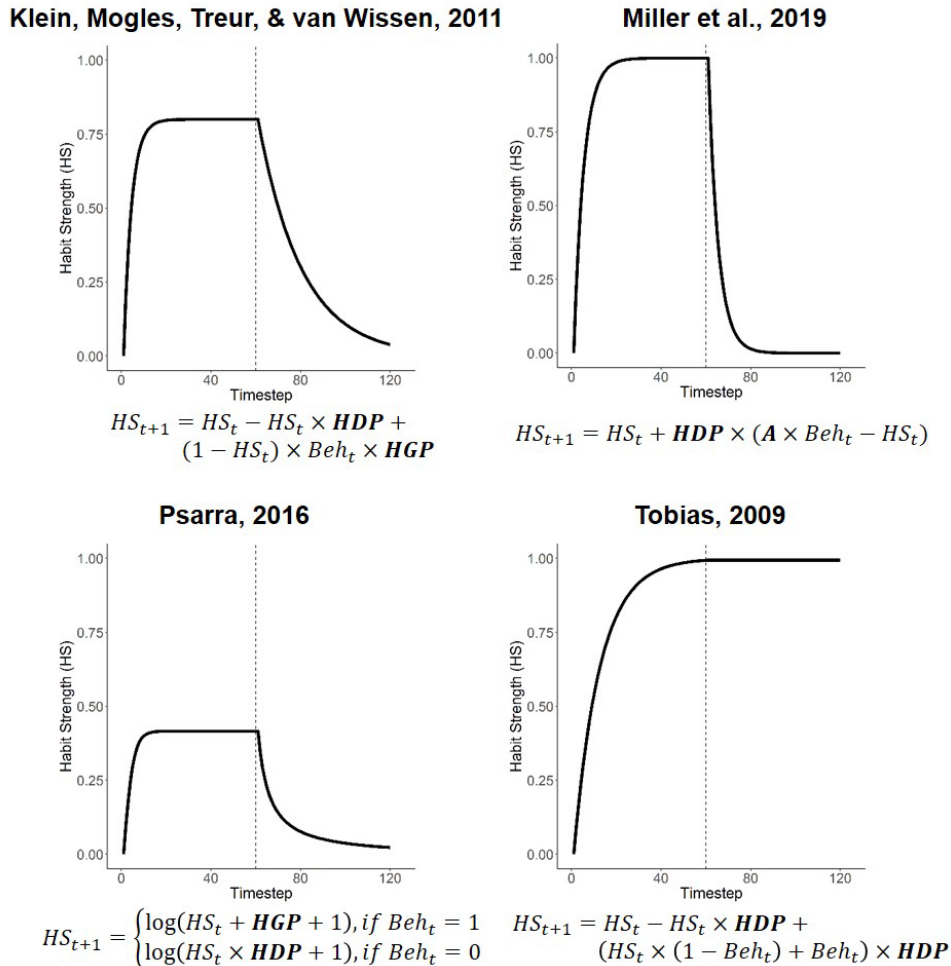


Figure 4.1 Temporal dynamics of habit strength in different computational models (HS = habit strength; Beh = behavior; HDP = habit decay parameter; HGP = habit gain parameter; A = scaling parameter in Miller et al., 2019).

There are also some important differences among the simulation results, especially in terms of at what value habit strengths reach an asymptote and at what rate they decay. The equation by Klein, Mogles, Treur, & van Wissen (2011) and Miller et al. (2019) are very similar, except that the former uses separate parameters for the rates of habit gain and decay, while the latter uses a single parameter for both rates. Because of this difference, habit strength in Miller et al. (2019) is bounded between 0 and 1, but it does not allow the empirically observed difference between the growth and decay rates (Lally et al., 2010). The model by Tobias (2009) is also bounded between 0 and 1, but habit strength under its formulization won't decay even if behavior is not performed, which contradicts with theories, observations, and common sense about habits. Finally, Psarra (2016)'s model uses a logarithmical transformation between previous and current habit strengths, so the value range depends on the parameters. Also, separate equations are used for cases where a behavior is performed or omitted, and thus the parameter for decay rate is not consistently applied to the two cases.

As discussed in Chapter 2, the process of option generation is influenced not only by habit values, but also by the activation values of options themselves. Conceptually, as with any memory process, activation values of behavioral options fluctuate in time due to many factors (e.g., memory rehearsal, external reminders), while habit is a cue-dependent mechanism to increase activation values. In his work on modeling recycling behavior in a field setting, Tobias (2009) modeled activation value or what he called *memory accessibility* computationally¹⁴. Figure 4.2 shows a typical trace of the memory accessibility of a behavior option according to Tobias's model – it decays at a fixed rate of 0.8 but is occasionally enhanced slightly by executing the behavior and greatly by receiving a reminder for the behavior. According to the model, for a behavioral option to be generated at a decision moment, its accessibility has to be higher than a threshold value, which decreases when the habit strength for the option is strong. This is mathematically equivalent with summing an option's accessibility value and habit strength value to obtain a total activation value, and then compare the total activation value to a constant threshold. Overall, the option generation model complements with our sequential sampling model of option evaluation in Chapter 3, and together they form a more complete model of daily lifestyle decisions.

¹⁴ Though not to be discussed in detail, Psarra (2016) proposed a construct called awareness of options, in addition to the construct of habit strength. This is similar to what we define as activation value here. However, awareness was defined by Psarra as cue-dependent, which make it less distinguishable from habit strength.




Figure 4.2 A typical trace of memory accessibility according to Tobias (2009)’s model.

4.1.2 The current investigation

In the current investigation, we aim to build on previous research to model the habit formation process in two field intervention studies of dental behavior change, focusing on both behavior explanation and prediction. For the explanation part, the studies were designed to address the limitations identified in the literature review by employing a multiple-week longitudinal design and an objective measure of toothbrushing behavior using sensors. More specifically, we intend to answer the following questions about the reciprocal relationships between habit strength, attitude, and dental behavior:

1. How do attitude and habit strength jointly influence toothbrushing behavior, both inter-individually and intra-individually?
2. Does strong habits attenuate the positive effect of attitude on toothbrushing behavior? (the habit-attitude interaction effect)
3. How do attitude and habit strength change over time? And are these temporal changes influenced by behavior repetition itself?

For behavior prediction, we are interested in whether the computational models of habit formation and option generation can improve the prediction of toothbrushing behavior. This can be seen as an application of the psychological computing approach outlined in Chapter 1. Not only useful by itself, improved prediction can also serve as a validation for the computed cognitive states, which may be used to inform interventions delivered by digital systems.

Because the two studies are similar in design, we discuss them together in the next section. Then we report in two separate sections about the behavioral results for answering the explanatory questions, and then the method and results of applying the theory-based computational models for behavior prediction.

4.2 Data collection

4.2.1 Participants

Study 1

Forty healthy university students or young workers were recruited through a local participant database and personal network. The main inclusion criterion was that they used to only brush their teeth once a day (or at least rarely brushing twice), and the criterion was checked by personal communication with the participants. The sampling consisted of 26 males and 14 females, and the average age was 24.48 ($SD = 3.13$, median = 24). Eight participants were randomly selected and awarded 25 euros. The study was reviewed and approved by the ethical review board of our department.

Study 2

Study 2 was conducted in collaboration with Philips Research. Seventy-nine adults were recruited through a recruitment agency contracted by Philips. A lenient main criterion was used that the participants used to brush only once a day, or they usually brushed less than two minutes for each session. Other criteria include that they were between 18 and 60 years old, understood Dutch, and were manual toothbrush users. The eventual sample consisted of 41 females and 37 males (1 chose “other”), with ages between 20 and 63 years old (mean = 39.63, median = 38, $SD = 10.97$). Most participants were healthy, except that one suffered from cystic fibrosis and one from narcolepsy. Participants were paid 80 euros by the recruitment agency. Compared with the highly-educated sample in Study 1, participants in Study 2 possessed more varied education levels, including high school (15%), secondary vocational education (middelbaar beroepsonderwijs or MBO in Dutch) (37.5%), Bachelor (28.75%), Master (15%), and Doctoral degree (1.25%). The study was reviewed and approved by the Internal Committee on Biomedical Experiments (ICBE) at Philips Research.

4.2.2 Design & procedure

Study 1

Participants were enrolled in a 4-week intervention program during which they were persuaded to change their oral health routine from brushing teeth once a day to brushing twice a day. The main behavioral variable was whether they complied with the new target brushing behavior (i.e., brushing also in the morning or in the evening) on each day during the study period. Figure 4.3 shows the events that occurred during the study. At the beginning, a face-to-face meeting was held between the experimenter and each participant, where the study information and intervention goal were explained, the sensor used for behavior monitoring was attached to the participant’s toothbrush, and a consent form was signed. After participants returned home, their toothbrushing behaviors were monitored by the sensors for 3

weeks, and at the end of the third week they returned the sensor to the experimenter. Reminders for the target brushing behaviors were sent daily in the first week using a self-programmed mobile app, every other day in the second week, and were dismissed in the third and fourth week. At five time points in between the weeks ($t_1 - t_5$), short surveys were sent using the same app to ask questions about attitude and habit strength. In case that participants missed a survey in the mobile app, the same survey was sent to them through e-mails.

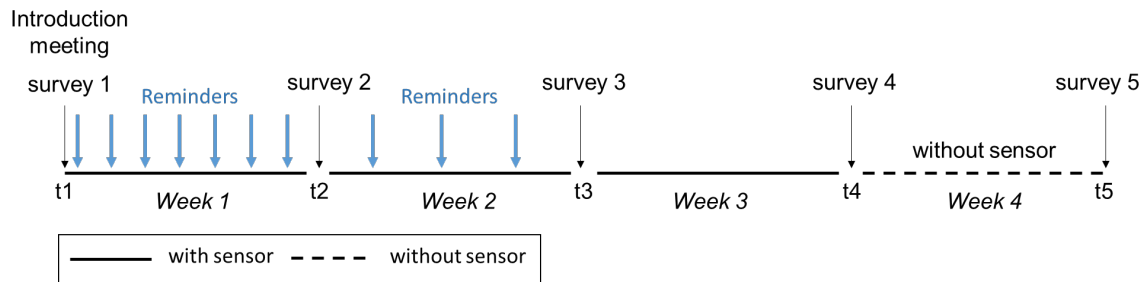


Figure 4.3 Timing of events in Study 1.

Study 2

Because a lenient criterion on participants' existing toothbrushing habits was used, the primary intervention target focused not only on daily frequency (i.e., twice a day), but also on the duration of brushing. The goal of the intervention was thus to develop an optimal oral health routine of two brushing sessions that last for at least 2 minutes (or at least a 4-minute brushing daily). Despite this difference, the main behavior variable used in data analyses was still whether participants complied with brushing twice a day or not. Besides a consideration of consistency, there were two other reasons why the behavior measure was based on occurrence rather than duration. First, habits, at least in the sense of the theories discussed in this thesis, were about decisions, or in other words the initiations of behaviors, rather than the execution characteristics, such as duration. It was not clear from the existing theories how habits as cue-behavior associations would affect the duration of a behavior, in addition to automatically triggering its initiation. Second, the sensors (accelerometers) used in the studies were accurate enough for detecting brushing episodes but not for estimating the durations of episodes.

The study procedure generally broke down to three phases: a baseline calibration period, a laboratory intervention session, and a follow-up monitoring period (Figure 4.4). At the beginning, participants came to the lab in groups of 10-15 for an introduction session, in which general study information and procedure were explained, but not the specific intervention (so they only knew it was about oral health). Also in the meeting, participants were offered new manual toothbrushes with sensors attached, and were asked to sign a consent form and to complete the first survey. After the baseline period of about 5-10 days, they were invited back to the lab for the intervention session individually. They were shown presentations

4.2 Data collection

about oral healthcare, and were exposed to the intervention target of brushing twice a day for at least 4 minutes. During the lab session, physiological data from the participants were recorded for purposes not related to this thesis. The second and third survey with mostly identical questions was completed by the participants before and after the lab session. After the lab session, participants returned home and were monitored for a follow-up period that led to a total of approximately 3 weeks. Two additional surveys were sent by e-mails in the middle and at the end of the follow-up period.

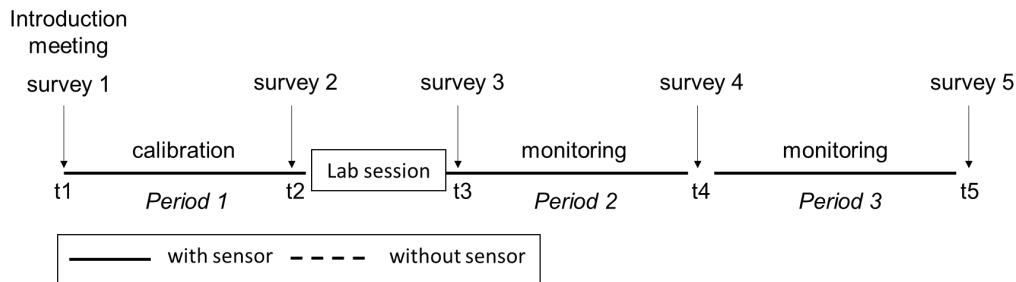


Figure 4.4 Timing of events in Study 2.




Figure 4.5 Number of days in each study period for each participant (each row).

Figure 4.5 shows the eventual number of days in each period of the study, separated by time when participants completed the 5 surveys (two surveys at t2, the day of the lab session). It was clear that the number of days in each period varied moderately, which was mainly due to logistic reasons imposed by the individual lab sessions. Overall, the design of Study 2 was not

as optimal as in Study 1, due to its embedment in a larger project at Philips research, but it provided a larger and more diverse sample.

4.2.3 Intervention techniques

Although this Chapter does not focus on evaluating intervention techniques, different combinations of behavior change techniques were used in the two studies. Therefore, it is worth describing these techniques in some detail, in order to better interpret and compare the results later on.

Study 1

During the short face-to-face meeting, a combination of health education and implementation intention was used. First, participants were shown an educational booklet and read through the reasons why they should switch to brushing teeth twice a day (e.g., “to prevent gum disease”, “to maintain a fresh breath”, and “to reduce chances of getting a heart attack or stroke”). Next, participants were told the benefits of implementation intentions for forming new habits, and were asked to complete the sentence “*when _____, I will go and brush my teeth*” at the back of the booklet with environmental cues they chose for themselves, such as “*when I get off my bed in the morning*” or “*before I go to sleep*”. They were also asked to imagine the written scenario in their minds for 2 minutes, and then to recite the sentence aloud.

The reminders sent during the first two weeks were of two kinds, depending on a random assignment of conditions to the participants. Some received general reminders that were sent in the middle of the day, which simply asked them to remember performing the target brushing behavior at appropriate time. Others received just-in-time interventions either at their indicated wake-up time or 30 minutes before their personal sleep time, which suggested them to brush teeth right away (for functional differences between the two types of reminders, see Chapter 2).

Study 2

In the calibration week, participants were not told to change their dental behaviors. During the lab session, they were invited to a professional oral healthcare laboratory at Philips Research for the following procedures. First, they were asked to perform a plaque-disclosing test and reflected on the results (the disclosed dental plaque with colors), in order to increase their awareness of their dental health status and the self-relevance of the study. Next, they were randomly assigned to a gain-framed or a loss-framed persuasion condition (see Tversky & Kahneman 1981), and were shown 14 informational slides about optimal oral healthcare routine, with motivational messages focusing on either the positive effects of a good oral healthcare routine (e.g., clean and strong teeth) or the negative consequences of a bad routine (e.g., cavities and tooth decay). While processing the persuasive information in the slides,

4.2 Data collection

their physiological responses, including ECG, skin conductance, and blood pressure, were measured. The comparison between the two persuasion conditions and analyses of the physiological data were only of interests to the related larger research project at Philips. No other intervention techniques (e.g., reminders) were used.

4.2.4 Measurements¹⁵

Toothbrushing behavior

Participants' toothbrushing behavior was measured by the Axivity AX3 sensors attached to the lower-end of their toothbrush grips (see Figure 4.6). The Axivity AX3 sensor is a 3-axis accelerometer developed by Newcastle University specifically for scientific research in human movements (e.g., Doherty et al., 2017). Constrained by the memory space of the device, sampling frequency was set at 50 Hz to ensure the storage of data for three weeks. The sensitivity range for accelerations was set at $\pm 8g$. The sensor was water-proof, and a fully-charged sensor could work for 3 weeks without additional charges. Participants in both studies also self-reported on how many days of the previous week they brushed their teeth in the morning/evening (Study 1) or brushed teeth twice a day for at least 2 minutes each time (Study 2).




Figure 4.6 An example of a sensor attached to a toothbrush.

Habit strength

Habit strength was self-reported using the 4-item SRBAI with 7-point response scales. It assessed behavioral automaticity by prompting participants to rate their agreements with descriptions of performing a target behavior (e.g., “*Behavior X is something...*”), including “*I do automatically*”, “*I do without having to consciously remember*”, “*I do without thinking*”, and “*I start doing before I realize I am doing it*”. The target behavior in Study 1 was “*brushing teeth in the morning*” or “*brushing teeth in the evening*”, depending on which behavior was

¹⁵ Measures that were administrated but are not of interests to the current chapter include demographics (both studies), trait self-control (both studies), behavioral activation system and behavioral inhibition system scales (Carver & White, 1994; Study 2), and behavioral intention (Study 2).

not performed by each participant before the study. In Study 2, because of the lenient inclusion criterion, the behavior was more generally phrased as “*brushing teeth twice a day and in total at least 4 minutes*”. Internal reliabilities of the SRBAI were very high in both Study 1 (Cronbach’s $\alpha = 0.95$) and Study 2 (Cronbach’s $\alpha = 0.94$). These items were translated into Dutch in Study 2.

Attitude

Attitude was measured using 7-point semantic differential scales that were typically used in studies that followed the Theory of Planned Behavior (e.g., Verplanken et al., 1994). Four items were used in Study 1 (bad – good, useless – useful, harmful – beneficial, unpleasant – pleasant), while in Study 2 three more items were added (foolish – wise, unhealthy – healthy, difficult – easy). We also made a common distinction between instrumental attitude and affective attitude (e.g., see Tobias, 2009), because inter-item correlations and factor analysis clearly suggested two separate factors. Instrumental attitude focused on how a behavior satisfied instrumental goals, such as health benefits in the context of dental behaviors, while affective attitude taps more onto the emotional aspects of the experience relating to the behavior (e.g., comfort of brushing, effort spent on brushing). Affective attitude score was based on a single item in Study 1 (unpleasant – pleasant) and the average score of two items in Study 2 (unpleasant – pleasant, difficult – easy). Internal reliabilities (Cronbach’s α) for instrumental attitude were 0.94 and 0.93 for the two studies, while affective attitude also had a satisfying internal reliability of 0.71 in Study 2. The attitude items were translated into Dutch in Study 2.

4.2.5 Pre-processing

Pre-processing was performed to transform the raw 3-axis accelerometer data to behavioral data at the day-level (i.e., brushing twice or not on a specific day). The same procedure was used in both studies, which included the following steps: converting 3-axis signals to signal vector magnitudes (SVM), extracting brushing episodes, and classifying episodes to match day-level data.

Converting 3-axis signal to SVM

The first step was to compute SVM based on the raw three-axis accelerometer data, according to the equation below:

$$SVM = \frac{1}{n} \sum_{i=1}^n \left| \sqrt{x_i^2 + y_i^2 + z_i^2} - 1 \right|$$

SVM provided a summarized movement magnitude measure by combining the acceleration information from the x, y, and z axis, and down-sampling the 50 Hz raw data to magnitude measured at 1 Hz ($n = 50$ in the equation above). Figure 4.7a shows one participant’s data

after SVM-transformation, where each data point (dot) represents the average movement magnitude in each 1-second time window. This processing was done using a built-in SVM algorithm Open Movement v1.0.030, the default software for the Axivity AX3 sensor.

Extracting brushing episodes

From Figure 4.7a, it was clear that brushing episodes could even be visually identified (the spikes) when the data were clean, but not when there was noise caused by other movements. Given this problem, a threshold-based algorithm was first used to scan the data sequentially to efficiently extract all potential brushing episodes, and then a manual check was performed to exclude “invalid” episodes. The algorithm included the following steps:

1. While scanning the data, it detected the onset of a brushing episode when a data point exceeded a sample-level threshold on SVM. Because there were inter-individual differences in average movement magnitude (i.e., some people moved more intensively when brushing teeth), the sample-level threshold was personalized to be 3 times the mean SVM of the whole data set of each participant.
2. For detecting the end of an episode, the algorithm ignored small gaps that were less than 30 seconds. When a data point below the sample-level threshold was detected, the next 30 data points were checked. If any of the 30 data points exceeded the threshold, the integration of the current episode continued. Otherwise the integration was concluded and the search for the next episode started.
3. When a brushing episode was identified, certain features about the episode could be computed, including start time, duration, mean SVM, standard deviation of SVM, and even spectral information. Only start time, duration, and mean SVM were used for filtering and classifying the episodes in the current research.
4. After identifying all the episodes, two criteria were used to filter out some “invalid” episodes. First, only episodes with durations between 30 and 1000 seconds were included. Second, a participant-level threshold was used to exclude episodes with very weak movements, i.e., mean episode SVM less than 15 times the average SVM of all data points.

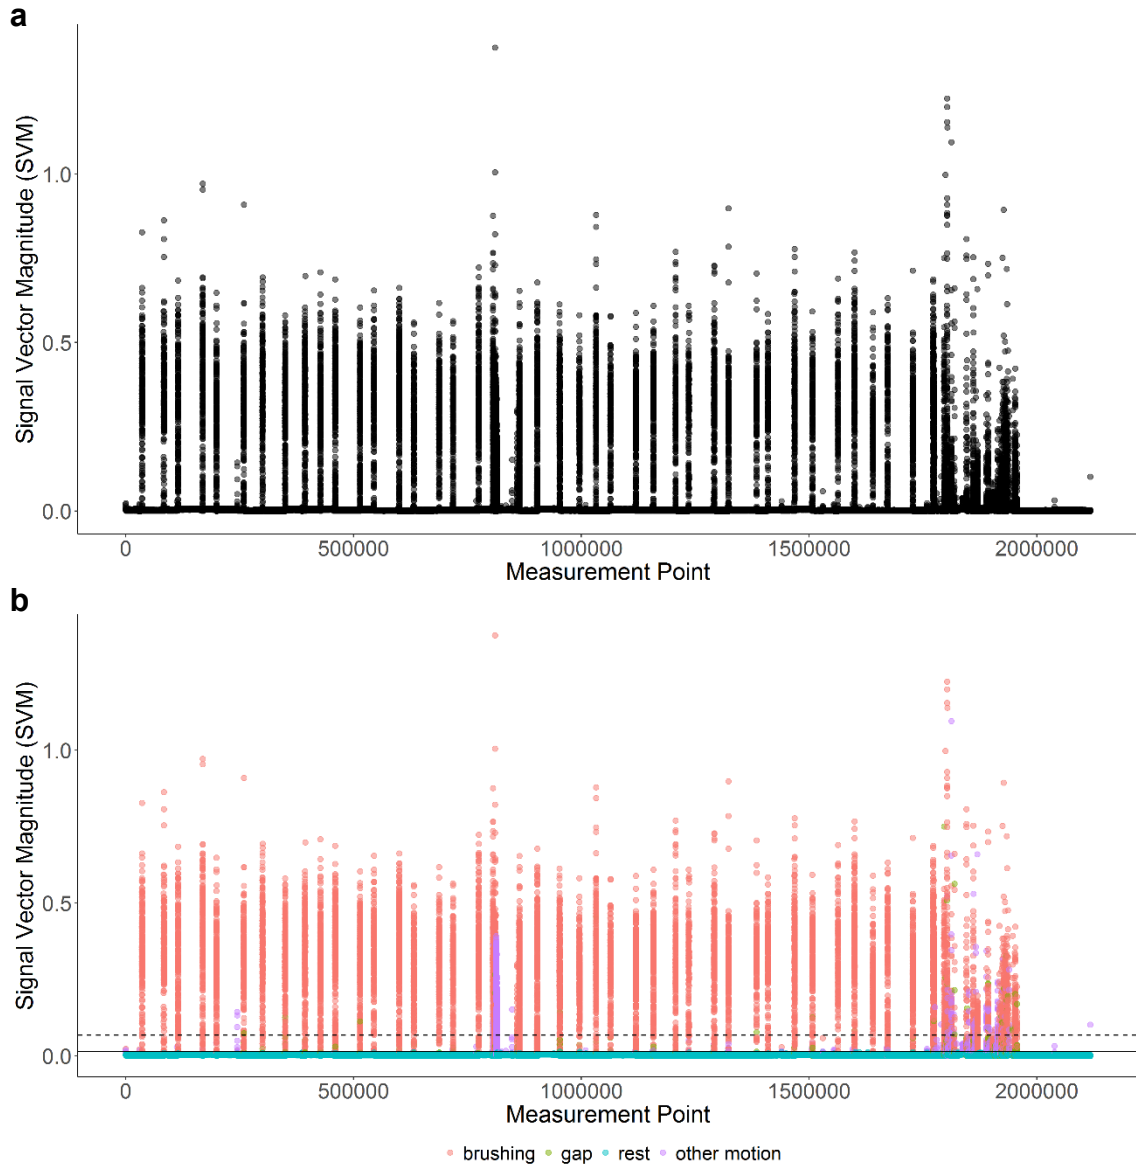


Figure 4.7 (a) Sensor data in signal vector magnitude; (b) Sensor data classified after applying the thresholding algorithm (solid and dashed horizontal lines represent the levels of sample-level and episode-level thresholds respectively).

After running the algorithm (see Figure 4.7b for the classified data at this step), a manual check was performed to remove false positives that might occur for several reasons: (1) long-duration movements other than toothbrushing, due to either known events (e.g., when participants brought the sensors back to the lab) or unknown events (e.g., vibrations due to a running washing machine); (2) short-duration episodes around true brushing episodes (with gaps larger than 30 seconds), probably reflecting preparation and cleaning behaviors; (3) some rare episodes at typical brushing time, but with very different characteristics from typical brushing episodes not captured by the thresholds used in the algorithm. The manual check removed 587 out of 1965 episodes identified from 40 participants (29.9%) in Study 1, and 1485 out of 4145 episodes

(35.8%) in Study 2. Although there was no way to formally assess identification accuracy without self-reports, we were confident that the procedure was quite accurate in terms of measuring the daily occurrence of brushing behavior, and the procedure was unbiased because the same criteria were used for all participants before any statistical analyses. Of course, as noted earlier, estimation of brushing duration was more problematic so it was not used in statistical analyses.

Classifying episodes to create day-level data

The remaining episodes were then classified into 6 categories based on the starting time of the episodes: *morning* (5:00 – 12:00), *morning-afternoon* (12:00 – 15:00), *afternoon* (15:00 – 19:00), *afternoon-evening* (19:00 – 21:00), *evening* (21:00 – 24:00), and *overnight* (0:00 – 5:00). The final episode-level data may contain more than one episode for each time category on each date. At the data level, two variables – *morning brushing* and *evening brushing* – were created, and their values (0 or 1) were determined by searching in the relevant categories on the same date to see if any episode existed. For *morning brushing*, category *morning* was searched first, and if no episode was found, category *morning-afternoon* was searched. For *evening brushing*, categories *evening* and *overnight* were searched first, and if no episode was found, category *afternoon-evening* was searched. When there were known or unknown events that caused noise in the data in a certain period, the values for the two brushing variables were coded as missing data. Eventually, at the day level, a dichotomous indicator (0 or 1) for *the target brushing behavior* and for *brushing twice* were used as the behavior measures in Study 1 and Study 2 respectively. Weekly (Study 1) or period-based (Study 2) *behavior frequency* and *behavior rate* (frequency divided by the number of days) were then computed to obtain the behavioral variables used in the statistical analyses.

4.3 Behavioral results

4.3.1 Results of Study 1

Data quality check

Response rate for the weekly surveys was 75% on average (see Figure 4.8a for the histogram). We also checked on how many days each participant had brushing data missing due to movement noise or technical faults, and the average proportion of days missing was 9%, but most participants had complete data (see Figure 4.8b). Based on the limited self-reported behavior data¹⁶, self-reported and sensor-measured behavior rates correlated strongly but not perfectly ($r(56) = 0.88$, Figure 4.8c), suggesting relatively good validities of both measures.

¹⁶ Only in 56 records both self-reported and sensor-based behavior rate were available, due to an administration error that the self-report item was not included in the web version of the surveys sent to the participants who missed the surveys in the mobile app.

Based on the quality-check, 4 participants were excluded for further analyses: 2 of them did not respond to any weekly surveys, 1 had missing data for the whole 1st and 3rd week, and 1 had sensor data with severe movement noise.

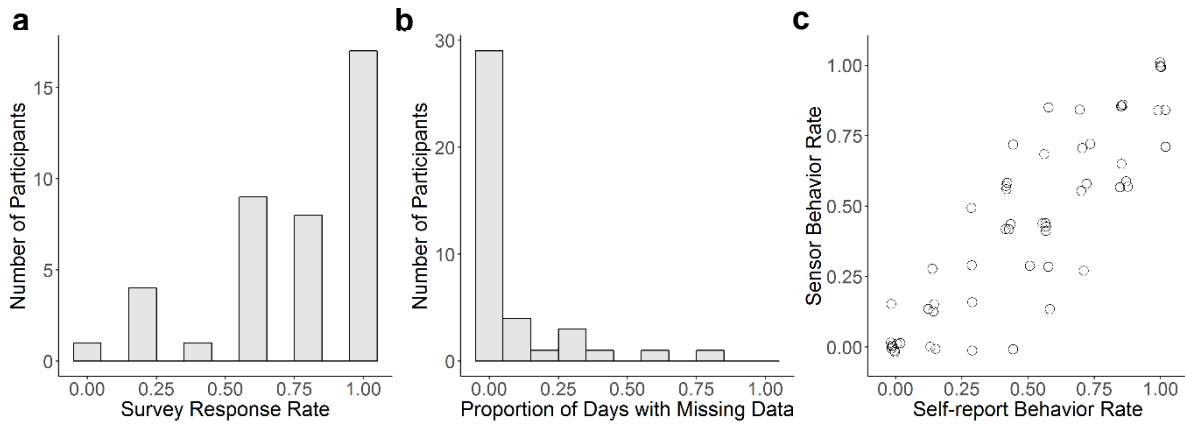


Figure 4.8 Data quality check for Study 1: (a) Histogram of survey response rate; (b) Histogram of proportion of days with missing sensor data; (c) Correlation between self-reported and sensor-measured behavior rate.

Description and modeling of changes

Before examining the aggregated change patterns over all participants, some insight into the data could be gained by plotting the change pattern of each participant. Figure 4.9 provides two examples of such plots for two individual participants, showing a change of behavior at both day and week level, and changes of attitude and habit strength at week level. Note that participant 38, who had very positive and stable attitude, performed the target brushing behavior consistently in the early period and self-reported habit strength increased from 1 to 3. In contrast, participant 39, who had more neutral attitude, behaved more randomly and no clear habit was formed.

4.3 Behavioral results

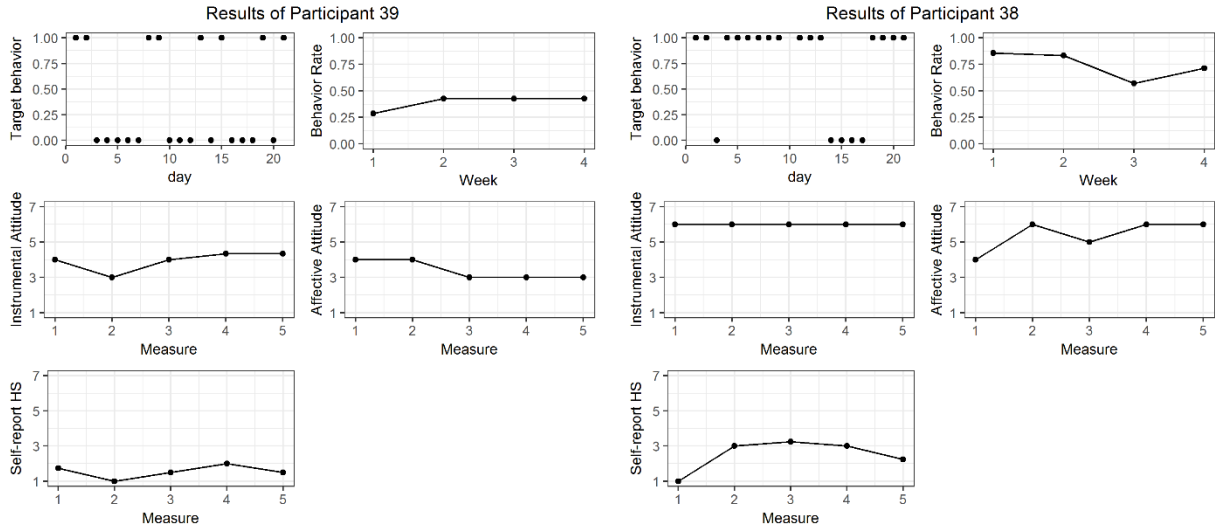


Figure 4.9 Examples of results for two individual participants in Study 1. (In each panel: top-left: change of target brushing behavior at day level (0 or 1); top-right: change of behavior rate at week level; middle: changes of instrumental and affective attitude at week level; bottom: change of self-reported habit strength at week level).

In order to analyze the aggregated change pattern of behavior rate, objective measures were used for the first three weeks, while self-report measure were used for the fourth week, in which the sensors were removed according to the study design. Besides describing the change patterns with line plots and histograms, for each variable, we also modeled a latent growth curve underlying its change by using a multi-level regression with time predictors (Raudenbush & Bryk, 2002), the 1st and 2nd order polynomials of week index or measurement index. Since adding quadratic terms did not improve model fit over the linear models for any of the variables, we report below only the linear change pattern for each variable.

As in Figure 4.10, the temporal trend of behavior rate indicated that participants complied less with the new toothbrushing behavior over the weeks. The multilevel modeling suggested that the variance in behavior rate was mostly between-person ($ICC = 0.755$), and it confirmed the decreasing trend with a significant and negative effect of the time predictor ($B = -0.04$, $95\% CI = [-0.07, -0.01]$, $p = .004$).

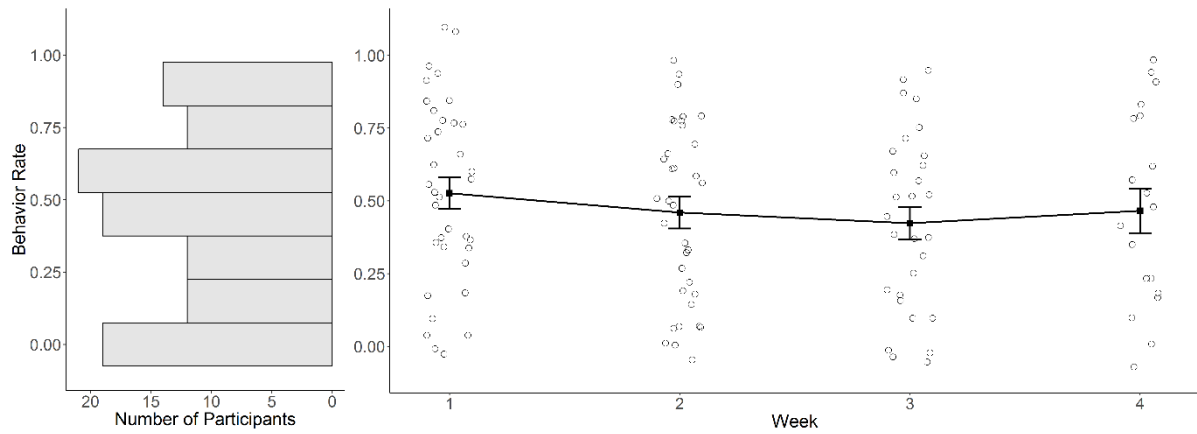


Figure 4.10 Histogram of behavior rate and its change over the weeks in Study 1 (errorbars represent one standard error).

Participants reported to hold generally positive instrumental attitude about brushing teeth for a second time (always above 5 on average), and there was a large individual difference between the participants in terms of their instrumental attitude ($ICC = 0.732$). This instrumental attitude also became slightly more positive over the weeks (Figure 4.11), and was confirmed by a significant and positive effect of time in the multilevel model ($B = 0.15$, $95\% CI = [0.07, 0.22]$, $p < .001$). In contrast, people's affective attitude about brushing twice a day was more neutral on average (around 4; Figure 4.12), and the individual difference in affective attitude was much smaller ($ICC = 0.483$). Affective attitude did not change significantly over time ($B = 0.11$, $95\% CI = [-0.02, 0.24]$, $p = .10$).




Figure 4.11 Histogram of instrumental attitude and its change over the weeks in Study 1 (errorbars represent one standard error).

4.3 Behavioral results




Figure 4.12 Histogram of affective attitude and its change over the weeks in Study 1 (errorbars represent one standard error).

Finally, self-reported habit strength increased clearly over the weeks (from 2 to 3 on average), and the change was the strongest among the modeled variables ($B = 0.23$, 95% CI = [0.13, 0.34], $p < .001$; also ICC = 0.607). The model only quantified a linear growth of habit strength, but from visual inspection there was a tendency of faster growth in the early compared to the later period (Figure 4.13). This resembled previous results of modeling habit growth in the real world (Gardner et al., 2014; Lally et al., 2010).




Figure 4.13 Histogram of self-reported habit strength and its change over the weeks in Study 1 (errorbars represent one standard error).

Effects of attitude and habit strength on behavior

The repeated-measured design allowed us to disaggregate of between-person and within-person effects (see Curran & Bauer. 2011). First, for each variable of interest, a multilevel null model was built to obtain its between-person component – the estimates of person means. Compared with a strategy of computing person means without modeling, the multilevel approach assumed that person means were drawn from a normal distribution, and it weighted participants with more observations more heavily than participants with fewer observations

to reduce variance due to small samples. Second, each variable's within-person component was extracted by subtracting the estimated person means from the raw data. Third, using the disaggregated components, between-person and within-person correlations between pairs of variables were computed to understand the individual effects of habit and attitude on behavior. Finally, the potential interaction effects of habit and attitude on behavior, both between-person and within-person, were estimated in multiple regression models.

Table 4.1 shows the between-person and within-person correlations for pairs of variables. At the between-person level (coefficients below the diagonal), all three behavior determinants – instrumental attitude, affective attitude, and habit strength correlated substantially with both behavior rates, while the correlations with the self-reported behavior rate were slightly higher than the ones with the objective behavior rate. However, these effects were completely inter-individual, because intra-individually, no correlations were found between the behavioral determinants and the behavior rates (coefficients above the diagonal). As for the relationships among the determinants, the two attitudinal variables correlated strongly as expected, while habit strength correlated more strongly with the affective dimension than with the instrumental dimension, both inter-individually and intra-individually.

Table 4.1 Between-person and within-person correlations between pairs of variables in Study 1.

	1	2	3	4	5
1. Behavior rate (sensor-based)	1	0.52 ^{***}	0.13	0.05	-0.08
2. Behavior rate (self-reported)	0.64 ^{***}	1	-0.19	-0.12	0.14
3. Instrumental attitude	0.51 ^{***}	0.57 ^{***}	1	0.33 ^{***}	0.38 ^{***}
4. Affective attitude	0.46 ^{***}	0.61 ^{***}	0.64 ^{***}	1	0.44 ^{***}
5. Habit strength	0.55 ^{***}	0.74 ^{***}	0.52 ^{***}	0.63 ^{***}	1

Note: Between-person and within-person correlations are shown below and above the diagonal. Significance indicators were adjusted for multiple test ($p < .05^*$, $p < .01^{**}$, $p < .001^{***}$).

When interaction effects were examined in multiple regression models (see Table 4.2 & 4.3), there was no indication of the moderating role of habit strength on the positive effects of either instrumental or affective attitude on behavior. Due to the small sample size in Study 1 (between-person $N = 36$), it is possible that the analyses were not powerful enough to detect the interaction effect. For the same reason, although all behavioral determinants correlated strongly with behavior individually at the between-person level, most coefficients estimated from the multiple regression models were not statistically significant.

4.3 Behavioral results

Table 4.2 Regression models of attitude, habit strength, and their interaction on objective behavior rate (Study 1).

	Between-person		Within-person	
	<i>B</i>	95% CI	<i>B</i>	95% CI
Instrumental attitude (IA)	0.11 ⁺	[-0.00, 0.22]	0.03	[-0.03, 0.09]
Affective attitude (AA)	-0.01	[-0.12, 0.10]	0.01	[-0.03, 0.04]
Habit strength (HS)	0.07	[-0.02, 0.16]	-0.03	[-0.07, 0.01]
IA <i>By</i> HS	0.07	[-0.03, 0.16]	-0.02	[-0.08, 0.04]
AA <i>By</i> HS	-0.06	[-0.15, 0.04]	0.02	[-0.03, 0.06]
Adjusted <i>R</i> ²	0.320		-0.016	

Note: Significance levels were indicated as $p < .10^+$, $p < .05^*$, $p < .01^{**}$, $p < .001^{***}$.

Table 4.3 Regression models of attitude, habit strength, and their interaction on self-reported behavior rate (Study 1).

	Between-person		Within-person	
	<i>B</i>	95% CI	<i>B</i>	95% CI
Instrumental attitude (IA)	0.12 ⁺	[-0.00, 0.23]	-0.05 ⁺	[-0.10, 0.002]
Affective attitude (AA)	0.00	[-0.12, 0.12]	-0.01	[-0.04, 0.02]
Habit strength (HS)	0.15 ^{**}	[0.05, 0.25]	0.04 ⁺	[-0.002, 0.07]
IA <i>By</i> HS	0.08 ⁺	[-0.01, 0.18]	0.01	[-0.07, 0.08]
AA <i>By</i> HS	-0.10 ⁺	[-0.20, 0.00]	0.04 ⁺	[-0.01, 0.09]
Adjusted <i>R</i> ²	0.589		0.069	

Note: Significance levels were indicated as $p < .10^+$, $p < .05^*$, $p < .01^{**}$, $p < .001^{***}$.

Effects of behavior on the changes in attitude and habit strength

The reciprocal effects of behavior on the changes in attitude and habit strength were examined by regressing the difference scores¹⁷ of the dependent variables between two successive

¹⁷ It has been highly debated in the literature on the appropriate method for analyzing change (see e.g., Lord's paradox, Lord, 1967). For randomized experiments, it is clear that covariate methods (ANCOVA or regression with baseline as covariate) are preferred to the use of difference or change scores (e.g., Maris, 1998; van Breukelen, 2006). However, for experiments with existing samples or correlational studies, there is some consensus that the difference-score method is less biased and less susceptible to spurious correlations than covariate methods (e.g., Allison, 1990; Glymour, Weuve, Berkman, Kawachi, & Robins, 2005; Maris, 1998). A simple simulation with the habit formation model discussed in this Chapter showed that when there was no effect on change, covariate methods erroneously produced significant correlations while difference-score method did not.

measurement points to either objective or self-reported weekly behavior rate. Because the variances in the different scores were completely within-person, simple linear regressions were used instead of multilevel models. Results indicated no dependence of attitude change on behavior rate in the preceding week – neither for the instrumental (objective: $B = 0.45$, 95% CI = [-0.16, 1.06], $p = .15$; subjective: $B = 0.29$, 95% CI = [-0.14, 0.73], $p = .19$) nor the affective dimension (objective: $B = 0.39$, 95% CI = [-0.66, 1.44], $p = .46$; subjective: $B = 0.86$, 95% CI = [-0.08, 1.79], $p = .08$), but a positive relationship was found between the change in habit strength and sensor-based behavior rate ($B = 1.40$, 95% CI = [0.60, 2.20], $p < .001$, adjusted $R^2 = 0.14$) and marginally on self-reported behavior rate ($B = 0.65$, 95% CI = [-0.03, 1.32], $p = .06$, adjusted $R^2 = 0.03$). Thus, participants who performed the new brushing behaviors more often in a week seemed to show more habit growth afterwards.

4.3.2 Results of Study 2

Study 2 had a very high survey response rate of 98%, and most participants did not miss any of the 5 surveys (see Figure 4.14a). As with Study 1, the proportion of days with missing sensor data averaged across participants was 9%. Two participants did not register any sensor data due to technical faults, so only self-reported behavior rate was used for them. The self-report and objective measures of behavior rate also correlated strongly ($r(213) = 0.65$; see Figure 4.14c), although this correlation was substantially lower than the 0.88 in Study 1. There were several differences that might account for the lower correlation. First, the self-report measure concerned not only with the occurrence of brushing twice, but also with the duration that was supposed to be longer than 2 minutes for each brushing session. Participants might have had more difficulty reporting the number of occurrences with a specific duration than simply reporting the number of occurrences. If participants did take duration into account, then some data points in the top-left corner of Figure 4.14c could be explained by the fact that although they brushed twice (high objective behavior rate), they self-reported very low rates because many brushing sessions were less than 2 minutes. Second, as mentioned in the method section, the objective measure was per period, which differed in length (i.e., between 4 and 10 days), while the subjective measure always referred to the previous week. This temporal mismatch between the two methods could have caused the lower correlation. Third, the older and less-educated sample of participants might have had more trouble with self-reporting. For example, for participant 27, although sensor data clearly suggested this person always brushed at most once a day, the self-reported behavior rates for the 4 surveys were 6, 0, 7, 0, showing very unlikely fluctuations in behavior.

4.3 Behavioral results

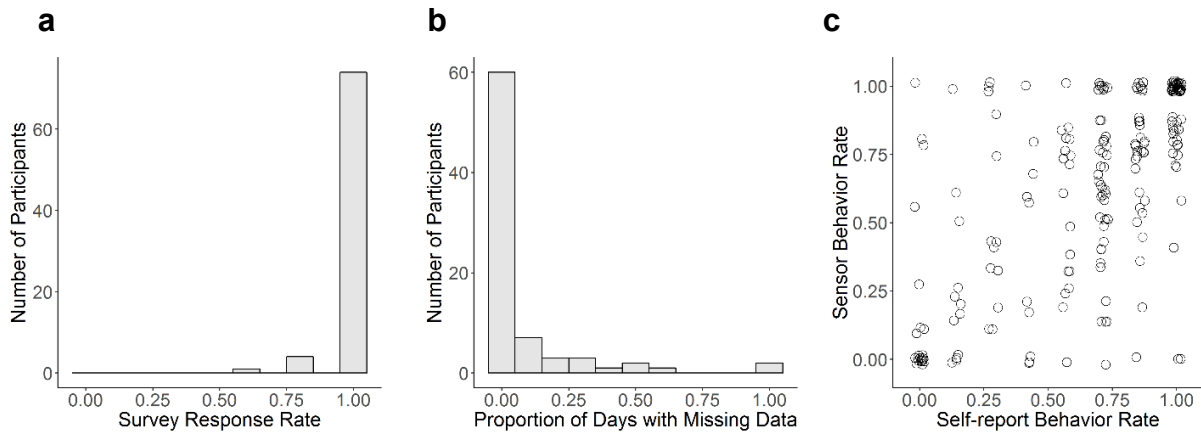


Figure 4.14 Data quality check for Study 2: (a) Histogram of survey response rate; (b) Histogram of proportion of days with missing sensor data; (c) Correlation between self-report and objective behavior rate.

Description and modeling of changes

As in Figure 4.15, the distributions and temporal changes of sensor-measured and self-reported behavior rate were pretty similar, except for two small differences. First, based on multilevel null models, the variance in sensor-measured behavior rate was more heavily attributed to individual difference ($ICC = 0.720$) than the variance in self-reported behavior rate ($ICC = 0.610$). Second, while the objective measure showed a small decline in behavior rate ($B = -0.03$, $95\% CI = [-0.07, -0.01]$, $p = .021$), the self-reported behavior rate stayed stable over time ($B = 0.02$, $95\% CI = [-0.01, 0.05]$, $p = .22$).

Compared to Study 1, instrumental attitude was also clearly more positive than affective attitude on average (Figure 4.16 & 4.17), but in Study 2 there were relatively more between-person variations in affective attitude ($ICC = 0.688$) than in instrumental attitude ($ICC = 0.550$). Over the 5 measurement points, both attitudinal dimensions showed small peaks at measurement 3, quantified by the positive effects of a linear time predictor (instrumental: $B = 0.58$, $95\% CI = [0.41, 0.75]$, $p < .001$; affective: $B = 0.57$, $95\% CI = [0.33, 0.80]$, $p < .001$) and the negative effects of a cubic time predictor (instrumental: $B = -0.08$, $95\% CI = [-0.11, -0.05]$, $p < .001$; affective: $B = -0.09$, $95\% CI = [-0.13, -0.05]$, $p < .001$). This abrupt increase in attitude was probably due to the lab intervention session that happened between measurement 2 and 3 (these two measurements were on the same day, before and after the intervention).




Figure 4.15 Histogram of behavior rate and its change over the weeks in Study 2 (errorbars represent one standard error).




Figure 4.16 Histogram of instrumental attitude and its change over the weeks in Study 2 (errorbars represent one standard error).




Figure 4.17 Histogram of affective attitude and its change over the weeks in Study 2 (errorbars represent one standard error).

4.3 Behavioral results




Figure 4.18 Histogram of self-reported habit strength and its change over the weeks in Study 2 (errorbars represent one standard error).

The most evident difference between Study 1 and Study 2 was that in Study 2 participants had much higher habit strengths overall, and especially even at the beginning of the study (Figure 4.18). Probably due to the lenient inclusion criterion, many participants were not learning a completely new habit of brushing twice a day, although they were aware of the recommended duration during the lab intervention session. As a result, habit strength differed mostly between-person ($ICC = 0.769$) and showed a much smaller growth over the weeks ($B = 0.08$, $95\% CI = [0.0001, 0.15]$, $p = .049$).

Effects of attitude and habit strength on behavior

As shown in Table 4.4, the results of the between-person correlations between pairs of variables largely coincided with the results in Study 1. Participants who had more positive instrumental or affective attitudes, or had stronger habits of brushing teeth twice a day, indeed brushed twice a day more frequently on average. The correlation between habit strength and brushing behavior was particularly strong when the behavior was measured by self-reports. Inter-individually, habit strength again correlated more strongly with affective attitude than with instrumental attitude.

Intra-individually, however, there were two differences compared to Study 1. First, instrumental attitude also had a slight within-person effect on self-reported behavior rate, meaning that if a person happened to consider brushing twice more beneficial or healthy, they would also be more likely to comply with the behavior. Second, inconsistent with the results of Study 1, habit strength did not correlate with instrumental attitude, nor with affective attitude within-person.

Table 4.4 Between-person and within-person correlations between pairs of variables in Study 2.

	1	2	3	4	5
1. Behavior rate (sensor-based)	1	0.40 ^{***}	0.02	0.04	-0.16
2. Behavior rate (self-reported)	0.73 ^{***}	1	0.24 ^{***}	0.17	-0.01
3. Instrumental attitude	0.42 ^{***}	0.35 ^{***}	1	0.41 ^{***}	0.09
4. Affective attitude	0.47 ^{***}	0.63 ^{***}	0.34 ^{***}	1	0.15
5. Habit strength	0.45 ^{***}	0.73 ^{***}	0.18	0.55 ^{***}	1

Note: Between-person and within-person correlations are shown below and above the diagonal. Significance indicators were adjusted for multiple test ($p < .05^*$, $p < .01^{**}$, $p < .001^{***}$).

With a sample size twice as large, Study 2 was certainly more informative about the interaction effect between habit and attitude in determining toothbrushing behavior (see Table 4.5 & 4.6). There was some suggestive evidence that inter-individually, there was a negative interaction effect between affective attitude and habit strength – for those who had stronger habits of brushing teeth twice, the positive influence of affective attitude on behavior rate was attenuated. There was also a trend at the within-person level that habit strength could have similarly moderated the effect of affective attitude on sensor-measured behavior, though the marginal effect would require replications with more observations. Besides the interaction effect, when habit strength was a predictor alongside with attitude at the within-person level, it had unexpected negative effects on sensor-measured behavior rate. Finally, as with the results of Study 1, it was clear that when the three behavioral determinants were considered at the same time, the effect of affective attitude tended to diminish, even though it correlated strongly with behavior by itself.

4.3 Behavioral results

Table 4.5 Regression models of attitude, habit strength, and their interaction on objective behavior rate (Study 2).

	Between-person		Within-person	
	<i>B</i>	95% CI	<i>B</i>	95% CI
Instrumental attitude (IA)	0.23**	[0.09, 0.37]	0.0002	[-0.04, 0.04]
Affective attitude (AA)	0.03	[-0.05, 0.11]	0.01	[-0.02, 0.05]
Habit strength (HS)	0.07**	[0.02, 0.11]	-0.04*	[-0.07, -0.01]
IA <i>By</i> HS	0.08	[-0.04, 0.19]	-0.02	[-0.08, 0.05]
AA <i>By</i> HS	-0.05 ⁺	[-0.11, 0.001]	-0.05 ⁺	[-0.10, 0.01]
Adjusted <i>R</i> ²	0.342		0.029	

Note: Significance levels were indicated as $p < .10^+$, $p < .05^*$, $p < .01^{**}$, $p < .001^{***}$.

Table 4.6 Regression models of attitude, habit strength, and their interaction on self-reported behavior rate (Study 2).

	Between-person		Within-person	
	<i>B</i>	95% CI	<i>B</i>	95% CI
Instrumental attitude (IA)	0.10*	[0.01, 0.20]	0.07**	[0.02, 0.12]
Affective attitude (AA)	0.05 ⁺	[-0.06, 0.16]	0.03	[-0.01, 0.07]
Habit strength (HS)	0.08**	[0.05, 0.11]	-0.01	[-0.05, 0.03]
IA <i>By</i> HS	0.0005	[-0.08, 0.08]	-0.002	[-0.07, 0.07]
AA <i>By</i> HS	-0.05*	[-0.09, -0.01]	-0.003	[-0.06, 0.06]
Adjusted <i>R</i> ²	0.566		0.045	

Note: Significance levels were indicated as $p < .1^+$, $p < .05^*$, $p < .01^{**}$, $p < .001^{***}$.

Effects of behavior on the changes in attitude and habit strength

For the reciprocal influences, results again indicated no dependence of attitude change on behavior rate in the preceding week – neither for the instrumental (objective: $B = 0.09$, 95% CI = [-0.13, 0.31], $p = .41$; subjective: $B = 0.07$, 95% CI = [-0.15, 0.28], $p = .54$) nor the affective dimension (objective: $B = 0.11$, 95% CI = [-0.23, 0.45], $p = .53$; subjective: $B = 0.12$, 95% CI = [-0.23, 0.46], $p = .50$). In contrast with Study 1, Study 2 with a larger sample did not reveal any correlation between behavior rate in the previous week and change in habit strength over the week (objective: $B = 0.07$, 95% CI = [-0.31, 0.44], $p = .71$; subjective: $B = 0.21$, 95% CI = [-0.17, 0.58], $p = .27$).

Because participants' habit strengths were much higher at baseline in Study 2 than in Study 1, we checked the same effects of behavior rate on change in habit strength for a sub-group of

participants who had weaker habit strengths than the mid-point of the scale (3.5). Figure 4.19 shows that for the sub-group, the distribution and temporal change of habit strength were very similar to those in Study 1, and there was some fluctuation of habit strength over the weeks based the significant coefficients of linear and cubic time predictors in multi-level growth curve modeling (linear: $B = 1.22$, 95% CI = [0.59, 1.85], $p < .001$; cubic: $B = -0.18$, 95% CI = [-0.31, -0.06], $p = .004$). For this sub-group, again no significant correlation was found between behavior rate and change in habit strength, even though the estimates were clearly larger than those in the whole sample (objective: $B = 0.33$, 95% CI = [-0.21, 0.87], $p = .23$; subjective: $B = 0.47$, 95% CI = [-0.07, 1.01], $p = .09$).




Figure 4.19 Histogram of self-reported habit strength and its change over the weeks for participants with weak baseline habits in Study 2 (errorbars represent one standard error).

4.4 Predicting toothbrushing behavior with theory-based models

4.4.1 Computational models of cognitive variables

We computed two cognitive variables – habit strength (HS) and memory accessibility (Acc) – to examine whether they can improve prediction for toothbrushing behavior on the next day (1-step forecasting). We used the equation from Klein, Mogles, Treur, & van Wissen (2011; henceforth as Klein’s equation) for modeling habit strength, which was similar to the equation from Miller et al. (2019) used in Chapter 3. In Chapter 3, the focus was on decision-making rather than learning, so Miller’s equation was preferred for its simplicity and its desirable numeric property that habit strength was bounded between 0 and 1 regardless of its free parameters. However, while Miller’s equation was controlled by one parameter of learning rate, Klein’s equation with two free parameters allowed different rates in habit growth and decay, which was more consistent with the empirical findings about human habit formation (e.g., Lally et al., 2010). The equation with a habit decay parameter (HDP) and a habit gain parameter (HGP) is the following:

$$HS_{t+1} = HS_t - HDP \times HS_t + HGP \times (1 - HS_t) \times Beh_t \times Cue_t$$

The equation models habit strength to decay naturally but to grow with decreasing speed when the corresponding behavior (Beh) is performed under the corresponding cue. Habit strength at the start of the study (HS_0) was initiated by scaling the self-reported behavioral automaticity to the interval between 0 and 1, and later habit strengths were computed by the equation. Because we did not measure whether the same cues were always encountered when participants brushing their teeth (e.g., whether the locations were always their bathrooms at home), the variable cue was assumed to always take the value 1. We considered it as a reasonable simplification for a first evaluation of our approach.

Memory accessibility of options was modeled using the equation in Tobias (2009). Accessibility decays naturally as a natural memory process, but can be enhanced by behavior executions and external reminders. The equation controlled by three free parameters – accessibility decay parameter (ADP), accessibility gain parameter with behavior execution (AGP_{beh}), and accessibility gain parameter with reminder (AGP_{rem}), is as the following:

$$Acc_{t+1} = Acc_t - ADP \times Acc_t + (AGP_{beh} \times Beh_t + AGP_{rem} \times Rem_t) \times (1 - Acc_t)$$

Accessibility of the target brushing behavior was assumed to be 1 (maximum) at the beginning of the studies, as participants were either just told to perform the behavior (Study 1) or were highly aware of the optimal dental routine (Study 2) in face-to-face meetings. Accessibility on the date of the lab session in Study 2 was also set to 1 for all participants, as they were told how to brush optimally. For simplicity, both receiving reminders (Study 1) and receiving notifications or e-mails for answering surveys (Study 1 & 2) were considered as the same type of reminder, modulated by a single parameter AGP_{rem} .

4.4.2 Model comparison method

We set to evaluate the predictive value of computed habit strength and memory accessibility by comparing the performance of statistical models with and without these variables. The target for prediction was the brushing behavior on the next day, with the occurrence of brushing as the *negative cases* and the absence of brushing as the *positive cases*. They were coded in this way because in applications a more important goal would be to detect the positive cases, i.e., the days on which the brushing behavior was likely to be omitted. The theory-based computational approach would be considered valuable if it led to models that performed better than models based simply on past behavior and on weekly self-reported variables. Specifically, models with 4 different feature sets were compared:

- *Survey model*: The primary features in the survey model were the variables measured by weekly surveys, including *instrumental attitude*, *affective attitude*, and *self-reported behavioral automaticity*. In addition, the *occurrence of lab sessions* (including the introduction meeting in Study 1) and the *occurrence of reminders* (including notifications and e-mails for surveys) were also included as features.

- *Past-behavior model*: The primary feature in this model was the *past behavior rate* until the day of the last observation. For example, if the brushing behavior on the 11th day was to be predicted, the brushing rate in the last 10 days (e.g., 0.8) would be the value for this variable. For the first day, past behavior rate was set to 0 in Study 1, as participants self-reported to rarely brush twice. In Study 2, the self-reported behavior rates in the previous week were used for the initial values. Again, the *occurrence of lab sessions* and the *occurrence of reminders* were also included as features.
- *Theory-based model*: This was the model of our interest that includes only computed *habit strength* and *accessibility* as features.
- *Combined model*: The combined model included features in both the past-behavior model and the theory-based model. It was used mainly to evaluate whether combining past behavior and computed cognitive features could further boost prediction performance.

For each model type, three common statistical learning algorithms were used, including logistic regression, support vector machine, and random forest. In total, this resulted in 12 models (4 model types \times 3 algorithms) to be trained and tested.

Two different approaches were used to compare model performance. First, a two-level hierarchical k -fold cross-validation procedure was used on each of the two data sets separately (see Figure 4.20). For each data set, all observations were divided into k non-overlapping groups (with the restriction that one participant's data were always in only one group), so that 1 group was reserved for model testing, and the remaining $k-1$ groups were used for training in each round (the outer loop). Because tuning was needed for both the free parameters in the equations of HS and Acc and the hyperparameters for support vector machine and random forest, the training set in each round was further divided, with 1 group reserved as the test set for parameter tuning and the remaining $k-2$ groups as the training set for parameter tuning (the inner loop). For each free parameter in the theory-based equations, a 3000-step random search was used, and in each step a random value was drawn from a uniform distribution between 0 and 1. For the hyperparameters, grid-search was used to swipe the parameter space as defined in Table 4.6. These parameter values were optimized to obtain the best overall prediction performance in the inner cross-validation loop, indicated by area under curve (AUC) in receiver operating characteristic (ROC) curves for logistic regression and random forest and by Matthews correlation coefficient (MCC) for support vector machine. Due to the sample size difference between the two studies, 9 folds were used for Study 1 (4 participants in each group) and 5 folds were used for Study 2 (15 participants in each group), in order to have sufficient data for training.

4.4 Predicting toothbrushing behavior with theory-based models

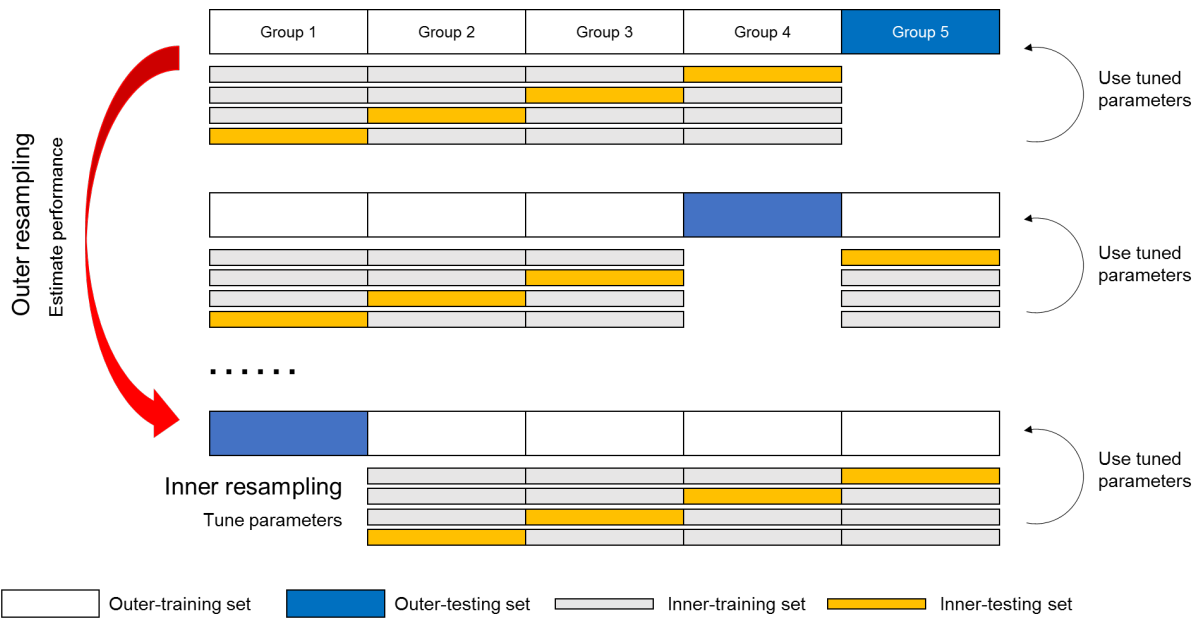


Figure 4.20 An illustration of the nested cross-validation procedure used (it shows the 5-fold scenario for Study 2, but the same idea applies to Study 1).

Table 4.6 Range of values used in hyperparameter tuning.

Support vector machine	
C (regularization parameter)	0.1, 1, 10, 100, 1000
γ (width parameter of the Gaussian kernel)	0.0001, 0.001, 0.01, 0.1, 1, 10
Random forest	
<i>ntree</i> (number of trees)	500
<i>nodesize</i> (minimum observations in the terminal nodes)	1, 4, 16, 64, 256
<i>mtry</i> (number of features used for node split)	1, 2, ..., n_{feature}

Since we had two similar data sets, in a second approach, we evaluated the ability of each model type to predict new data. This approach was used to measure the generalizability of the models, in particular the theory-based model. Specifically, one of the two data sets was used to train the models, and the resultant models were used to predict the observations in the other data set. When parameter tuning was required, a k -fold cross-validation was used on the whole training data set, with the same search methods indicated above. Again, 9-fold or 5-fold cross-validation was used when Study 1 or Study 2 was used as the training data set respectively.

For model comparison, we primarily focused on AUC and MCC, since they provided a more balanced evaluation of prediction performance for both positive and negative cases. Overall prediction accuracy, F1-score, and other case-sensitive performance measures (e.g., true and

false positive rates) were also used for comparison. All analyses were performed in R statistical programming environment (version 3.3.3), with the help of the *mlr* (machine-learning R, version 2.1.3) package (Bischl et al., 2016).

4.4.3 Model comparison results

Cross-validation within each data set

Study 1 included 711 non-missing observations for the prediction task, with 376 positive cases (non-brushing) and 335 negative cases (brushing). Thus, the baseline prediction accuracy was 53% if a null model predicted all positive cases. Figure 4.21 shows the testing ROC curves of different models, and Table 4.7 compares additional testing performance measures of the models (aggregated over cross-validation iterations). Overall, various performance measures indicated that the theory-based models were better than the survey models, but were slightly worse than the past-behavior models. It was also clear that combining the features of the theory-based and past-behavior models did not improve performance any further. All models were able to achieve prediction accuracy substantially higher than the baseline, ranging between 63% and 70%. If detecting positive cases was the main interests, it seems that the theory-based models were more reluctant in predicting positive cases, which was also reflected in the relatively high precision but low true positive rate (TPR). In terms of learning algorithms, their results were largely the same, although more complex algorithms (SVM and random forest) showed more decline of performance from training to testing set, suggesting some overfitting in training. Results of Study 1 did not show any support for the benefits of using the theory-based computational approach.




Figure 4.21 Model comparison results of Study 1 based on ROC curves for different models and algorithms.

4.4 Predicting toothbrushing behavior with theory-based models

Table 4.7 Comparison of model performances in predicting testing data (Study 1).

		AUC	MCC	Acc	TPR	FPR	Precision	F1-score	NPV
<i>Logistic regression</i>	Survey	0.660	0.261	0.632	0.644	0.382	0.654	0.649	0.607
	PB	0.758	0.394	0.698	0.702	0.307	0.719	0.711	0.674
	Theory	0.737	0.337	0.668	0.660	0.322	0.697	0.678	0.639
	Combined	0.750	0.367	0.684	0.681	0.313	0.709	0.695	0.657
<i>SVM</i>	Survey	0.655	0.277	0.640	0.662	0.385	0.659	0.660	0.619
	PB	0.753	0.338	0.671	0.721	0.385	0.678	0.698	0.662
	Theory	0.698	0.311	0.648	0.564	0.257	0.711	0.629	0.603
	Combined	0.740	0.391	0.689	0.614	0.227	0.752	0.676	0.641
<i>Random forest</i>	Survey	0.663	0.309	0.655	0.673	0.364	0.675	0.674	0.634
	PB	0.700	0.313	0.657	0.662	0.349	0.680	0.671	0.632
	Theory	0.704	0.274	0.633	0.574	0.301	0.681	0.623	0.594
	Combined	0.707	0.312	0.655	0.638	0.325	0.688	0.662	0.624

Note: Survey = survey model; PB = past-behavior model; Theory = theory-based model; SVM = support vector machine; Acc = accuracy; TPR = true positive rate; FPR = false positive rate; NPV = negative prediction value; MMC = Matthews correlation coefficient.

Study 2 included 1508 non-missing observations for the prediction task, with 557 positive cases (non-brushing) and 951 negative cases (brushing). Thus, the data were less balanced and the baseline prediction accuracy was 63% if a null model predicted all negative cases. Figure 4.22 shows the testing ROC curves of different models, and Table 4.8 compares additional testing performance measures of the models in Study 2. In contrast with Study 1, the theory-based models performed much better than the survey models, and also slightly better than the past-behavior models. The models with combined features was arguably the best, although the improvements over the theory-based models were very small. Since the data were more unbalanced (more negative cases due to a higher brushing rate) compared with Study 1, all models were able to predict more accurately, with average accuracy between 67% and 78%. When examining the case-sensitive measures, the advantages of the theory-based models seemed to be mainly driven by more accurate predictions for the negative cases. Thus, consistent with Study 1, the theory-based models were able to predict some more negative cases (days on which participants would not forget to brush twice) than the past-behavior models, as the latter produced more false positives. Again, differences between the three algorithms used were very small, but again logistic regression was less prone to overfitting in our prediction task.




Figure 4.22 Model comparison results of Study 2 based on ROC curves for different models and algorithms.

Table 4.8 Comparison of model performances in predicting testing data (Study 2).

		AUC	MCC	Acc	TPR	FPR	Precision	F1-score	NPV
<i>Logistic regression</i>	Survey	0.676	0.278	0.684	0.372	0.132	0.622	0.465	0.702
	PB	0.792	0.447	0.752	0.521	0.113	0.730	0.608	0.760
	Theory	0.815	0.504	0.776	0.533	0.082	0.792	0.637	0.771
	Combined	0.816	0.514	0.781	0.555	0.087	0.788	0.651	0.778
<i>SVM</i>	Survey	0.678	0.282	0.687	0.354	0.118	0.638	0.455	0.700
	PB	0.743	0.410	0.736	0.506	0.129	0.696	0.586	0.751
	Theory	0.757	0.491	0.771	0.531	0.089	0.777	0.631	0.768
	Combined	0.778	0.488	0.769	0.528	0.089	0.776	0.628	0.767
<i>Random forest</i>	Survey	0.661	0.217	0.660	0.339	0.152	0.566	0.424	0.687
	PB	0.772	0.441	0.746	0.585	0.160	0.682	0.630	0.776
	Theory	0.791	0.484	0.767	0.555	0.108	0.750	0.638	0.774
	Combined	0.804	0.498	0.773	0.576	0.111	0.752	0.652	0.782

Note: Survey = survey model; PB = past-behavior model; Theory = theory-based model; SVM = support vector machine; Acc = accuracy; TPR = true positive rate; FPR = false positive rate; NPV = negative prediction value; MMC = Matthews correlation coefficient.

Predicting new data

Results of the models' abilities in predicting new data are summarized in Figure 4.23 and Table 4.9. As would be expected, performances of all models dropped significantly when new data were predicted, leaving an accuracy around 65% at best for Study 1 and around 70% at

4.4 Predicting toothbrushing behavior with theory-based models

best for Study 2. Encouragingly, although the theory-based models included a few free parameters to be estimated from the training data before features could be computed, their performance drops were no worse than the other model types. Consistent with the results in the last section, the theory-based models and past-behavior models performed substantially better than the survey models, while the differences between them were quite small. When category-sensitive measures were considered, it was again shown that the past-behavior models were better at predicting positive cases, while the theory-based models were better at predicting negative cases.

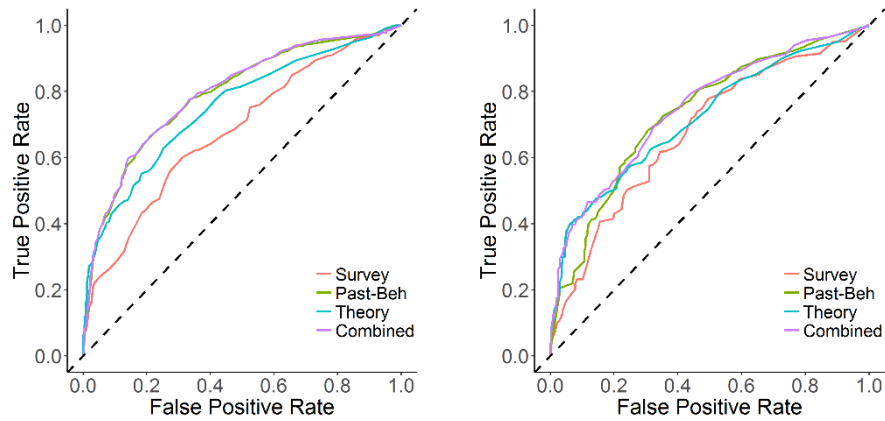


Figure 4.23 Model comparison results in terms of predicting new data, based on ROC curves of different models (Left panel: predicting Study 2's data using models trained on Study 1's data; Right panel: predicting Study 1's data using models trained on Study 2's data).

Table 4.9 Comparison of model performances in predicting new data.

		AUC	MCC	Acc	TPR	FPR	Precision	F1-score	NPV
<i>Predicting data set 2</i>	Survey	0.680	0.195	0.570	0.715	0.515	0.448	0.551	0.744
	PB	0.793	0.421	0.710	0.749	0.313	0.583	0.656	0.823
	Theory	0.753	0.362	0.702	0.601	0.239	0.596	0.599	0.765
	Combined	0.795	0.422	0.711	0.745	0.309	0.585	0.656	0.822
<i>Predicting data set 1</i>	Survey	0.677	0.247	0.605	0.431	0.200	0.707	0.536	0.556
	PB	0.733	0.311	0.637	0.473	0.179	0.748	0.580	0.581
	Theory	0.718	0.365	0.648	0.423	0.099	0.828	0.560	0.582
	Combined	0.749	0.364	0.647	0.418	0.096	0.831	0.556	0.580

Note: Survey = survey model; PB = past-behavior model; Theory = theory-based model; Acc = accuracy; TPR = true positive rate; FPR = false positive rate; NPV = negative prediction value; MMC = Matthews correlation coefficient.

Lastly, for theoretical interests, we examined the optimal parameter values for the free parameters in the theory-based equations of habit strength and accessibility. For parameters governing the dynamics of habit strength, optimal ranges of parameter values could be found, and the results were similar regardless of the data set used (see Figure 4.24). To achieve best performance based on AUC, the optimal value for the habit decay parameter (*HDP*) was in the range of 0.15 and 0.2, while the optimal value for the habit gain parameter (*HGP*) was in the range of 0.1 and 0.2.

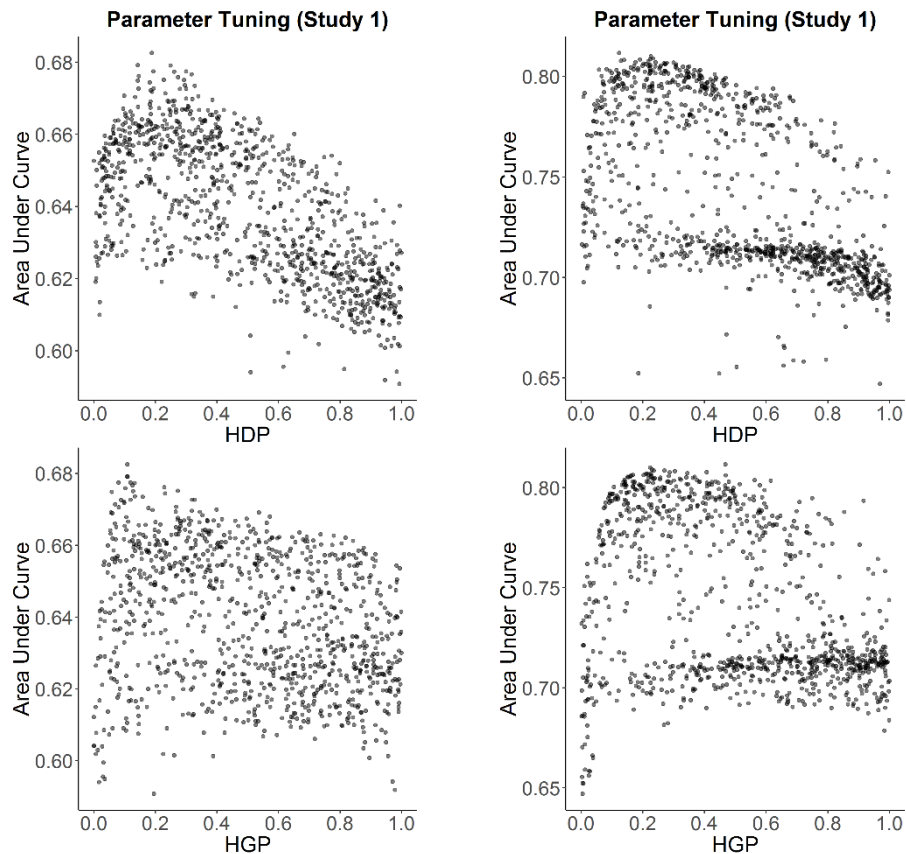


Figure 4.24 Tuning results for parameter HDP and HGP in the computational model of habit strength, shown as the relationship between parameter values (x-axis) and model performance (area under curve, y-axis).

In contrast, for parameters that determine the dynamics of accessibility, there seemed to be no relationships between their values and model prediction performance (see Figure 4.25). If one examined the individual features in the theory-based models, the feature habit strength contributed to most of their predictive powers, while the feature accessibility did not contribute as much.

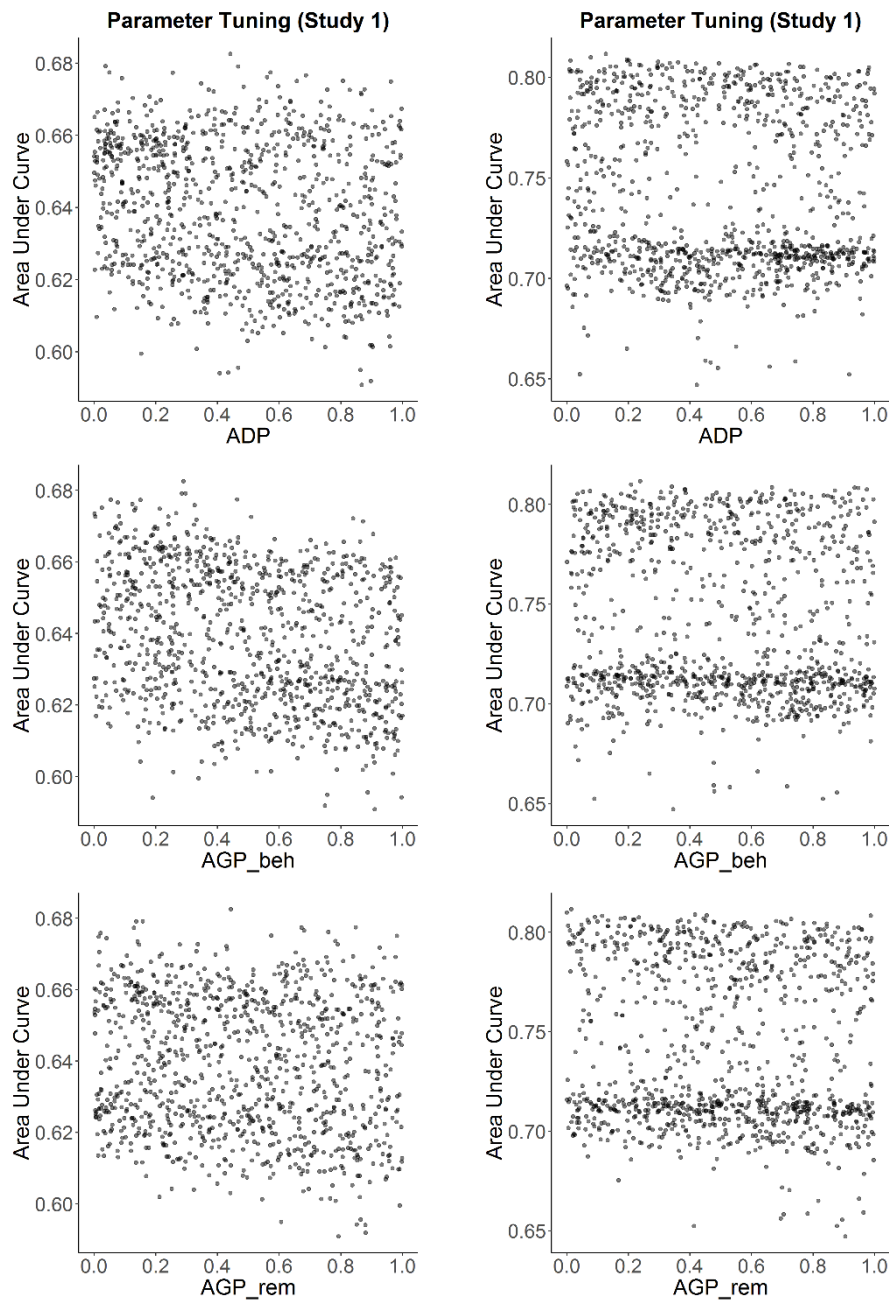


Figure 4.25 Tuning results for parameter HDP and HGP in the computational model of habit strength, shown as the relationship between parameter values (x-axis) and model performance (area under curve, y-axis).

4.5 General discussion

In this chapter, we investigated habit formation in the real-world through two field intervention studies of promoting optimal dental routines. The theoretical construct of habit strength was operationalized both as self-reported behavioral automaticity and as a computed variable

based on its theoretical relationship with behavioral repetition. Adding to the growing literature on habit formation in health-related behaviors, we examined the reciprocal relationships between self-reported habit strength, attitude, and objectively-measured toothbrushing behavior. In the second half of the chapter, we ventured a bit further to explore whether a computational approach to habit formation could assist behavior prediction in intervention trials.

4.5.1 The reciprocal relationships between habit strength, attitude, and brushing behavior

First, our results suggest a clear and moderately strong positive relationship between habit strength and brushing behavior, in addition to the positive effects of instrumental and affective attitude on behavior. This finding is consistent with a series of empirical findings on the role of habit strength on behavior execution in health-related behaviors (Gardner, 2015). As we measured brushing behavior and the behavioral determinants multiple times, we could quantify the strengths of the inter-individual and intra-individual effects separately. It is somewhat surprising that both the effects of attitude and habit strength on brushing behavior are found to be only between-person but not within-person. This finding, if generalizable, would imply that people with an overall more positive attitude and stronger habit also engage in the corresponding behavior more frequently, but for one specific person a positive change in attitude or habit strength is not associated with increased behavior frequency. Such an implication also contradicts with psychological theories that treat habit formation as a within-person process. However, we consider this interpretation as premature, because the lack of within-person effects may be attributed to measurement errors. Unlike the more reliable person means used for computing between-person effects, the within-person scores (deviations from person means) inherently include a component of random variations from one measurement point to another. Given that the variations in these within-person scores are already smaller than the individual differences (ICC between 50-75%), the inherent measurement errors may mask any small but meaningful within-person effects. More longitudinal studies with better controls for measurement error are needed to study habit formation within-person.

Second, our results provide some support for the theoretical idea that strong habit attenuates the influences of attitude on behavior (e.g., Triandis, 1977; Verplanken et al., 1994), and the effect also applies to toothbrushing behavior. This moderation effect of habit strength on the relationship between attitude and behavior was found in Study 2, but not in Study 1. Given that Study 2 had a sample size twice as large and that habit strength in Study 2 was much more varied between-person than Study 1, the significant finding in Study 2 should be weighted more, yet more replications are clearly needed to verify the effect. Even if the moderation effect exists, it is smaller and perhaps less robust than the main effects of habit

strength and attitude. This may explain why some null results have been reported in the literature (Murtagh et al., 2012; Norman, 2011; Gardner et al., 2012), although significant effects were reported more frequently (see Gardner, 2015).

Third, we found mixed results regarding the reciprocal effect of behavior execution on the growth of habit strength. A strong positive correlation was found between weekly behavior rate of toothbrushing and the change in habit strength measured before and after the week in Study 1, but the effect was estimated to be much smaller and non-significant in Study 2. According to the computational models of habit formation (e.g., Klein, Mogles, Treur, & van Wissen, 2011; Miller et al., 2019; Psarra, 2016; Tobias, 2009) and its empirically observed dynamics (Lally et al., 2010), the positive influence of behavior execution on habit strength should gradually decrease as habit strength increases. This may explain the negative result in Study 2, since the average self-reported habit strength in the second sample was already quite high.

While more research is needed to examine the effect, we did provide a more rigorous test when compared with previous studies (de Bruijn et al., 2014; Fleig et al., 2003; Gardner & Lally, 2013; Wiedemann et al., 2013) by improving two aspects of study design: (1) Change in habit strength was perfectly matched to the preceding behavior frequency in our repeated-measure design; (2) Difference score was used as an unbiased way to control for baseline habit strength. For baseline adjustment in our studies, the use of difference score is preferable to the use of analysis of covariance (ANCOVA) or regression with baseline habit strength as covariate (see Verplanken & Melkevik, 2008), because the former is less biased than the latter in correlational designs (see Allison, 1990; Glymour et al., 2005; Maris, 1998). The superiority of difference score was further confirmed in a simulation study using the same habit formation model presented in 4.6.1. Specifically, even if a behavioral determinant (e.g., attitude) is not influenced by behavior execution but changes randomly across measurement points, using its baseline as a covariate would produce a strong but spurious correlation between behavior frequency and the behavioral determinant, as long as there is a causal link from the behavioral determinant to behavior execution. Such strong correlations were indeed evident in our data when using the covariate-method, even though there is no strong theory on the change of attitude based merely on behavior repetition. The difference-score method was shown to be immune to this problem through the simulation. It should be noted that although the difference-score method is less biased, it may lack the power to detect small effects due to measurement error.

Some of the controversies in our findings and in the empirical literature overall (e.g., lack of within-person effects, the mixed results regarding the moderation effect of habit) might eventually be resolved if researchers start to distinguish process theories of habit formation on one hand and statistical models of empirical data on the other. While statistical models describe mostly linear relationships at the weekly level (or other aggregated temporal level), the

actual causal influences of habit strength and other determinants on behavior and vice versa are likely to be non-linear and operate at the level of individual decisions. Thus, rather than literally mapping statistical relationships to causal processes (e.g., habit and attitude interacts literally in a way described by interaction effect in linear regression models), it may be more fruitful to generate more precise statistical hypotheses at different temporal scales using computational models of the underlying behavioral processes.

To further appreciate the difference between statistical models and process theories, one can consider the reciprocal relationship between habit strength and behavior as an example. As discussed already at length, habit strength can influence behavior either through its contribution to the activations of behavioral options (see Chapter 2) and as a mechanism to shift one's baseline preferences for the habitual behavioral options (see Chapter 3). Reversely, habit strength increases through behavior execution nonlinearly as described in 4.6.1. Given these mechanisms, although people with stronger habit obviously tend to perform the behavior more often, the corresponding within-person effect may not be straightforward. Suppose that a person is learning to brush their teeth in the evening as a completely new habit, this person might comply with the behavior every day in the first week even though habit strength the beginning of the week is near zero. For the second week, because the preceding habit strength is much higher, if the person somehow omits brushing for two evenings, a negative within-person correlation between habit strength and behavior might be produced if the two week's data are analyzed statistically.

In addition, the within-person effect of habit strength on behavior is likely to be nonlinear, in the sense that at a certain level of habit strength the effect is saturated but habit strength itself can still increase to its possible maximum. If this assumption is correct, then the complex interaction between habits and goals at the cognitive level may also not lead to a straightforward statistical prediction of a negative interaction term as in linear regression models. Moreover, while a statistical interaction effect assumes relative independence between the predictors (e.g., the distribution of attitude does not change at different levels of habit strength), in natural habit formation processes the combination of strong habit and very negative attitude is rarely to be found. Thus, to really test the moderation role of habit strength, experimental manipulations might be needed to change the attitude of people with strong or weak habits, in a way similar to the devaluation paradigm used in laboratory learning experiments (e.g., Dickinson, 1985).

4.5.2 Implications for objective behavior measurements in habit research

One particular methodological contribution of our studies is the use of sensor-based objective measures of behavior in habit research. Although objective behavior monitoring is becoming more accessible due to the advancements in sensor technologies, applications of sensor-based

measures in psychological research on habits is still very scarce (see Gardner, 2015). The success of our application is evident in the strong correlation between sensor-based and self-reported weekly brushing rate, and in the comparable results for the main research questions of interests when using the two types of measures. Our data also give an indication of the common-method bias, as the correlations of interests were almost consistently higher when self-reported behavior rate was used than when sensor-based behavior rate was used.

The strong correlations between self-reported and sensor-based behavior rates may lead some researchers to question the benefits of using the more complex technology-based approach. If the information provided by the two types of measures is redundant, would it be more convenient just to use retrospective self-reports? Here we defend the use of objective measures for a few reasons. First, even though the two type of measures correlated highly, sensor-based measures may still be more accurate and potentially lead to more precise estimates of effect size. This would be especially true when some participants in a study have memory problems or difficulties in understanding rating scales. For example in our second study with some older and less educated participants, there was clear evidence that on a few occasions the self-report measures were erroneous. The advantage of sensors over human memory systems is likely to be continuously enlarged by technology developments. Second, even if weekly self-reports are accurate, day-level behavioral data or even continuous behavioral monitoring are only feasible with sensors in real digital interventions. Self-reporting every occasion of a target behavior is a lot of burden for users and it may affect the primary research on behavior change in unintended ways. Third, although not explored in this chapter, sensor-based methods can actually collect much more information than the occurrence of behavior. For example, from our accelerometer data, brushing duration, intensity, and even temporal orders of brushing locations can be obtained, which are even more cumbersome for self-reporting. Bearing these advantages in mind, our approach can be further strengthened by calibrating the preprocessing algorithms with labeled brushing data (ground truth) and by reducing the physical salience of the sensors in the future.

4.5.3 Potential of using computed habit strength for behavior prediction

Our predictive modeling work shows some promises for using computed habit strength to improve behavior prediction. Through a nested cross-validation procedure on the data from both studies, the theory-based models that include computed habit strength and option accessibility as features performed better than the survey models that used self-reported behavioral determinants as features. With Study 2's data, the theory-based models also performed better than the models that are based simply on past behavior rate, although this advantage was not shown when model performance was measured in terms of predicting new data. Although the differences in prediction accuracies between the theory-based and past-behavior models are quite small, the two types of models actually make somewhat distinct predictions: The theory-based models tend to predict more negative cases (brushing) and thus produces

higher true negative rates at the cost of lower true positive rates. Conversely, past-behavior models predict more positive cases (non-brushing) and thus achieves higher true positive rates at the cost of lower true negative rates. If one examines their prediction patterns, it can be found that the above difference is mainly attributed to the early phases of individuals who often alternate between brushing and non-brushing. Because of the modeled dynamics of habit formation and the very small habit decay parameter (0.05 - 0.2), the occasional omissions of brushing (positive cases) do not affect habit strength very much, so the theory-based models seem to “forgive” these positive cases more easily. In contrast, when the total number of observations is still small at the early phase, omissions of brushing can lower past behavior rate quite substantially and thus more predictions of positive cases are made.

Besides the interests in behavior prediction, the parameter estimation procedure used in our studies also has implications for the theoretical understanding of habit formation. The optimal values tuned for the habit gain parameter are very close to the corresponding values obtained through a statistical modeling of the temporal dynamics of self-reported habit strength or behavioral automaticity (0.19 in Lally et al., 2010), although the optimal values for the habit decay parameter were larger than the 0.007 from Lally et al. (2011). In general, these results speak to the theoretical meaningfulness of the computational model of habit strength used for prediction. In contrast, the parameters in the equation of accessibility did not seem to have optimal values, which casts doubts onto the validity of modeling memory accessibility in the current way. The contrasting result is not surprising given that the computational modeling of habit strength has more theoretical foundations and has been researched more extensively.

We initially thought that the computed cognitive variables would increase the predictive power of model based on other commonly used features, such as behavioral determinants (e.g., attitude) and contextual factors (e.g., emotional states, environmental cues). Therefore, we were surprised that combining the computed habit strength with either instrumental or affective attitude did not perform better than the theory-based models alone (not reported in the result sections). Of course, the measurements of attitude were on the weekly level, so it is yet to know whether knowing more immediate contextual factors would further increase prediction accuracy of brushing behavior (e.g., sleepiness of the person in the evening, behavior of the partner, etc.). Without knowing these information, the current prediction accuracy of around 65 - 75% might be the limit.

Although the equation of habit strength was motivated by theories (e.g., Klein, Mogles, Treur, & van Wissen, 2011; Miller et al., 2019), the computed variable also represents a specific summary of past behavior. The similarity between the theory-based models and the past-behavior models was also reflected in the fact that they seemed to provide similar information, since adding these features together did not improve performance much further. Without knowing the true effect size of habit strength’s influence on behavior, optimizing the parameters in the

equation to produce the best prediction performance might return parameter values that transform the computed variable to capture other influences other than habit strength on behavior. Compared with past-behavior models that weight each behavior in the past evenly, the equation of habit strength weights behavior at different time point in the past in a more sophisticated way. Given the habit decay parameter, the contributions of behaviors that are far in the past to current habit strength are discounted in an exponential way, given by the decay parameter to the power of n (HDP^n), where n denotes the number of time steps to the past. Behaviors in the later stage of habit formation also tend to have increasingly smaller immediate contributions to the current habit strength because the habit gain parameter is modulated by the term $1 - HS_t$. For the purpose of behavior prediction alone, it would be interesting to examine more closely the mathematical properties of the equation and to explore whether other ways of weighing past behaviors could result in better prediction performance.

4.5.4 Implications on using theory-based computational models in digital interventions

Although more research is clearly needed, the theory-based computational approach used in our studies can be potentially used in digital intervention systems. As long as behaviors are observed by the system's sensors and parameter values are estimated from existing data, such a system can update its representations of users' cognitive states after every daily decision, bypassing the need of asking users to report their cognitive states repeatedly. This type of information can be used in two ways. First, tracking a user's habit strength of a newly trained behavior may give the system a better idea about the progress of behavior change. For example, even when the target behavior is already consistently performed, a habit strength less than its maximum would suggest some rooms for improvement, before the current coaching can be stopped at a minimum risk of relapse. Second, the computed cognitive states may indeed assist the system's behavior prediction, which can then be used to inform intervention decisions. For instance, although sending reminders did not cost much economically, psychologically too many reminders may irritate users. As a result, reminders are better to be saved for the occasions that the system has confidence in anticipating behavior omissions from the users. In conclusion, a better understanding and modeling of the dynamics of habit formation should motivate more intelligent digital intervention systems in the near future.

Chapter 5

Evaluating Mouse-tracking as a Technique to Reveal Self-Control Processes

5.1 Introduction

The self-control problem in many daily lifestyle decisions can be described as a decision dilemma (see Berkman et al., 2017). Typically, a decision-maker is faced with two (almost) equally attractive options, but they are attractive for opposite reasons. For example, one option gratifies the immediate desires of the person, but it comes with the potential cost of impeding their long-term health, and thus the impulse towards this option is to be “controlled”. The other option, although providing less or no instant gratification, complies nonetheless with their conscientious of living a healthy life and with the norms and expectations from today’s society. There is a strong trade-off to be made in making such decisions, as accepting either of the two desirable aspects means rejecting the other. People face these trade-offs all the time in daily lives, such as choosing between tasty French fries and healthy salads, or deciding to make the extra effort to floss teeth or not. With today’s digital technologies, people’s actual decisions in these situations can be tracked by sensors (see Chapter 4), their personal health standards and social norms can be self-reported, but their cognitive processes crucially involved in representing and resolving this decision dilemma remain difficult to be revealed.

It is without doubt that revealing these cognitive processes is of crucial importance for advancing the theoretical understanding of self-control. For instance, one may ask the question why people tend to be spontaneously attracted more to the gratifying but unhealthy options, to the point that the tendency has to be “controlled”. Many studies have shown show, for example, the aspect of taste has much larger impact than the aspect of health on people’s dietary choices (e.g., Sullivan et al., 2015; Hare et al., 2009; this Chapter). If one accepts the general assumption in most decision theories that people compare options based on their

¹⁸ This chapter is partly based on Zhang, C., Willemsen, M., & Lakens, D. (2018). Can Mouse-tracking Reveal Attribute Processing Speeds in Dietary Self-control? Commentary on Sullivan et al. (2015) and Lim et al. (2018) with A Simulation Study. *PsyArXiv*. doi: 10.31234/osf.io/725vp.

attributes, the answer to the above question must lie in how people represent attribute healthiness differently than the other attributes, such as tastiness and effort. Several hypotheses have indeed been proposed, for example, that the value of health is discounted due to the delay of gratification (e.g., Ainslie, 1975), that healthiness is more abstract and processed more slowly (Sullivan et al., 2015), or healthiness does not capture automatic visual attention (Motoki, Saito, Nouchi, Kawashima, & Sugiura, 2018). But testing these hypotheses requires a method to reveal unobservable cognitive processes.

Going beyond merely measuring behavior also has additional values for improving digital lifestyle interventions. For example, an e-coaching system may observe that two users choose an apple over a chocolate bar on a given day, but their underlying cognitive processes of re-making the decision might be very different. One person might choose the healthy option without any hesitation, while the other barely managed to resist the temptation of the chocolate bar. Similarly, before a user reverses her choices at the behavioral level, the conflict-resolving process at the cognitive level might already be shifted gradually. Thus, knowing the amount of conflicts experienced and self-control effort mobilized by users could inform the system more about the progress of behavior change and interventions needed for consolidating new habits. Besides, although some self-control situations can sometimes be identified based on the profiles of the choice options, this is not always the case given individual differences in perceptions, for example, in taste preferences. One interesting question is how to automatically identify decision situations with strong health-pleasure trade-offs from the situations where decisions are difficult merely because the options are equally good in both aspects.

In this chapter, we evaluate a mouse-tracking technique in the context of self-control, as it has been used for continuously tracking dynamic cognitive processes (for reviews, see Freeman, 2018; Song & Nakayama, 2009; Stillman, Shen, & Ferguson, 2018). Compared with other techniques that have also been applied to self-control, such as eye-tracking (Motoki et al., 2018) and neuroimaging (e.g., Hare et al., 2009; Maier, Makwana, & Hare, 2015), mouse-tracking has the unique advantages that it is cheap and readily applicable to digital systems in people's daily environments. Before outlining specific research aims and questions, a review of the mouse-tracking technique, applications, and its theoretical assumptions is provided.

5.1.1 The mouse-tracking technique and paradigm

In typical mouse-tracking experiments, the experimental task itself is rather simple and can be illustrated with the following example, where participants are asked to categorize the gender of face stimuli shown at the center of the screen (e.g., Freeman & Ambady, 2009; see Figure 5.1). Instead of pressing keys associated with the two gender categories as in traditional reaction-time tasks, participants indicate their responses by moving the cursor at the

5.1 Introduction

bottom-center of the screen to the category labels (“Female” or “Male”) in the top-left and top-right corners of the screen using a computer mouse. The response labels are usually shown after participants click the “start” button and initiate a movement of the cursor, in order to prevent that they make decisions before any movements. When one of the response labels is clicked, a response is recorded, as well as the temporal-spatial profile of the movement trajectory. The raw data of a single trial thus contain the positions of the cursor (x and y-coordinates) and the associated timestamps of the positions sampled at a certain frequency (e.g., 60 Hz).

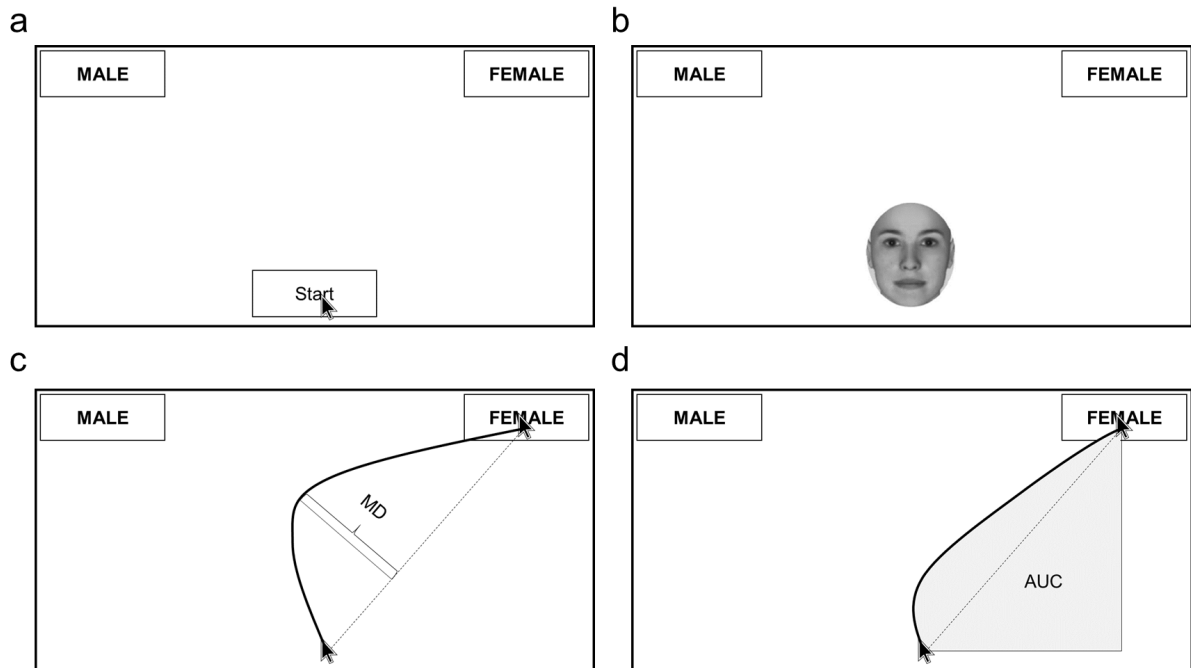


Figure 5.1 An example of a mouse-tracking task (a: click the button to start; b: a face stimulus shown briefly; c & d: two exemplar mouse trajectories from the starting position to the response labels).

Different types of parameters can be extracted from the raw data to describe the characteristics of individual movement trajectories, and to examine the correlations with task manipulations. Most commonly used parameters describe each trajectory as a whole, focusing on either spatial or temporal information (Freeman & Ambady, 2010; Hehman, Stolier, & Freeman, 2015). *Area under curve* (AUC) and *maximum deviation* (MD) are two popular spatial parameters that measure the general tendency of a trajectory to deviate from the straight line connecting the starting position of the cursor and the positions of the response labels (see Figure 5.1c & d), or the tendency to be “attracted” to the non-chosen response. In addition to the general deviation tendency, spatial complexity of the trajectories is measured by *x-flip* and *y-flip*, the number of directional reversals along the x- or y-axis during the whole movement. Because x- and y-flip count all the reversals regardless of magnitude, a different parameter called *number of commitments* was developed recently to measure how many times

a decision-maker actually “changes their minds” (Szasz, Palfi, Szollosi, Kieslich, & Aczel, 2018). This number is computed by counting the number of times the cursor enters the immediate spatial areas surrounding the two response labels, called the “areas of interests”. Aggregated temporal parameters, though used to a lesser degree, include *response time*, the *maximum speed*, and the *maximum acceleration* during the movement. Response time tends to correlate with AUC or MD, but the correlation is often weak, which justifies the use of mouse-tracking to provide non-redundant information. Finally, there are more advanced analysis methods that look at the temporal dynamics of trajectories rather than the aggregated trial-level measures (e.g., Scherbaum, Dshemuchadse, Fischer, & Goschke, 2010; Calcagni, Lombardi, & Sulpizio, 2017).

The most robust and omnipresent finding in mouse-tracking studies is that regardless of what categorization, judgment, or decision-making tasks are used, the more difficult the task, the longer the response time and the larger the spatial deviation from the hypothetical straight movement. The effect itself is not surprising, but the implications of the effect proposed by pioneers of the method have led to the practical popularity and presumed theoretical significance of the paradigm. The earliest application of mouse-tracking was in the field of language comprehension (Spivey, Grosjean, & Knoblich, 2005) and social categorization (Freeman, Ambady, Rule, & Johnson, 2008), where the researchers were interested in comparing two opposing theories of cognition. The traditional view of human cognition, as rejected by these researchers, posits that processing of either spoken language or social information takes discrete stages, and the outputs of a completed earlier stage (e.g., cognition) are passed as inputs to the next stage (e.g., motor-control). The contrasting view is that the processing of competing stimuli involves parallel activations of mental representations (e.g., two different words or social categories). These mental representations continuously compete with each other, and the intermediate states of the competition (i.e., partial information) are cascaded onto the subsequent stages.

Two characteristics of the associated mouse-tracking data were used as support for the parallel partial activation account. First, the initial movements along the midline between the two response labels and the distinguishable onsets of departures from the midline depending on task difficulty (e.g., phonological similarity between the two words, or the gender ambiguity of the face) are believed to reflect a continuous pulling on the cognitive system by the two competing activations or attractors. Second, distributions of spatial parameters (e.g., MD) were found to be closer to normal distributions (though with higher kurtosis) than bimodal distributions, and this pattern was used to reject the stage-based account that completing motor responses are initiated and the incorrect ones are later corrected. The later account would predict a pattern where some trajectories follow direct lines, while others are with extremely large changes of directions. Note that these mouse-tracking results were used to support both the cognitive theories and at the same time the assumption that mouse-tracking

data reflected the dynamic cognitive processes. We shall argue in section 5.14 that such an interpretation of mouse-tracking data is questionable, and the theoretical assumption that movement trajectories are affected continuously by ongoing cognitive process has never been tested independently. Nonetheless, the mouse-tracking technique has since been applied to many other research topics, including judgment and decision-making (e.g., Dshemuchadse, Scherbaum, & Goschke, 2013; Koop, 2013; McKinstry, Dale, & Spivey, 2008), based on the same untested assumption.

5.1.2 Applications of mouse-tracking in dietary self-control research

All applications of mouse-tracking in health-related self-control research have focused on dietary choices. This is not surprising, because compared with other lifestyle decisions (e.g., physical activities or dental behaviors), dietary choices can be easily studied in laboratory tasks without losing much of the ecological validity. Despite the common goal of studying decision conflicts in dietary self-control, several variations of decision tasks with mouse-tracking were developed for specific research aims, including: (1) evaluating or categorizing food stimuli as positive or negative (Gillebaart, Schneider, & de Ridder, 2015); (2) deciding whether or not to eat particular food items (Ha et al., 2016; Lim, Penrod, Ha, Bruce, & Bruce, 2018); (3) or choosing between pairs of food items based on personal preferences (Sullivan et al., 2015) or experimental instructions (e.g., to choose healthy food, Stillman, Medvedev, & Ferguson, 2017).

By asking participants to evaluate food items as positive or negative, Gillebaart et al. (2015) extracted MD from mouse-tracking data as a measure of conflict magnitude, and response time and timing of MD as measures of temporal aspects of conflicts. When comparing trials with self-control success (evaluating healthy food as positive and unhealthy food as negative) and trials with self-control failure (evaluating healthy food as negative and unhealthy food as positive), MD did not differ but success trials were associated with faster response time and earlier timing of MD. Besides mouse-tracking data, self-reported conflicts and trait self-control were also measured. Results indicated that higher trait self-control was not correlated with more conflict as indexed by MD. However, in self-control success trials, participants with higher trait self-control responded faster and their decision conflicts seemed to be resolved faster as indicated by earlier timing of MD.

Ha et al. (2016) did a similar study with school children, with a modification that the participants directly indicated with mouse-movement whether they preferred to eat food items or not. Instead of comparing self-control success trials and self-control failure trials, they compared trials with health foods as targets and trials with unhealthy foods as targets. Results suggested more conflicts – indicated by larger MD, slower responses, and later timing of MD – when children made decisions on unhealthy foods than on healthy foods. The conflicts were

especially strong when they had to reject unhealthy foods. They also reported that children's body mass index (BMI) correlated positively with larger MD in trials with unhealthy food.

Stillman et al. (2017) adapted the mouse-tracking paradigm to real choice tasks, where participants chose between healthy and unhealthy food items (self-control trials), between two healthy food items (comparison trials), or between one healthy food and one non-food item (control trials). Decision conflict, as measured by AUC, was largest in comparison trials, followed closely by self-control trials, and then by control trials with minimum conflicts. This seems to suggest that choosing between two similar options invokes more conflicts than choosing between two options with strong trade-offs, although the effect was limited to healthy options. Moreover, inconsistent with Gillebaart et al. (2015), trait self-control was found to correlate negatively with conflict magnitude (measured by AUC), and this correlation only existed in self-control trials but not in comparison trials. Finally, similar to the earlier studies on language comprehension and social categorization (e.g., Spivey et al., 2005; Freeman et al., 2008), Stillman and colleagues (2017) also analyzed the distribution of trajectory AUC, and used the absence of bimodality in distributions as a support for a continuous conflict-resolving theory of self-control. As they argued, if instead a dual-processing account of self-control is true (e.g., Hofmann et al., 2009), binary trajectory patterns would have been found – some trials with direct movement (no impulse at all) and others with extreme reversals of movement directions (impulse and then inhibition).

The studies above contributed to the exploration and description of the relationships between mouse-tracking parameters, task characteristics, and person characteristics, but they did not follow strongly from any concrete models of self-control, nor did they test directly any mechanism of self-control. In contrast, Sullivan et al. (2015) used mouse-tracking data to test a specific mechanism of dietary self-control following a concrete computational model. If self-control is conceptualized as a process of value-based decision-making (Berkman et al., 2017), the process can be modeled by sequential sampling models or accumulation models, and in particular by a drift diffusion model (e.g., Ratcliff & Rouder, 1998). To model a dietary choice between, for example, an apple and a chocolate bar, the drift diffusion model proposes that the competition between the two options is resolved in time, while the process can be described as a value-signal drifts stochastically between two decision boundaries or thresholds, before eventually the signal exceeds one of the boundaries. The direction and magnitude of the drifts are determined by the attributes of healthiness and tastiness of the two options plus some random noises. Sullivan et al. (2015) also assumed that both healthiness and tastiness take time to be processed, and only after processing they can be integrated into the drift rates of the value-signals. Based on the overall setup of the model, the question why people weigh tastiness more than healthiness in dietary choices can be explained by two mechanisms: (1) Healthiness is processed slower than tastiness (as being more abstract), so it has less opportunity to influence the accumulation process (latency difference); (2) Healthiness is weighted

less than tastiness (e.g., due to temporal discounting) when integrated into the drift rate of the value signal (slope difference).

By analyzing the temporal unfolding of movement trajectories in dietary choice tasks with mouse-tracking, Sullivan et al. (2015) claimed to find evidence for the mechanism of latency difference. Raw trajectory data were normalized to 101 timesteps, and then self-reported tastiness and healthiness differences were correlated with trajectory angle (i.e., the angle measuring the momentary cursor position relative to the starting position) at all the timesteps. Processing speeds or latencies of the two attributes were estimated as the earliest timesteps at which their correlations with angle became significantly larger than zero. Results consistently show that the latency estimates for healthiness are larger than the latency estimates for tastiness. Using the same method, a more recent study claimed to find that by providing participants with calorie information, the latency difference between processing tastiness and healthiness was reduced (Lim et al., 2018). If these claims are valid, the results would have profound implications for self-control theories and interventions.

5.1.3 Theoretical assumptions underlying the mouse-tracking method

Most mouse-tracking studies are based on the assumption that movement trajectories are affected continuously by the ongoing cognitive process of resolving the competition between response options. For example, Spivey et al. (2005) concluded that “our present findings virtually project the ongoing output of the language comprehension process onto a two-dimensional action space in which the potential goal objects act like attractor points and the manual movement serves as a record of the mental trajectory traversed as a result of the continuously updated interpretation of the linguistic input” (p. 10398). Similarly, Stillman et al. (2017) introduced mouse-tracking as a method that “captures, in a nonobvious and unobtrusive manner, the real-time temporal profile of conflict during a successful self-control choice” (p. 1). Surprisingly, this strong assumption has never been directly tested and the general confidence in this assumption might be attributed to a logical fallacy passed on from the earliest studies (Spivey et al., 2005; Freeman et al., 2008). In these two studies, the basic logic was that given the assumption that mouse-tracking provided a continuous measure of cognitive processes, certain data patterns (e.g., correlations between MD and task difficulty, a lack of bimodality in MD distribution, etc.) would strongly favor the parallel partial activation theory over the stage-based theory. However, it seemed that when the expected data patterns were found, the authors used the results not only as a support for the theory, but also for the assumption itself.

The continuous-mapping assumption is not the only assumption that can explain the basic effect that trajectory deviation correlates positively with task difficulty. We consider three different assumptions here, ordered by the strength of the coupling between cognition and motor-control required by them. The *strong assumption* of a continuous mapping between

cognition and motor-control is shared by most mouse-tracking researchers. It can be described more concretely by sequential sampling models: the momentary positions or directions of trajectory movements are partially determined by the status or drifts of the value-signal in the cognitive process. For example, if a value signal is momentarily attracted mainly by the right response, the hand movement is causally biased rightwards. Not all but some sequential sampling models (e.g., decision field theory, Roe et al., 2001) also assume that the drift rate at one time point is only affected by only one attribute, e.g., either by healthiness or tastiness. With this additional assumption, the strong assumption should predict clear differences between trajectories in trials with two different kinds of difficulty – a strong trade-off between options (*trade-off trials*, or self-control trials), or two options similar in all aspects (*similar trials*, or comparison trials). Trade-off trials should invoke large momentary drifts to both directions that may be cancelled out over time, while similar trials should lead to only small momentary drifts. Mapping this difference to movement trajectories, trade-off trials should be more complex than similar trials, as measured by e.g., x-flips. Thus, finding the differences between these two types of trials would provide stronger empirical support for the strong assumption than the basic difficulty-deviation relationship.

A weaker assumption, or the *moderate assumption*, also demands simultaneous motor movements when a decision-making process (e.g., a sequential sampling) is ongoing, but it assumes that the movements are always upwards before decisions are made (Sullivan, Hutcherson, Harris, & Rangel, 2019). This strategy is functionally efficient because it reduces the distances to both response targets. Only when choices are committed, people move in a straight line to the response targets. In this way, decision-making influences the distance of the upwards movement, but it does not affect the momentary movement direction or cursor position as in the strong assumption. It is nonetheless sufficient to produce the basic effect that more difficult tasks lead to larger and later maximum deviations. However, the moderate assumption does not predict different trajectories for trade-off and similar trials, because regardless of what causes the conflicts, equally difficult decisions will lead to the same amount of upwards movements.

The weakest one of the three, or simply the *weak assumption*, does not posit any coupling between cognition and motor-control. Participants in mouse-tracking tasks may simply make their choices in mind and then start to move, so the size of trajectory curvature is entirely determined by some motor-control properties (e.g., people naturally move with some curvatures rather than completely straight). In addition, if participants do change their minds in some difficult trials and reverse the movements to the opposite response targets, the same positive correlation between task difficulty and deviation can be found. Note that the commonly used strategy of detecting bimodality does not rule out this possibility. The number of reversal trials could be very small, and empirically this is consistent with the often very positively skewed distributions of MD or AUC (e.g., see Stillman et al., 2017). These patterns do

not necessarily produce bimodality, but statistical tests of no difference between the distribution means will yield significant effects.

5.1.4 Open issues before applying mouse-tracking to digital interventions

One important consideration for applying the mouse-tracking technique to digital lifestyle interventions is not theoretical but pragmatic. If one uses it to study decision conflicts in daily environments, the tasks have to be adapted to mobile devices, usually with much smaller displays and touch-control instead of a computer mouse. Although the adaptation of mouse-tracking to different screen sizes and various control devices is sometimes discussed in methodological papers (e.g., Freeman & Ambady, 2010), we could only identify two previous studies that used the mouse-tracking technique in non-traditional settings. In a study by Buc Calderon, Verguts, and Gevers (2015), a mouse-tracking task was performed on a 17-inch Wacom LCD tablet, placed on tables with a 30° orientation, and with a cordless pen to control digital widgets. In an even more “mobile” setting, Wirth, Pfister, & Kunde (2016) asked participants to perform “mouse-tracking” by dragging a cursor on iPads to response targets in portrait mode. However, even in the latter study, there was no direct comparison between a traditional physical-mouse condition and a touch-screen condition. Given the drastically different screen sizes and different motor-control constraints embedded in moving a mouse (e.g., Phillips & Triggs, 2001) and in moving one’s fingers (e.g., Dillen, Phillips, & Meehan, 2005), it is necessary to directly compare the trajectory profiles and test whether the differences influence the paradigm’s ability to reveal theoretically interesting effects or not.

Two use cases for the mouse-tracking technique in digital intervention systems can be identified: (1) to measure the strength of decision conflicts as a proxy of tracking behavior change progress; (2) to distinguish the type of decision conflicts, e.g., trade-off or similar, faced by users in making lifestyle decisions. Both cases require moderate to strong correlations of mouse-tracking parameters with conflict strength, and with different decision scenarios. In most theory-driven mouse-tracking studies, researchers primarily care about explanation so they are content with any statistically significant differences between conditions or correlations with task factors. However, for the applied purpose of prediction or measurement, the sizes of effects are crucial, which have not been well documented in the literature yet. Moreover, although Stillman et al. (2017) found a small MD difference between experimentally defined trade-off trials versus comparison trials, more systematic investigation is needed to cover a larger range of trade-off trials and similar trials as perceived by users, and to document effect sizes for different mouse-tracking parameters.

5.1.5 The current investigation

The overall goal of this chapter is to evaluate the usefulness of the mouse-tracking technique in revealing cognitive mechanisms underlying self-control and its applicability to measuring decision conflicts in digital intervention systems. We report a food-choice experiment with

both touch-screen and physical mouse conditions to address the two open issues mentioned above.

1. Whether the mouse-tracking technique can be adapted to touch-screen devices.
2. Whether mouse-tracking parameters are sensitive enough to measure decision conflicts or to classify different decision scenarios.

Because the results of the first experiment had casted more doubts to the default strong theoretical assumption in the field, additional work as done to scrutinize the different theoretical assumptions mentioned earlier. First, a replication of Sullivan et al. (2015) and a simulation-based re-evaluation of their method of revealing self-control mechanisms are reported. Also, a second experiment was conducted to further differentiate the three theoretical assumptions underlying the mouse-tracking technique.

5.2 Study 1: food-choice experiment with touch-screen and physical-mouse

5.2.1 Method

Participants & design

Forty-three students (23 males and 20 females) from Eindhoven University of Technology participated in the experiment as part of a course fulfillment. Their age varied from 19 to 31 years-old, with a mean of 21.5 years and a standard deviation of 1.9 years. Their BMIs were between 17.2 and 27.8 (mean = 22.3, $SD = 2.57$). According to the questionnaire used in the study, three participants reported to be vegan, one reported to be vegetarian, and one reported to have a low-carbohydrate diet. In addition, three participants indicated that they had hay fever, Celiac disease, or were lactose intolerant.

In a within-subject design, 180 food-choice trials were evenly divided into two blocks – one with a laptop and a mouse and one with a tablet. The order of the blocks was counterbalanced between participants. Moreover, as we were also interested in whether nudging could affect movement trajectories, in some trials additional information was shown underneath the food images to emphasize either the healthiness or tastiness difference between the two food items. Because this was beyond the scope of the current chapter and no effect was found (nor did it affect our main results), data analyses for our main objectives treated all trials as equal. Finally, 10 filler trials in each block were used to measure participants' natural movement profiles when no food-choice was involved.

Food stimuli

Food images were selected from the Standardized Food Image Database provided by the Image Sciences Institute, UMC Utrecht (Charbonnier, van Meer, van der Laan, Viergever, & Smeets, 2016). This database has the advantage that all food images were presented in a standardized manner, eliminating effects from visual differences in presentation. Based on the healthiness ratings of the food items from a Dutch sample ($N = 136$, Charbonnier et al., 2016) and their tastiness rated by a sample of Germans and North Americans ($N = 1988$, Blechert, Meulie, Busch, & Ohla, 2014), we selected 5 food items that were expected to score high on healthiness but low on tastiness, and 5 food item that were expected to score high on tastiness but low on healthiness. Each pair of food items was presented 4 times in the experiment and positions of the items (top-left or top-right) were counterbalanced. In filler trials, one of the two images was an empty plate without food, photoshopped from the food images, and the task was always to select the image with a food item.

Apparatus and measurements

Because the mouse-tracking paradigm was adapted for Android tablets, we developed our own application using Kivy - a Python library for building multi-touch applications across platforms (<http://www.kivy.org/>). In the touch-screen condition, the application was installed on Samsung Galaxy Tab S3 SM-T210 tablets, with a 7-inch screen (1024 x 600 pixels). In the laptop condition (with mouse), the application was installed on Window systems and was presented with a window size that matched exactly with the full screen size of the tablets. Therefore, the window sizes in both conditions were smaller than the one used in the original study.

During each food-choice trial, a blue round cursor first appeared at the bottom-center of the screen, and only when the participants started to drag it, the two images of the food items were shown at the top-left and top-right (see Figure 5.2). Additional nudges, if any, would appear underneath the food images to emphasize the healthiness or tastiness difference between the two items. Positions (x and y coordinates) of the movement trajectories and timestamps were tracked at a frequency of around 67 Hz. When the cursor was moved to the area of one of the food images and released, actual choice of that trial was recorded. Healthiness and tastiness of each food item were rated by the participants on 7-point scales using the same application before the food-choice task. After the food-choice task, basic information, including age, gender, BMI, special diet, and allergies, was measured using the application.

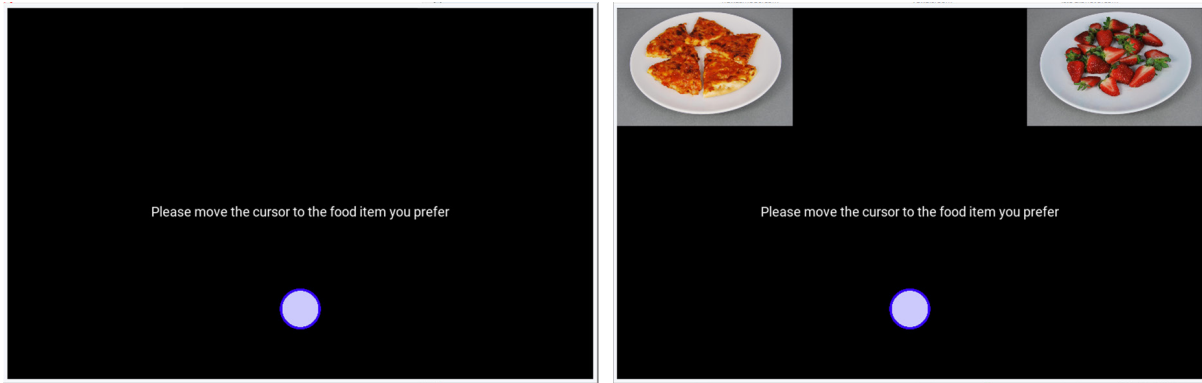


Figure 5.2 The mouse-tracking task in the food-choice experiment. When participants click and move the blue cursor (left), the food images are shown (right).

Procedure

Participants were invited to the lab for the experiments in a group of 4. After introduction and signing of consent forms, they were assigned to one of the four desks to perform all the tasks independently. They were not able to see the behaviors of other participants and they were told not to interact with each other. A table in the middle of the room was filled with real food items, the same ones as in the food images shown in the application. They were told by the experimenters that they would be given one food item to eat, randomly chosen from all the food items they chose during the main task. Participants performed the attribute rating task first, and then continued with the two blocks of food-choice tasks using the laptops and tablets. When tablets were used, they were vertically placed against the screen of the laptop in order to control screen orientation. There was a 60-second break between the two blocks. Finally, participants answered the short questionnaire about the general information mentioned above. At the end, they were debriefed and were asked to eat the randomly chosen food item, unless they had very strong reasons not to eat.

Data preprocessing

A time and space normalization was used on the raw trajectory data, following the standard processing procedure for mouse-tracking data (e.g., Freeman & Ambady, 2010). Specifically, trials with different durations were normalized to 101 time points and the x and y-coordinates at all time points were computed based on the raw coordinates. The new coordinates were then shifted and normalized to a coordinate space of two squares, from [0, 0] to either [-1, 1] (left item chosen) or [1, 1] (right item chosen). After normalization, problematic trials were removed based on predefined criteria. Firstly, trials in which the participants released the cursor prematurely were removed (6%). Secondly, 9 trials with technical faults, indicated by extremely large x and y-coordinates or negative area under curve (AUC), were removed (0.1%). Thirdly, trials with (log-transformed) reaction times larger than three standard deviations from the grand mean were removed (1.2%). Fourthly, trials in which the participants changed their minds more than twice were removed (0.8%; detected using areas of interests,

see Szaszi et al., 2018). Finally, one participant's data from the laptop condition were missing, and one participant's data were unusable due to an unfinished attribute rating task. The final dataset for analysis consists of 6932 food-choice trials and 794 filler trials from 42 participants.

5.2.2 Results

Comparison of trajectory parameters between the two control conditions

Three different types of statistics were used for the comparison between the two control conditions: (1) the density distributions of the trajectory parameters; (2) the correlations between trajectory parameters; (3) and the percentage of variance in these parameters accounted by participant (as intra-class correlation estimated from multilevel null models). We focused on the most commonly used parameters, including AUC, MD, and x-flip as spatial parameters, and response time (RT)¹⁹, maximum velocity, and maximum acceleration as temporal parameters. Because the sample size at the trial-level was very large, even very small differences between the two conditions would be statistically significant. Thus, we mainly evaluated whether the effect sizes of the differences or any substantial qualitative differences (through visualizations) would affect the validity of using mouse-tracking on touch-screen devices.

Figure 5.3 shows the density plots of the trajectory parameters for the two control conditions separately. Except for x-flip and response time, the distributions of all parameters in the touch-screen and mouse conditions overlapped greatly. The small existing differences were mostly not in the central tendencies of the distributions (e.g., the modes), but in the distribution tails that represented extreme values. For example, for the two spatial parameters AUC and MD and for maximum velocity, there were more large extreme values in the mouse condition than in the touch-screen condition. In contrast, the touch-screen condition produced more extremely large maximum accelerations than the mouse condition. The distributions for x-flip and response time were more clearly separated, suggesting that more horizontal reversals were identified and responses were slightly slower in the touch-screen condition than the mouse condition. When the parameters were subjected to multilevel models with control condition as the predictor, results showed the same picture: although statistical significant differences were found for all parameters, the explained percentages of variance by control modes were extremely small in most cases (all marginal- $R^2 < 0.015$). Only for x-flip, control mode could account for a sizable of 32.0% of the variance measured by marginal- R^2 . On average, participants horizontally reversed movement directions around 1 more time when using a touch-screen than using a physical mouse ($B = 0.95, p < .001$).

¹⁹ Response time was measures as the time between touching on the cursor and releasing the cursor at the target positions.

Table 5.1 shows the zero-order correlations between pairs of trajectory parameters. Overall, the relationships among parameters were consistent between the two control conditions, and unsurprisingly the two deviation measures AUC and MD correlated highly and temporal parameters also correlated moderately to highly among themselves. Moreover, in both conditions, the correlations between response time and the deviation measures were relatively small, which indicated that AUC or MD provided different information than response time. The major differences between the two control modes were again in x-flip: it correlated more with the deviation measures in the mouse condition, but more with response time in the touch-screen condition. Furthermore, it seemed that maximum velocity and maximum acceleration correlated positively with the deviation measures more strongly in the mouse condition than in the touch-screen condition.

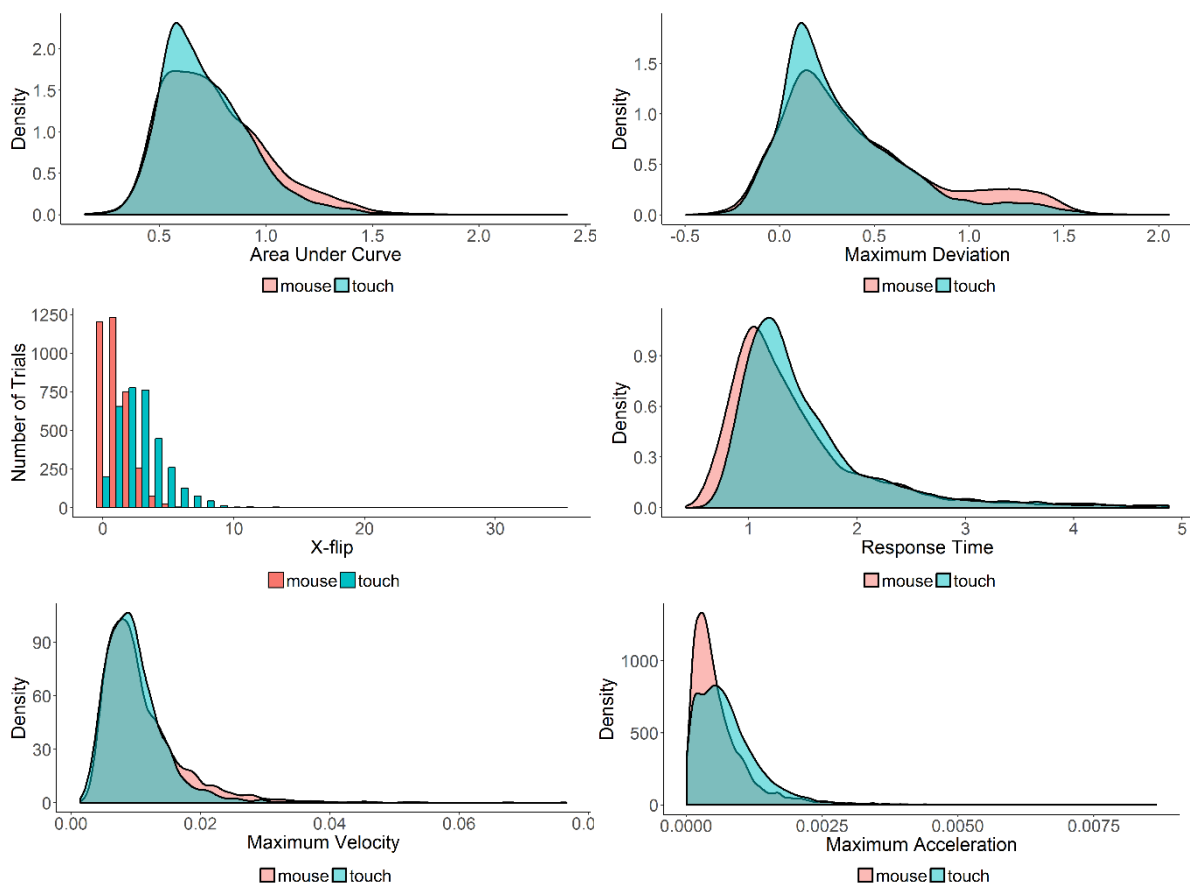


Figure 5.3 Density plots for the trajectory parameters (histogram for x-flip).

5.2 Study 1: food-choice experiment with touch-screen and physical-mouse

Table 5.1 Zero-order correlations between pairs of parameters (coefficients in the mouse and touch-screen conditions were shown below and above the diagonal respectively).

	1	2	3	4	5	6
1. AUC	1	0.89	0.18	0.17	0.03	-0.07
2. MD	0.86	1	0.24	0.27	0.00	-0.13
3. x-flip	0.41	0.55	1	0.49	-0.12	-0.25
4. RT	0.08	0.18	0.26	1	-0.50	-0.64
5. MV	0.30	0.39	0.23	-0.23	1	0.91
6. MA	0.12	0.11	0.00	-0.48	0.78	1

Note: MV = maximum velocity; MA = maximum acceleration.

In terms of percentages of variance explained by inter-individual differences, the intra-class correlations (ICCs) for most parameters were in the range of 20 - 35%, and the differences between the two control conditions were very small (*AUC*: $ICC_{\text{mouse}} = 0.230$, $ICC_{\text{touch}} = 0.286$; *MD*: $ICC_{\text{mouse}} = 0.234$, $ICC_{\text{touch}} = 0.271$; *response time*: $ICC_{\text{mouse}} = 0.250$, $ICC_{\text{touch}} = 0.265$; *maximum velocity*: $ICC_{\text{mouse}} = 0.329$, $ICC_{\text{touch}} = 0.333$; *maximum acceleration*: $ICC_{\text{mouse}} = 0.315$, $ICC_{\text{touch}} = 0.340$). In contrast, inter-individual variation in x-flip was smaller in the mouse condition ($ICC_{\text{mouse}} = 0.160$) and was close to zero in the touch condition ($ICC_{\text{touch}} = 0.012$).

Choice level analyses and estimation of personal decision weights

In order to compute conflict strength and to categorize trials into different types (trade-off versus similar), decision weights for healthiness and tastiness were estimated for each participant first. A simple example can explain why using raw attribute ratings is not sufficient for the requirement computations. When a person rates apple to be much healthier than chocolate bar (7 and 1) but much less tasty (1 and 7), it may appear like a strong trade-off trial. However, if this person weighs tastiness as much more important than healthiness, then choosing between chocolate and apple is actually very easy because the difference in healthiness is not subjectively important for the person. Thus, subjective utility, as multiplication of attribute ratings and attribute weights, should be used instead.

Choice was modeled in a multilevel logistic regression with healthiness difference and tastiness difference between left and right food items as predictors, and random-slopes were included to allow for individual differences in how the two predictors influenced choice. Results indicated that tastiness ($B = 1.68$, $p < .001$) had a much larger impact on choice than health-

iness ($B = 0.17, p = .007$), and the two attributes together explained the majority of the variance in choice (marginal $R^2 = 0.64$). When models were fitted to the data of each individual participant, tastiness and healthiness together were shown to account for 13% to 94% of the variance (median = 55%). Figure 5.4 illustrates the estimated personal decision weights for healthiness and tastiness for each participant. It was evident that tastiness had positive influences on choice for everyone, while for 11 participants healthiness had negative or no influence on choice. Based on the estimated personal decision weights, two new variables were computed:

- *Utility difference* between two food options was calculated as the sum of the absolute differences in healthiness and tastiness weighted by the attribute weights: $|w_{health} \times (Health_L - Health_R)| + |w_{taste} \times (Taste_L - Taste_R)|$.
- *Tradeoff strength* was computed as the absolute difference between the difference scores in healthiness and tastiness weighted by the attribute weights: $|w_{health} \times (Health_L - Health_R) - w_{taste} \times (Taste_L - Taste_R)|$. When the difference scores in the two attributes are with opposite signs (e.g., one positive and one negative), tradeoff strength represents the intuitive meaning of tradeoffs, but it also extends to the situations where both difference scores are with the same sign (e.g., both favoring the left option). To estimate the influence of tradeoff strength on trajectory parameters, we included the utility difference by tradeoff strength interaction term, so that the main effect of tradeoff strength became more intuitively as the estimate would equal to the effect when utility difference between two food options was zero.




Figure 5.4 Distribution of personal decision weights for healthiness and tastiness.

Sensitivity of trajectory parameters to variables describing decision conflicts

Because of the very strong correlations between MD and AUC, and between maximum velocity and maximum acceleration, only MD, x-flip, response time, and maximum velocity were used for the rest of the analyses. For each parameter, three random-intercept models were built with – utility difference (model 1), utility difference, tradeoff strength, and their interaction (model 2), and utility difference, utility of the stronger option, and their interaction (model 3) – as predictors. Because x-flip is a count variable, Poisson distribution was assumed in its models. We included model 3 to explore whether the utility of the stronger option in a pair would actually influence decision process in addition to the utility difference between the options. Intuitively, decisions become easier not only when there are larger differences between two options, but also when (at least) one of the option is very attractive. This intuition is also consistent with a variant of sequential sampling models in which absolute utilities of options rather than the differences or contrasts between their utilities are used as value signals (Bhatia, 2013; see also our model in Chapter 3).

Table 5.2 summarizes the analysis results. First of all, model 1 made it clear that utility difference had significant effects on MD, x-flip, and response time – when the utility difference between two food options was smaller, trajectories showed larger spatial deviation and more reversals of directions, and response became slower. Their effects were nonetheless very small, as the percentage of variance explained was between 0.012 and 0.043. Secondly, adding tradeoff strength to the model (model 2) almost did not improve model fits at all, and the estimates for the effects of tradeoff strength when utility difference was zero were very close to zero. Thirdly, model 3 showed that adding the utility of the stronger option substantially improved the model fits, and very often the new variable had an even larger impact than utility difference. Among all mouse-tracking parameters, response time was slightly more sensitive to the variables describing decision conflicts than MD and x-flip. Finally, in most cases, data from the two control conditions yielded nearly identical results, further suggesting the high similarity between trajectories generated on a touch-screen and with a physical mouse. The only exception was that maximum velocity was negatively associated with conflict strength in the touch-screen condition, but not in the mouse condition.

We also looked at whether these associations were mainly driven by the more extreme trials where participants changed their minds at least once, by fitting model 3 to the data with these trials excluded (16.9% of all trials). For MD, x-flip, and response time, the effect sizes (marginal R^2) were indeed attenuated, with the effect size of MD suffering the most (MD: 0.088 to 0.019; x-flip: 0.062 to 0.033; response time: 0.11 to 0.087). There was also a surprising increase of effect size for maximum velocity from 0.005 to 0.036.

Table 5.2 Sensitivity of trajectory parameters to variables describing decision conflicts. Model estimates before and after “/” were based on data from the mouse and touch-screen condition respectively.

	MD	x-flip	RT	MV
<i>Model 1</i>				
Utility difference	-0.15 ^{***} /-0.10 ^{***}	-0.15 ^{***} /-0.08 ^{***}	-0.21 ^{***} /-0.20 ^{***}	0.02/0.10 ^{***}
marginal R^2	0.018/0.012	0.025/0.020	0.042/0.043	0.0004/0.013
<i>Model 2</i>				
Utility difference	-0.18 ^{***} /-0.14 ^{***}	-0.19 ^{***} /-0.10 ^{***}	-0.25 ^{***} /-0.23 ^{***}	0.02/0.10 ^{***}
Tradeoff strength	-0.002/0.0006	0.02/0.001	-0.01/-0.06*	0.009/0.006
Interaction	-0.03 [*] /0.04 ^{***}	0.03 [*] /0.02 [*]	0.04 ^{***} /0.05 ^{***}	0.0003/-0.01
marginal R^2	0.020/0.016	0.025/0.021	0.046/0.051	0.0004/0.014
<i>Model 3</i>				
Utility difference	-0.15 ^{***} /-0.10 ^{***}	-0.15 ^{***} /-0.10 ^{***}	-0.23 ^{***} /-0.22 ^{***}	0.007/0.10 ^{***}
Utility of the stronger option	-0.21 ^{***} /-0.27 ^{***}	-0.17 ^{***} /-0.09 ^{***}	-0.26 ^{***} /-0.20 ^{***}	-0.03/0.11 ^{***}
Interaction	0.05 ^{**} /0.05 ^{***}	0.05 [*] /0.05 ^{***}	0.07 ^{***} /0.06 ^{***}	0.03/-0.03
marginal R^2	0.069/0.114	0.066/0.054	0.115/0.097	0.002/0.035

Note: Significance levels were indicated as $p < .10^+$, $p < .05^*$, $p < .01^{**}$, $p < .001^{***}$.

Sensitivity of trajectory parameters to trial type


The above analyses showed the sensitivity of trajectory parameters to continuous variables describing decision conflicts was small, but one might expect stronger sensitivity to be found if distinctively different types of trials are examined. Therefore, we identified three types of trials – *tradeoff trials*, *similar trials*, and *dominant trials*, from the data based on the measure of utility difference and tradeoff strength (see Box 5.1 for the categorization logic). The categorization resulted in 1047 tradeoff trials (mean utility difference = 0.10; mean tradeoff strength = 3.57), 2207 similar trials (mean utility difference = 0.08; mean tradeoff strength = 0.67), and 1080 dominant trials (mean utility difference = 5.01) (see Figure 5.5 for a visual illustration).

Box 5.1 Logic for the trial categorization.

```

If (utility difference > 2.5):
  If (no real tradeoff presents):
    If (the stronger option was chosen):
      trial type = "dominant";
else if (utility difference <= 2):
  if (no real tradeoff presents):
    trial type = "similar";
  else:
    if (tradeoff strength < 1):
      trial type = "similar";
    else:
      trial type = "tradeoff";

```

**Figure 5.5** Visualization of the categorization of trials.

Results based on random-intercept models turned out to be very similar with the analyses using continuous variables (see Figure 5.6 for the means, 95% CIs, and distributions in the different types). Compared with dominant trials, tradeoff trials had slightly larger maximum deviation ($B = 0.35$, $p < .001$), more x-flips ($B = 0.29$, $p < .001$), and longer response time ($B = 0.41$, $p < .001$). This was not surprising given that the main difference between these two types was utility difference. However, when comparing tradeoff trials to similar trials, which

shared the same level of utility difference, all trajectory parameters did not differ at all (all p s > .394).

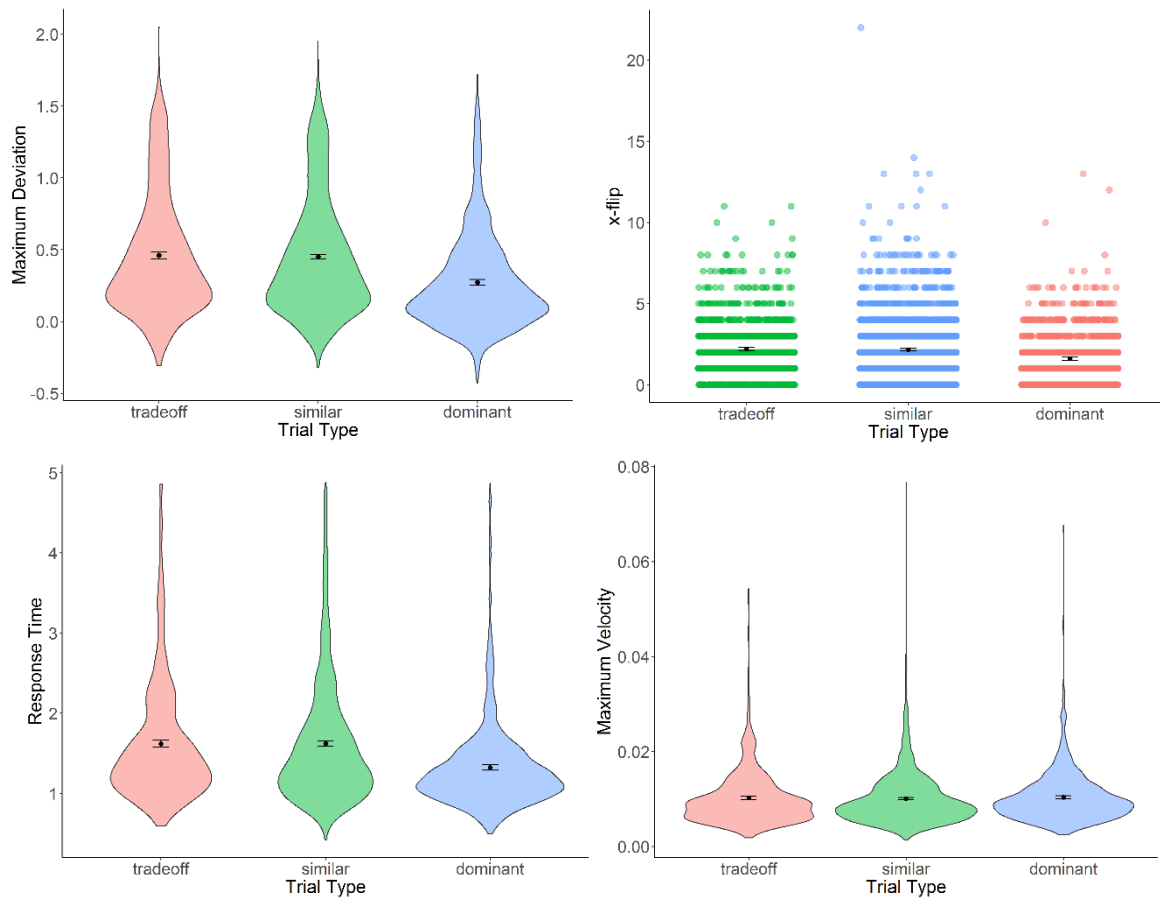


Figure 5.6 Comparison of trajectory parameters in the three different types of trials (error bars represent 95% CIs).

5.2.2 Discussion of Study 1

In summary, the results provide very clear answers to the research questions in Study 1. First, we could conclude that trajectory data generated from touch-screen devices were very similar to those generated using the traditional mouse-tracking setting, as indicated by a variety of parameters. Although some small differences existed, they did not have much impact on the estimations of the other effects of interests. The only sizable difference was that more horizontal direction reversals were found in the touch-screen setting. This was likely to be caused by a technical factor that the spatial resolution of detecting any movement was higher on the touch-screen device than on the desktop (given the different configuration of the Kivy app for Android and Windows system), rather than a human motor-control factor. Thus, if the same setting is used in the future, some filter should be used to exclude x-flips with very small distance.

In terms of the sensitivity of parameters to decision conflicts, our results replicated the general finding that more difficult decisions lead to larger deviation, more reversals, and longer decision time. However, the effect sizes were unlikely to be large enough to be used for accurately measuring conflicts strength at single-trial level, making them less useful for intervention applications than for theory-driven research. We also did not find any difference between the two types of difficulties, i.e., trials with or without tradeoffs. This casts some doubts on the strong assumption of the mouse-tracking paradigm, at least given sequential sampling models that assume sampling shifts among attributes. Motivated by the lack of strong relationship between movement trajectory and decision scenarios, we decided to replicate previous effects relating to self-control using the data from Study 1, and to further examine the different theoretical assumptions underlying the mouse-tracking technique.

5.3 Study 2: (conceptual) replications of previous findings

In Study 2, we attempted to replicate previous findings in studies that applied the mouse-tracking technique to self-control with our own data from Study 1. In section 5.3.1, we discussed several loosely related effects about self-control, including the differences between trajectories in trials with self-control success and self-control failure (Gillebaart et al., 2015; Ha et al., 2016), and the difference between self-control trials and comparison trials (Stillman et al., 2017). In section 5.3.2, evidence suggesting a processing latency difference in Sullivan et al. (2015) was closely replicated. Because at the trial-level the sample size was several thousand, even very tiny effects would be significant at an alpha level of 0.01. Thus, for brevity, we only report detailed statistics (e.g., exact effect sizes, confidence intervals, or p -values) when the effect sizes were at least equal to those found in Study 1 (e.g., marginal $R^2 \geq 0.01$).

5.3.1 Replicating associations between trajectory parameters and self-control

Comparison between trials with self-control success and self-control failure

Although slightly different response labels were used (*positive* versus *negative*, or *eat* versus *not to eat*), both Gillebaart et al. (2015) and Ha et al. (2016) compared if trajectories were different between the trials in which participants controlled themselves successfully (e.g., evaluated chocolate bar as *negative*) and in the trials in which they yielded to unhealthy foods (e.g., chose to *eat* French fries). The findings were inconsistent because while Gillebaart et al. (2015) found faster responses for self-control success trials, Ha et al. (2016) found slower responses and larger deviations for this type of trials (where participants rejected to eat unhealthy food items).

We categorized self-control success and self-control failure using two methods. In the first method, tradeoff trials were defined as in Study 1 as the trials in which weighted healthiness difference and weighted tastiness difference between two options had the opposite signs. The

second method, as used in the previous studies (Gillebaart et al., 2015; Ha et al., 2016), tradeoff trials were those that included one typical healthy food and one typical unhealthy food. For example, *pear* and *donut* would be considered as such a trial, regardless of the actual attribute ratings made by the participants. After tradeoff trials were identified from all trials, trials in which participants chose the healthier option were categorized as *self-control success* trials, and vice versa as *self-control failure* trials. Using both methods, random-intercept models with trial type (success versus failure) as the predictor did not reveal any meaningful difference for all four trajectory parameters considered (MD, x-flip, response time, and maximum velocity; all marginal $R^2 < 0.003$).

Comparison between different types of trials defined by the food stimuli

As discussed in the introduction, Stillman and colleagues (2017) found that choices between two comparable health foods led to more conflicts, as measured by trajectory MD, than choices with tradeoffs. It could be said that we did not replicate the effect in Study 1, as no difference in MD or other parameters was found between tradeoffs trials and similar trials. However, the trial categorization in Study 1 was based on attribute ratings and personal decision weights, rather than a task manipulation as in Stillman et al. (2017). Here we emulated their manipulation by categorizing trials by the typical categories the food options belonged.

Trials were thus categorized into *healthy trials* (two healthy foods), *unhealthy trials* (two unhealthy foods), and *tradeoff trials* (one healthy and one unhealthy food). Using this categorization variable as the predictor in random-intercept models, results did not reveal any significant differences between tradeoff trials and healthy trials in terms of all trajectory parameters considered. However, very small differences were found between healthy trials and unhealthy trials that the latter was associated with larger MD, more x-flips, and longer response time (though all marginal $R^2 < 0.003$). Taking a closer look at the utility differences underlying the health and unhealthy trials, it was evident that the small effects could be attributed to the smaller average utility difference underlying the latter type (see Table 5.3).

Table 5.3 The continuous decision conflict measures (as in Study 1) underlying difference types of trials categorized based on the food stimuli's typical categories.

	Healthy	Unhealthy	Tradeoff	Tradeoff - Success	Tradeoff - Failure
Utility difference	2.02	3.09	3.09	2.89	3.22
Utility of stronger option	10.26	10.50	10.68	10.67	10.68
Tradeoff strength	1.99	2.84	3.43	2.83	3.82

5.3.2 Replicating Sullivan et al. (2015) to test cognitive mechanisms of self-control

Replication of the results using regression coefficients

As with Sullivan et al. (2015), for all the trials with conflicts (one food item scored high on healthiness while the other scored high on tastiness), self-control success rate (SCSR) was computed as the percentage of trials in which the healthier item was chosen. Results showed that SCSR was between 0% to 43%, with a mean of 10% and a standard deviation of 11%. The participants in our study had lower SCSR rates, higher weights for tastiness, and lower weights for healthiness, compared to those the original study. Eleven participants even had negative weights for healthiness, while this was true for only four participants in Sullivan et al. (2015). It seems that our participants generally exerted less dietary self-control than those in the original study, possible because the health education session in Sullivan et al. (2015) was not used in our replication.

Next, trajectory angles at the 101 timesteps were computed as the angles between the lines connecting $[0, 0]$ and the current coordinates and the y-axis. For each participant and each timestep, the variable angle was regressed to healthiness difference and tastiness difference separately to estimate their correlations. Figure 5.7a shows the temporal unfolding of the standardized regression coefficients²⁰ of the attributes on trajectory angle. The general pattern was very similar to the original study, but the coefficients of healthiness had larger standard errors and did not eventually become significant. The reason for this can be seen in Figure 5.7b: The coefficients of healthiness for participants with low SCSR (median-split) actually became to be significantly negative over time. This was again possibly due to the fact that participants in our study were less motivated to execute self-control. Nonetheless, the order of the timesteps when the coefficients of the attributes for sub-groups became significant was exactly the same as in Sullivan et al. (2015). For participants with high SCSR, the earliest timestep for tastiness and healthiness to have significant and enduring influences were 43 and 54 respectively; for participants with low SCSR, the gap between the timesteps for the two attributes was much larger (0 and 79). In our sample, healthiness never had positive influence on trajectory angle for 57% of the participants (35% in Sullivan et al., 2015). For 31% of the participants, healthiness eventually had negative influence on trajectory angle.

For the analyses at the individual level, the 13 participants with eventual significant negative coefficients for healthiness were excluded²¹. For the remaining participants, a paired *t*-test

²⁰ Note that unstandardized regression coefficients were used in Sullivan et al. (2015) and Lim et al. (2018). We used standardized coefficients, which are easier to interpret because they equal correlation strengths between pairs of variables.

²¹ Unlike Sullivan et al. (2015), we only excluded participants with significant negative coefficients for healthiness, but not those with nonsignificant coefficients for healthiness. Including all participants would lead to a stronger effect (mean difference = 18.62, $t(41) = 4.66$, $p < .001$, $d_z = 0.72$).

revealed that trajectory angles were influenced by tastiness (mean = 57.07, $SD = 19.74$) at significantly earlier timestep compared with healthiness (mean = 74.03, $SD = 26.5$; mean difference = 16.97, 95% CI = [6.61, 27.32], $t(28) = 3.35$, $p = .002$, Cohen's $d_z = 0.62$). Moreover, as in the original study, individual estimates of processing speed difference between healthiness and tastiness correlated strongly with participants' SCSR ($\beta = 0.70$, $p < .001$), explaining 44.5% of the variance in SCSR ($R^2 = 0.39$ in Sullivan et al., 2015; see Figure 5.7c).

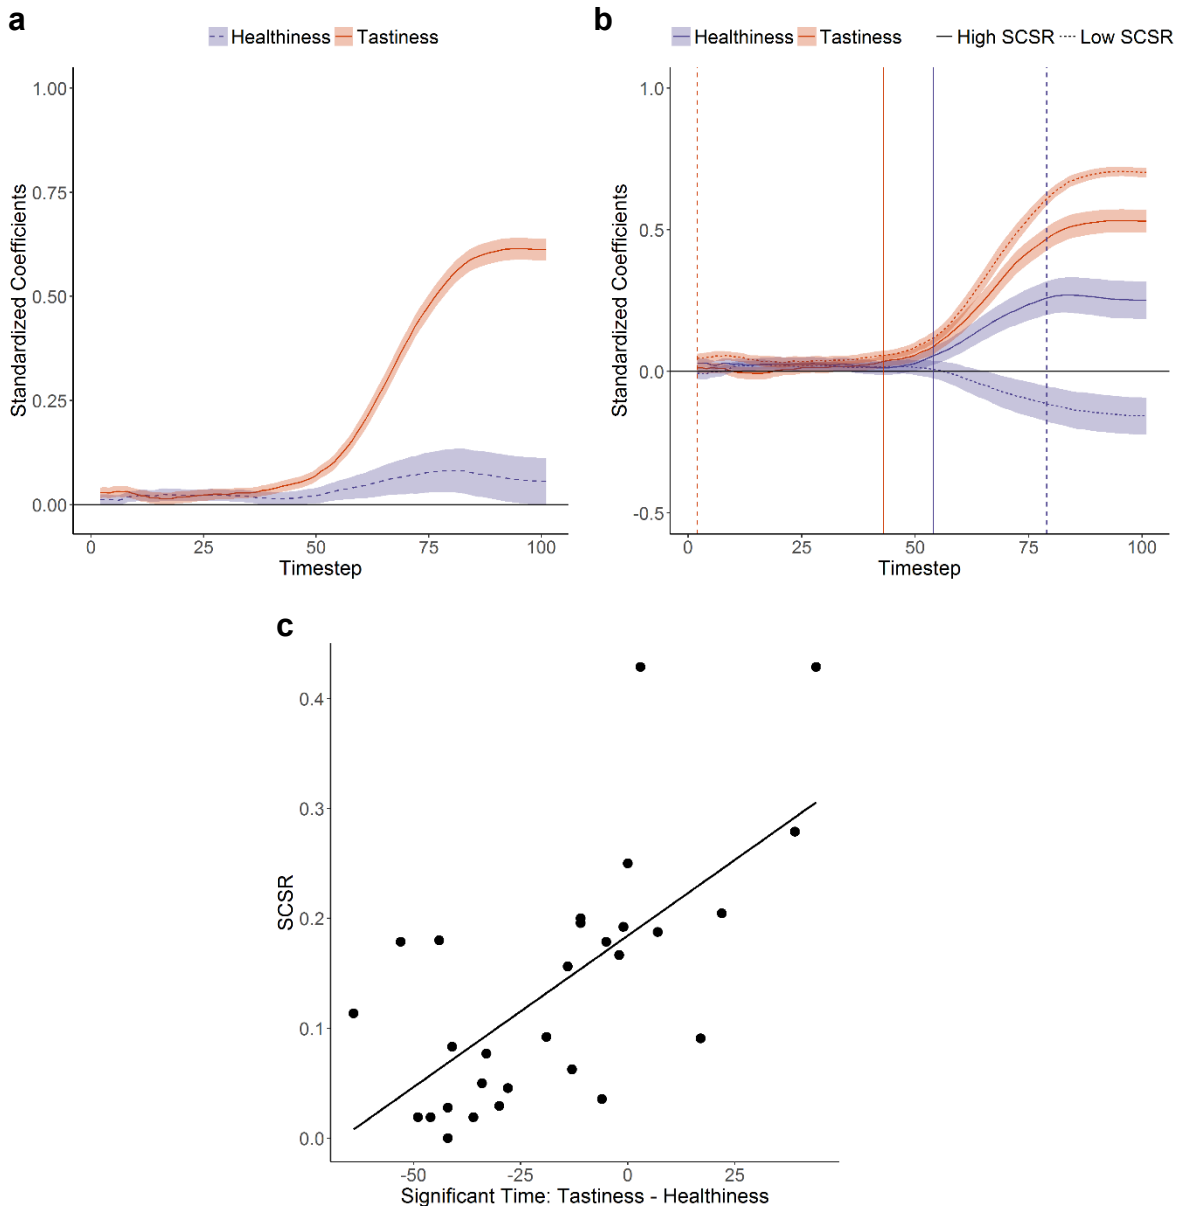


Figure 5.7 Temporal patterns of the regression coefficients of the attributes on trajectory angle for all participants (a), for SCSR sub-groups (b), and (c) correlation between difference in earliest significant timesteps and SCSR.

5.3 Study 2: (conceptual) replication of previous findings

Replication of the results using proportion curves

As noted by Sullivan et al. (2015), the above method of estimating processing speed suffers from a statistical issue. If the coefficients of both attributes increase from zero to their final weights and a similar level of noise applies to both curves, the coefficients of tastiness are destined to become significant earlier than the one of healthiness. Therefore, in a different method, instead of using the raw regression coefficients, timesteps at which the attributes' coefficients exceeded certain proportions (from 0.1 to 0.9) of their final weights were computed. Cubic polynomials were then fitted to the timestep data to indirectly estimate the timing of the two attributes' coefficients when they exceeded the proportion of 0.0 (intercepts of the proportion curves). The intercepts were used as an alternative measure of when healthiness and tastiness began and continued to influence trajectory angle.

Inconsistent with the original results, it can be seen from Figure 5.8 that healthiness was able to exceed certain proportions between 0.1 and 0.9 slightly faster than tastiness. When cubic polynomials were fitted to the data, results did show that the y-intercept of healthiness was slightly larger than the intercept of tastiness (44.5 and 36.4 respectively), but the difference was much smaller compared to the original study.




Figure 5.8 Estimated timesteps at which the weighing of the two attributes exceeded certain proportions of their final weights.

Trajectory curvature in filler trials and processing speed estimates

As filler trials were included in our experiment, we explored whether participants' average trajectory curvatures (MD) in filler trials correlated with their average MD in the food-choice trials, and whether they correlated with the estimated processing speeds. We thought that since no healthiness or tastiness information was relevant for making responses in the filler

trials, the trajectories in those trials were approximations of how participants naturally moved cursors from one point to another. Correlation analysis revealed that curvature magnitude in the filler trials correlated strongly with the curvature magnitude in the food-choice trials, explaining over 60% of the variance of the latter. Thus, any genuine influences from the two attributes on trajectories had to be superimposed onto the strong effects of motor-control.

Results also indicated a significant positive correlation between MD in the filler trials and estimated processing speed ($\beta = 0.39$, $p = .003$, $\Delta R^2 = 0.15$; see Figure 5.9), while controlling for decision weights. In other words, if participants moved hands with larger trajectory curvatures in the filler trials, their estimated processing speeds of the attributes appeared to be slower. This pattern was especially strong for tastiness ($\beta = 0.65$, $p < .001$, $\Delta R^2 = 0.39$), but weaker and nonsignificant for healthiness ($\beta = 0.24$, $p = .179$, $\Delta R^2 = 0.06$). The nonsignificant correlation for healthiness might due to the small sample size and the ceiling effect that healthiness did not correlate with trajectory angle significantly for 12 participants even at the last timestep.




Figure 5.9 Correlations between participants' average MD in the filler trials and the earliest significant timesteps of the two attributes.

5.3.3 Discussion of Study 2

The results of Study 2 were polarizing in terms of the two groups of findings we attempted to replicate. We did not replicate the associations between trajectory parameters and some trials characteristics used in earlier studies (Gillebaart et al., 2015; Ha et al., 2016; Stillman et al., 2017) - no difference was found between tradeoff trials with either self-control success or self-control failure, and between tradeoff trials and trials containing two healthy foods. The former effect was inconsistent previously as well (Gillebaart et al., 2015; Ha et al., 2016), and the

5.3 Study 2: (conceptual) replication of previous findings

later effect was very small in the original study (Stillman et al., 2017). Also considering the findings in Study 1, we suspect that trajectory parameters are more related to conflict measures based on individual attribute ratings and decision weights (i.e., utility difference, utility of stronger option). When trials are categorized more coarsely as in the previous studies, perceived utility difference may differ between trial types in unpredictable way and lead to the observed differences in trajectory parameters. Finally, as those studies are largely exploratory, without correction for multiple comparisons, the large number of mouse-tracking parameters combined with many ways to categorize trials may lead to some false positives.

Given the overall very small effects between trajectory parameters and several conflict measures, it was surprising that Sullivan et al. (2015)'s results of testing an even more subtle cognitive mechanism of self-control with a more novel analytic method could be almost perfectly replicated (except for the results based on proportion curves). Also the effect sizes were consistently large with correlations between key variables reaching $r = 0.6$, and Cohen's d_z between 0.6 and 1 for the main hypothesis tests. The analyses on the filler trials also seem to suggest that people's natural movement tendencies contribute to the estimation of processing latency, so at least these estimates cannot be interpreted in absolute. In the next study, we reanalyzed Sullivan et al. (2015)'s data and conducted a model-based simulations to find out whether the highly robust results could be interpreted as evidence for mechanism of self-control based on processing latency difference.

5.4 Study 3: re-analyses of Sullivan et al. (2015) and simulations

5.4.1 Re-analyses of Sullivan et al. (2015)'s data

Sullivan et al. (2015)'s method and results can be better understood by examining closely the interrelationship between attributes, choice, and trajectory angle (see a visual explanation in Figure 5.10a). The correlations between the attributes and trajectory angle (path x) may exist because they are both correlated with *choice* (path y and z in Figure 5.10b). On the one hand, both attributes influence choice (path y), and tastiness is usually weighed more than healthiness (e.g., Sullivan et al., 2015; Hare et al., 2009). On the other hand, the relationship between trajectory angle and choice is temporarily constrained (path z): the distributions of the angle for left and right choices overlap strongly from timestep 1 till around timestep 50, then start to diverge, until they approach the Bernoulli distribution of final choices (Figure 5.10c). The resulting temporal paths of the correlations between choice and angle must be an S-curve (Figure 5.10d, grey curve), as long as trajectories are curved towards the midline. Assuming that the attribute-angle correlations are mainly mediated by choice, even if a cognitive mechanism unrelated to processing speed causes the decision weight difference, these correlations

will unfold in similar S-curves, one consistently above the other (Figure 5.10d, red and blue curves), and thus an alleged processing speed difference would appear.

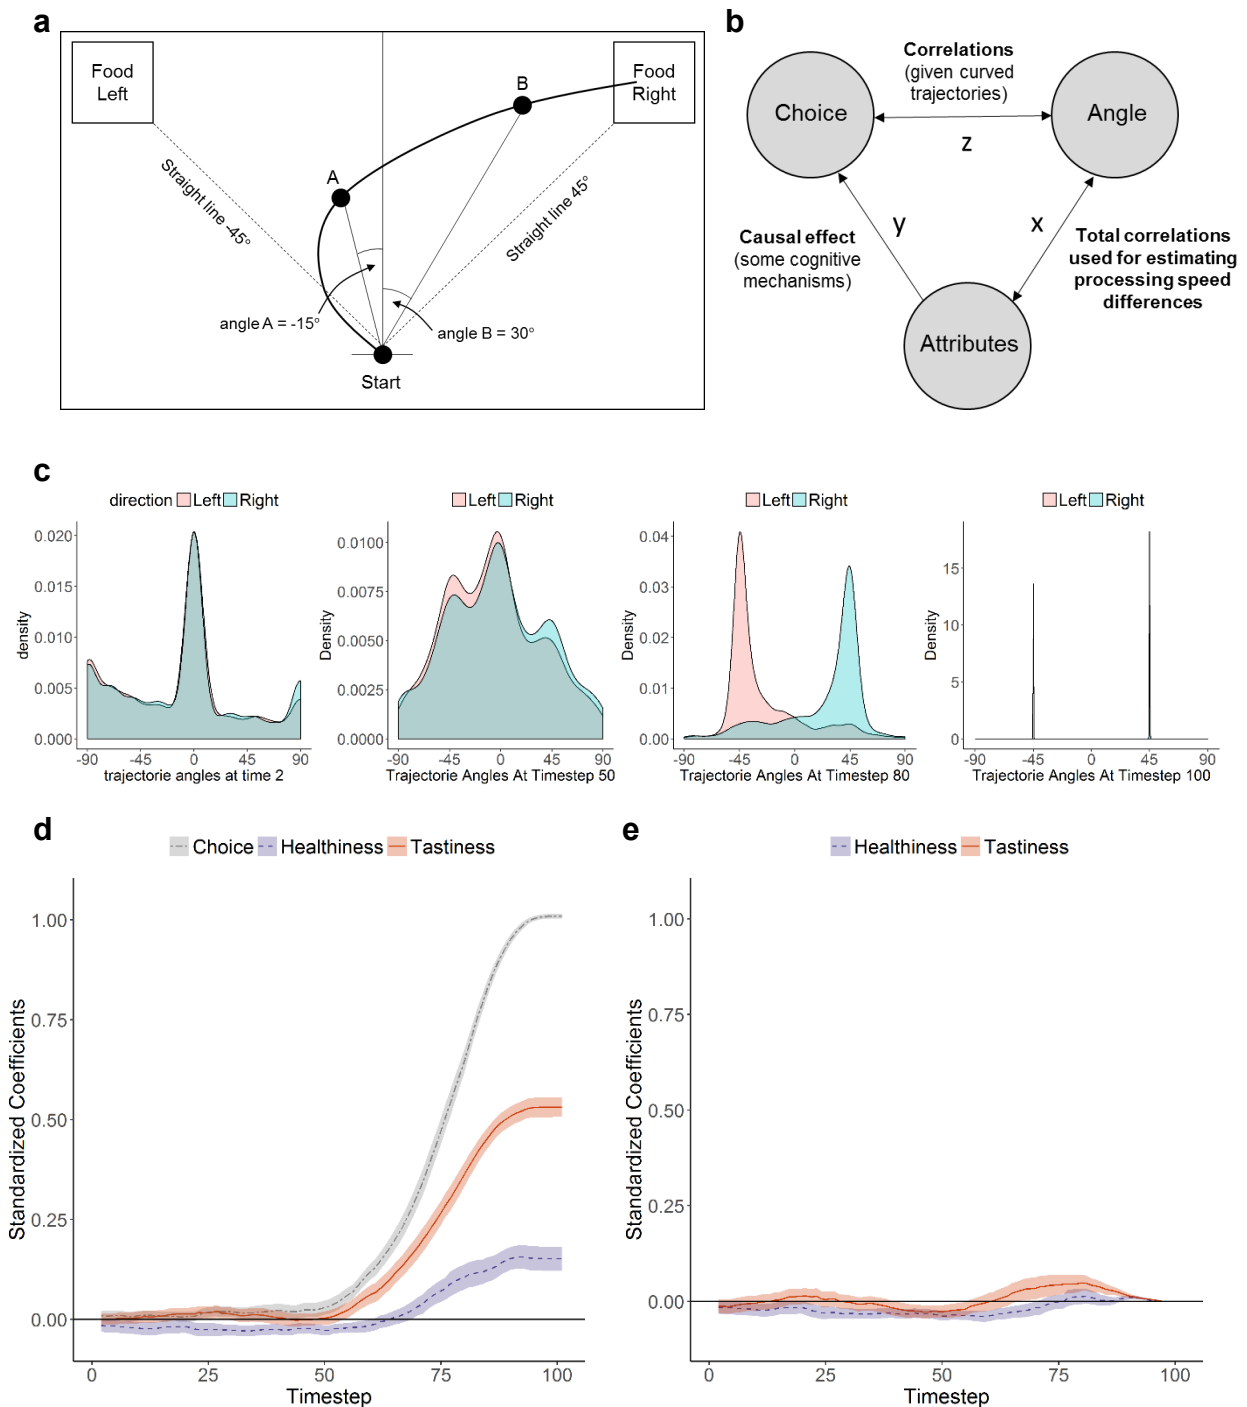


Figure 5.10 (a) Visual explanation of the variable *angle*; (b) *Choice* potentially mediating the correlations between *attributes* and *angle*; (c) Development of the distributions of angle for choices to the left and right; (d) Regression coefficient paths of the attributes mimicking the S-shaped regression coefficient paths of choice; (e) Regression coefficient paths of the attributes after controlling for choice.

The conceptual analysis above raises the question how much choice accounts for the attribute-angle correlations in the empirical data. This question can be answered by removing the common correlations with choice through a mediation analysis, i.e., to estimate the direct effects of attributes on angle while controlling for choice. This is equivalent to splitting the data based on choice direction or recoding attribute differences and angle into absolute values before analyses. If there are sizable effects unmediated by choice, even when directional information is removed, the magnitudes of attribute differences should correlate with the magnitude of angle. Results indicated that the total effects were attenuated by a factor of 11, and the remaining temporal patterns became noisier (Figure 5.10e). Although small significant correlations remained between timesteps 65 and 85 (all $\Delta R^2 < 0.021$), the data pattern does not allow one to confidently infer any processing speed difference.

In the next section, we report a simulation study to confirm the problem identified by the mediation analysis that choice accounts for the majority of the correlations between attributes and trajectory angle. The simulation also helps to address the question how this problem influences the ability of Sullivan et al. (2015)'s method to distinguish cognitive mechanisms.

5.4.2 Simulation study method

Models and assumptions

A simple drift diffusion model was used to model how people make decisions in the food-choice task (see Sullivan, Hutcherson, Harris, & Rangel, 2019). It assumes that people accumulate a value signal that measures the relative preference to one food item over the other, until the signal exceeds a threshold and the decision is made (see Figure 5.11a). The drift rate of the value signal at each accumulation step is partially determined by the attribute differences in tastiness (Δ_T) and healthiness (Δ_H), but only after the attributes are processed and integrated into a decision-making circuit (i.e., after t_T and t_H respectively), defined by the equations in Figure 5.11a. Besides a possible latency difference, a slope difference between S_H and S_T could provide an alternative mechanism that shifts the relative weights of healthiness and tastiness on choice.

For the relationship between decision-making and motor-control, three different assumptions were examined. In the *deterministic* scenario, the model assumes that people move upwards in straight lines (0°) before decisions are made, possibly as a way to reduce distances to both stimuli, and then move towards the chosen stimuli in straight lines (Sullivan et al., 2019; see Figure 5.11b). It is also assumed that the movement speed is a constant, so that the points at which trajectory turns (y_0) is completely determined by decision time. Finally, we followed Sullivan and colleagues (2019) to model angles after decisions to be either 45° or -45° . This might sound odd because even when people move straightly to the target, the angle

will gradually change from 0° to $\pm 45^\circ$. However, it can be observed in empirical data that people often stop at the targets for quite some normalized timesteps, so fixing angles at $\pm 45^\circ$ can be seen as a reasonable simplification. Despite its simplicity, the deterministic scenario could clearly reproduce the classical finding in mouse-tracking paradigms that more difficult decisions lead to larger curvatures (e.g., Dshemuchadse, Scherbaum, & Goschke, 2013; Freeman & Ambady, 2009; McKinstry et al., 2008; Spivey, Grosjean, & Knoblich, 2005; Stillman, Medvedev, & Ferguson, 2017). This is because when integrated attribute differences are small, decisions take longer to resolve and lead to larger y_0 . However, it should be noted that the relationship between attribute differences and y_0 is overly deterministic. Given an appropriate level of drift rate noise (σ)²², attribute differences as weighted by slopes could explain 16% of the variance in y_0 , while the marginal- R^2 estimated from empirical data was only 1.4% (Sullivan et al., 2015).

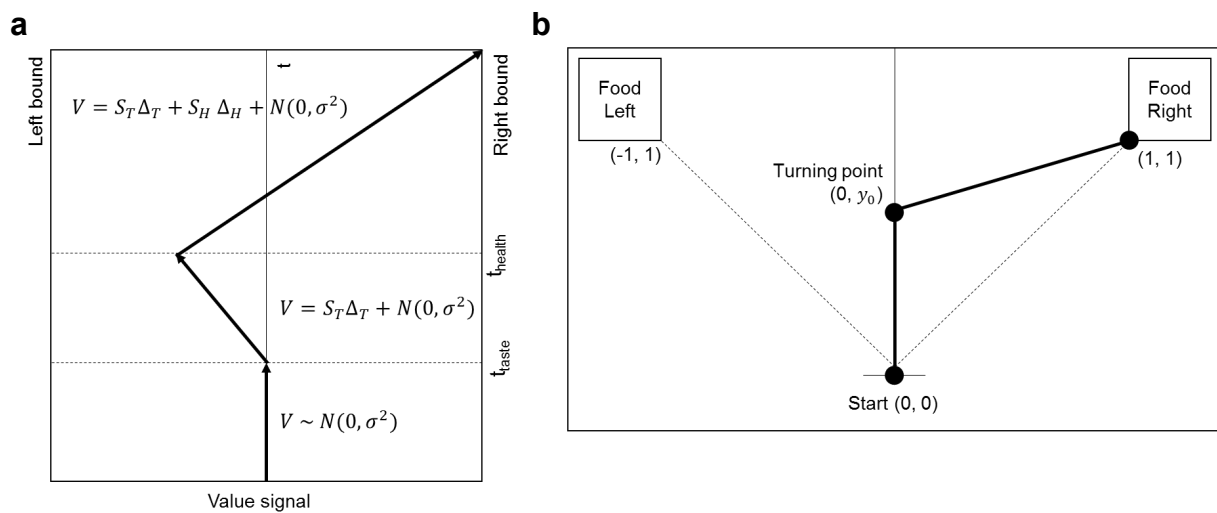


Figure 5.11 (a) A drift diffusion model of the decision-making process; (b) A simple computational model of movement trajectory in food-choice task, where people move upwards while making decisions, and then move directly to the target;

In the *random* scenario, we detached the decision-making process from motor movements. The same drift diffusion model produced choices and decision time, but the turning point y_0 was not influenced by decision time but randomly sampled from a beta distribution bounded between 0 and 1, with a mean of 0.5. Thus, any correlations between attributes and trajectory angle in this scenario are attributed only to the indirect path, i.e., their common correlations with choice.

²² Sigma σ was set to 0.035 so that attribute differences could account for about the same amount of variance in choice as with empirical data (marginal- $R^2 = 0.52$ in Sullivan et al., 2015).

Unlike the two extreme scenarios above, in the third and *realistic* scenario, we calibrated the model with Sullivan et al. (2015)'s data to resemble real mouse-tracking tasks. As with the previous scenarios, accumulation noise σ was set at 0.035, so that about half of the variance in choice was accounted by tastiness and healthiness differences (marginal- $R^2 = 52\%$ in Sullivan et al., 2015). However, the correlation between decision time and y_0 was attenuated by allowing trial-to-trial variations in movement speed v_1 during the upward movement. Based on empirical data, movement speed v_2 after decision-making was also assumed to vary and on average much faster than v_1 . More specifically, both v_1 and v_2 were drawn from beta distributions, while the distribution for v_2 had a larger mean (about 2.5 times). The exact parameter values for the two beta distributions were calibrated to ensure that y_0 or the maximum deviation of the trajectory would correlate weakly with the integrated and weighted attribute difference ($|w_H \Delta_H + w_T \Delta_T|$), approximating a marginal- R^2 of 1.4% in the empirical data (Sullivan et al., 2015). Finally, we relaxed the strong assumption of fixing angles after decision-making at $\pm 45^\circ$ by allowing 50% of the trajectory angles to change gradually after the turning.

Conditions for different cognitive mechanisms

In each of the three scenarios, the latency and slope parameters in the drift diffusion model were varied to create three different conditions of what cognitive mechanisms are at work: (1) equal latency and slope for healthiness and tastiness ($t_H = t_T = 250$; $S_H = S_T = 0.01$); (2) equal slope but larger latency for healthiness ($t_H = 420$, $t_T = 250$; $S_H = S_T = 0.01$); (3) equal latency but smaller slope for healthiness ($t_H = t_T = 250$; $S_H = 0.002$, $S_T = 0.02$)²³. The latter two conditions were critical for testing the competence of the commented method in different scenarios. If the method is competent, it should only detect latency differences when a true latency difference was implemented in the simulation, but not when a slope difference was implemented.

Data analysis

For each modeling scenario and each cognitive condition, 1000 trials were simulated to generate the data for analyses. Three different methods were compared. First, as in Sullivan et al. (2015), raw regression coefficient paths were produced by computing the correlations of trajectory angle to healthiness and tastiness difference at 100 timesteps. A latency difference could be identified by inspecting whether one of the paths arose above zero earlier than the other. Second, the alternative method of proportion paths in Sullivan et al. (2015) were used. As a method to normalize the raw correlation coefficients, the proportion paths estimated at

²³ The exact values for latencies and slopes were calibrated as to reproduce the large empirical difference in decision weights of the two attributes at choice level (e.g., $w_T = 1.17$ and $w_H = 0.33$ in Sullivan et al., 2015; also see Hare et al., 2009).

which timesteps the coefficients exceed certain proportion (from 0.1 to 0.9) of their corresponding endmost values (usually the maximums). The shapes of the proportion paths would indicate which one exceeded the proportion of zero first. Third, the mediation modeling in our re-analyses was used. Regression coefficient paths were again produced, but under the condition where variable choice was controlled in the regression models.

5.4.3 Simulation results

The deterministic scenario

In the deterministic scenario, results indicated that different mechanisms could often be differentiated by visually inspecting the raw regression coefficient paths (Figure 5.12a-c). When the slope difference was large (Figure 5.12c), it could come close to mimic a true latency difference in raw regression coefficient paths, but proportion paths would still sharply separate the two mechanisms (Figure 5.12e & 5.12f)²⁴. In addition, it is crucial to note that when the decision-making process truly influence movement trajectories, our method of controlling for choice direction preserved the critical features in the regression coefficient paths that could be used to differentiate the three conditions²⁵ (Figure 5.12g - 5.12i). Although the coefficients dropped to 0 when choice direction completely overlapped angle at the end, in a sizeable window the patterns were clean albeit attenuated.

Even when a latency difference can be detected in the deterministic scenario, one should realize that the processing speed or latency estimates are relative rather than absolute. As y_0 is determined by both decision time and movement speed, for people with faster movement speed, the absolute estimates for both tastiness and healthiness will increase, even though processing latencies for the two attributes remain the same.

²⁴ Following Sullivan et al. (2015), timestep in the plots of proportion curves is represented by the y-axis rather than the x-axis.

²⁵ Readers might be puzzled by the negative coefficients observed in the random value accumulation stage (before t_T and t_H) when choice direction was controlled. This can be more easily understood if one considers the effect of controlling for choice as analyzing trials with choices to the left and the right separately. When directional information was removed, the distributions of absolute attribute difference scores ($|\Delta_H|$ and $|\Delta_T|$) were severely right-skewed (i.e., a much larger sample for differences of 0 than differences of 4), as the raw attribute ratings were drawn from uniform distributions. Therefore, just by random accumulation, turning points were more likely to be reached in trials with smaller attribute difference due to the sample size differences.

5.4 Study 3: re-analyses of Sullivan et al. (2015) and simulations

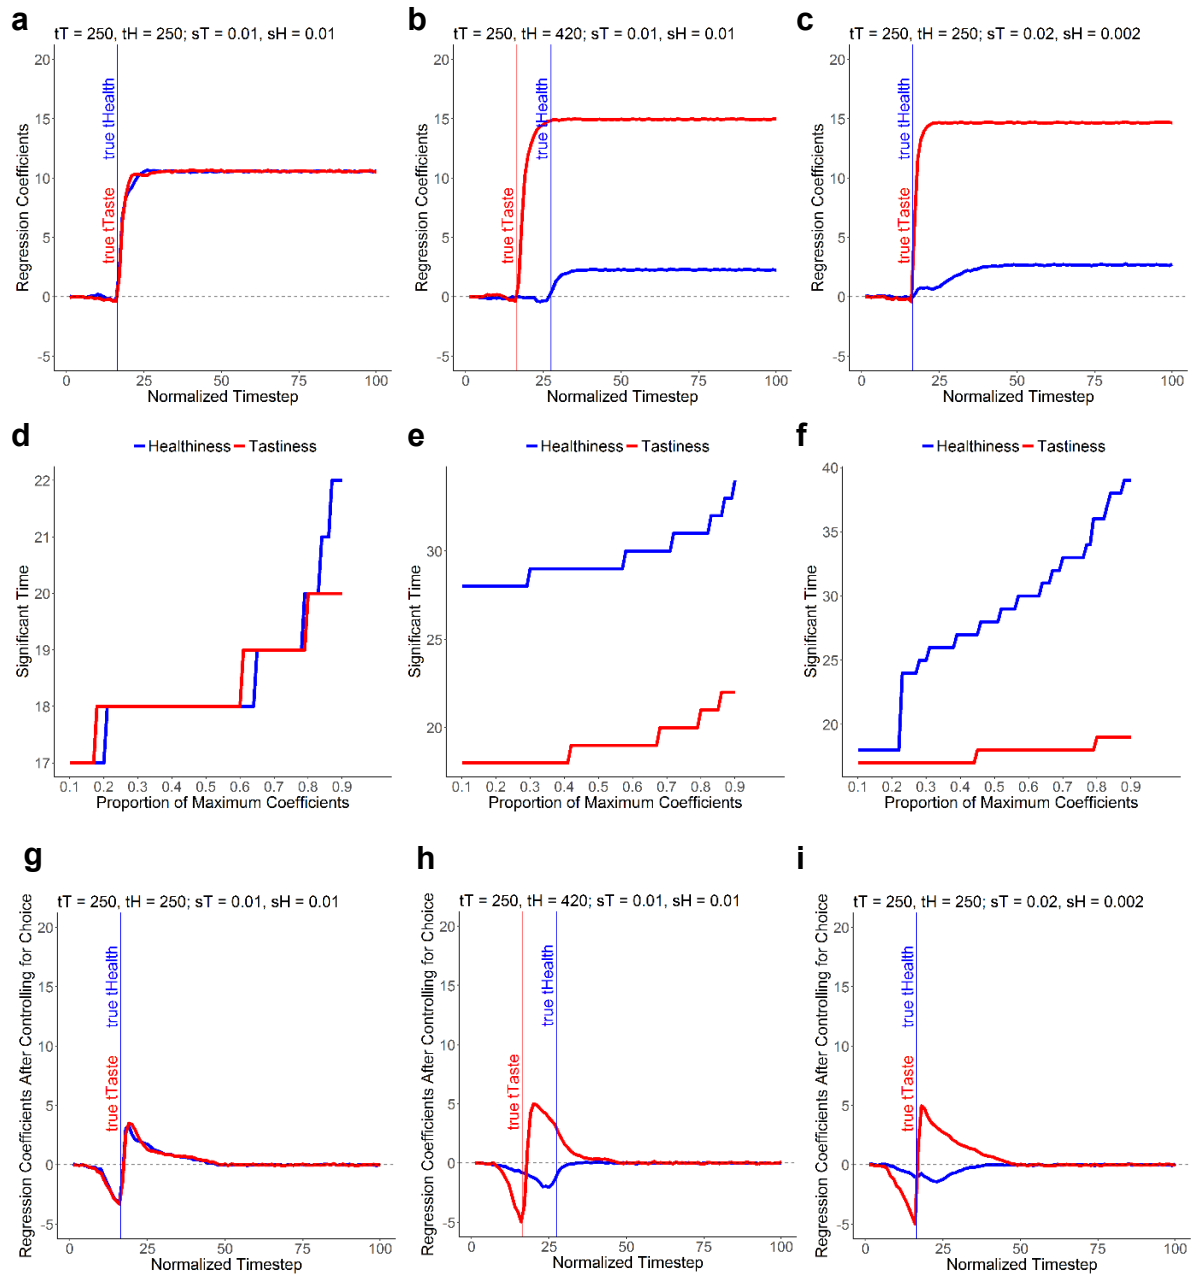


Figure 5.12 Simulations results in the deterministic scenario (Top: results based on regression coefficient paths for three conditions – (a) equal latency and equal slope , (b) latency difference and equal slope, and (c) equal latency and slope difference; Middle: Results of the same conditions when using proportion paths; Bottom: Results when controlling for choice.

The random scenario

Results in the random scenario showed that without controlling for choice, similar S-shaped paths for correlation coefficients were obtained in the same three conditions (Figure 5.13a - 5.13c). The commented method was clearly problematic here because even when a slope difference was the true mechanism, the pattern of the raw regression coefficient paths (Figure 5.13b & 5.13c) showed a latency difference, and its proportion paths were also hardly distinguishable from the ones generated by a true latency difference (Figure 5.13e & 5.13f). The results confirmed our conceptual analysis that the exact shapes of coefficient paths depended mainly on the decision weight difference between healthiness and tastiness at the choice level, regardless of whether a latency or a slope difference was the underlying mechanism.

Critically, when choice was controlled for, the previous regression coefficient paths completely collapsed, with only negligible random variations remained around timestep 50 (Figure 5.13g - 5.13i). These results show that the method of controlling for choice direction can be used to distinguish the condition where decision-making affected movement trajectory from the condition where decision-making and motor-control were unrelated. In other words, the mediation analysis could provide an estimate of the relative contribution of the direct causal influences of attributes on trajectory angle and the choice-mediated attribute-angle correlations. When examining the empirical patterns before and after controlling for choice direction (Figure 5.10d & 5.10e), it seemed that the empirical situation was closer to the random scenario than the deterministic scenario, which supported our mediation analysis that the attribute-angle correlations in the empirical data were mostly mediated by choice. However, note that unlike the proportion paths estimated in Sullivan et al. (2015), these paths in all conditions in the random scenario did not show clear latency difference (the paths overlapped greatly). This indicated the needs to evaluate the commented method in a more realistic scenario.

5.4 Study 3: re-analyses of Sullivan et al. (2015) and simulations

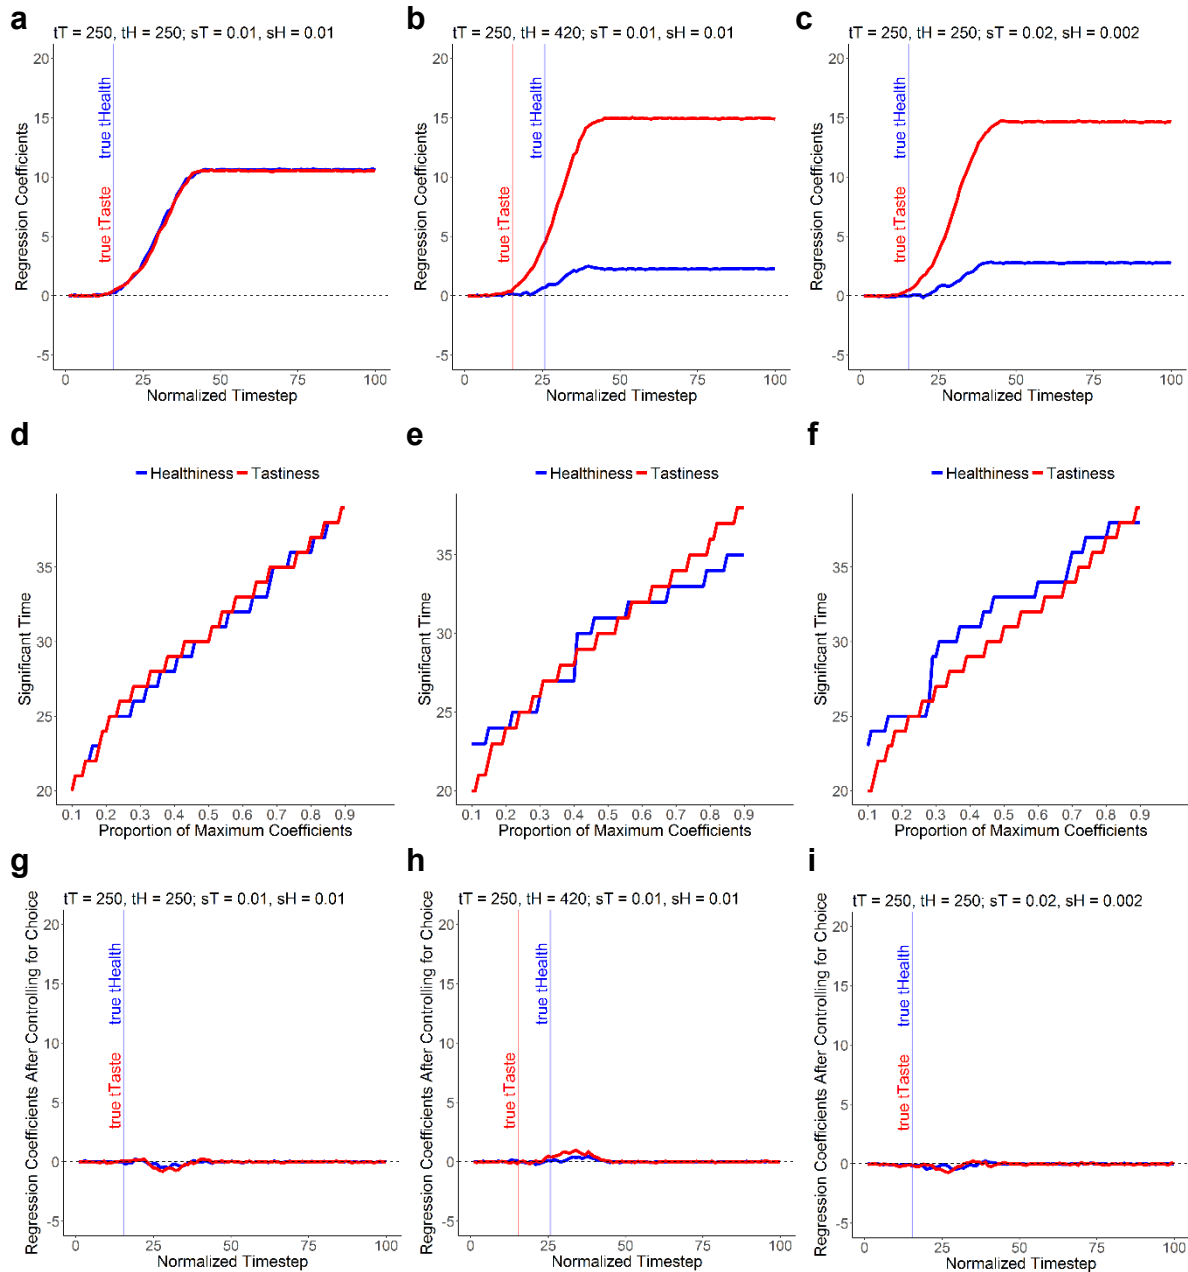


Figure 5.13 Simulations results in the random scenario (Top: results based on regression coefficient paths for three conditions – (a) equal latency and equal slope , (b) latency difference and equal slope, and (c) equal latency and slope difference; Middle: Results of the same conditions when using proportion paths; Bottom: Results when controlling for choice.

The realistic scenario

With the calibrated model in the realistic scenario, the regression coefficient paths appeared to be smoother and more similar to the empirical observations. For the raw regression coefficient paths (Figure 5.14a - 5.14c), results indicated that an apparent latency difference could be attributed to a true latency difference or a slope difference, and that after controlling for choice direction, the magnitudes of attenuated patterns were in between those in the deterministic and random scenarios (Figure 5.14g - 5.14i). For proportion paths (Figure 5.14d - 5.14f), a clear latency difference could be consistently observed under a true latency difference, but also sometimes (Figure 5.14e) though not always (Figure 5.14f) under a slope difference. Because of this stochastic property, we ran the same simulation for an additional 100 times, from which slope difference appeared as latency difference in both raw regression coefficient paths and proportion paths for 46 times. These findings question the competence of the commented method to differentiate between the two distinct mechanisms in the realistic scenario. The empirical proportion paths in Sullivan et al. (2015) seems at least equally plausible to be drawn from the varying patterns under the slope difference than from the clean and consistent patterns under the latency difference.

Overall, the simulation study strengthened our critics in two ways. First, it confirmed that our re-analysis method of controlling for choice was a sound and informative approach to separate the indirect effects via choice from the more meaningful direct effects of attributes on trajectory angle. Second, it answered the question under what conditions the commented method would be competent – it worked under the noise-free deterministic scenario, but would fail when motor-movement was unrelated to decision-making or when the relationship was very weak, as in the real mouse-tracking data on food-choice (e.g., Gillebaart et al., 2016; Lim et al., 2018; Stillman et al., 2017; Sullivan et al., 2015). Note that the models used in the simulation study were not meant to be the true cognitive and motor-control models underlying mouse-tracking, but were simplified but plausible models that helped to illustrate the limitations of the commented method.

5.4 Study 3: re-analyses of Sullivan et al. (2015) and simulations

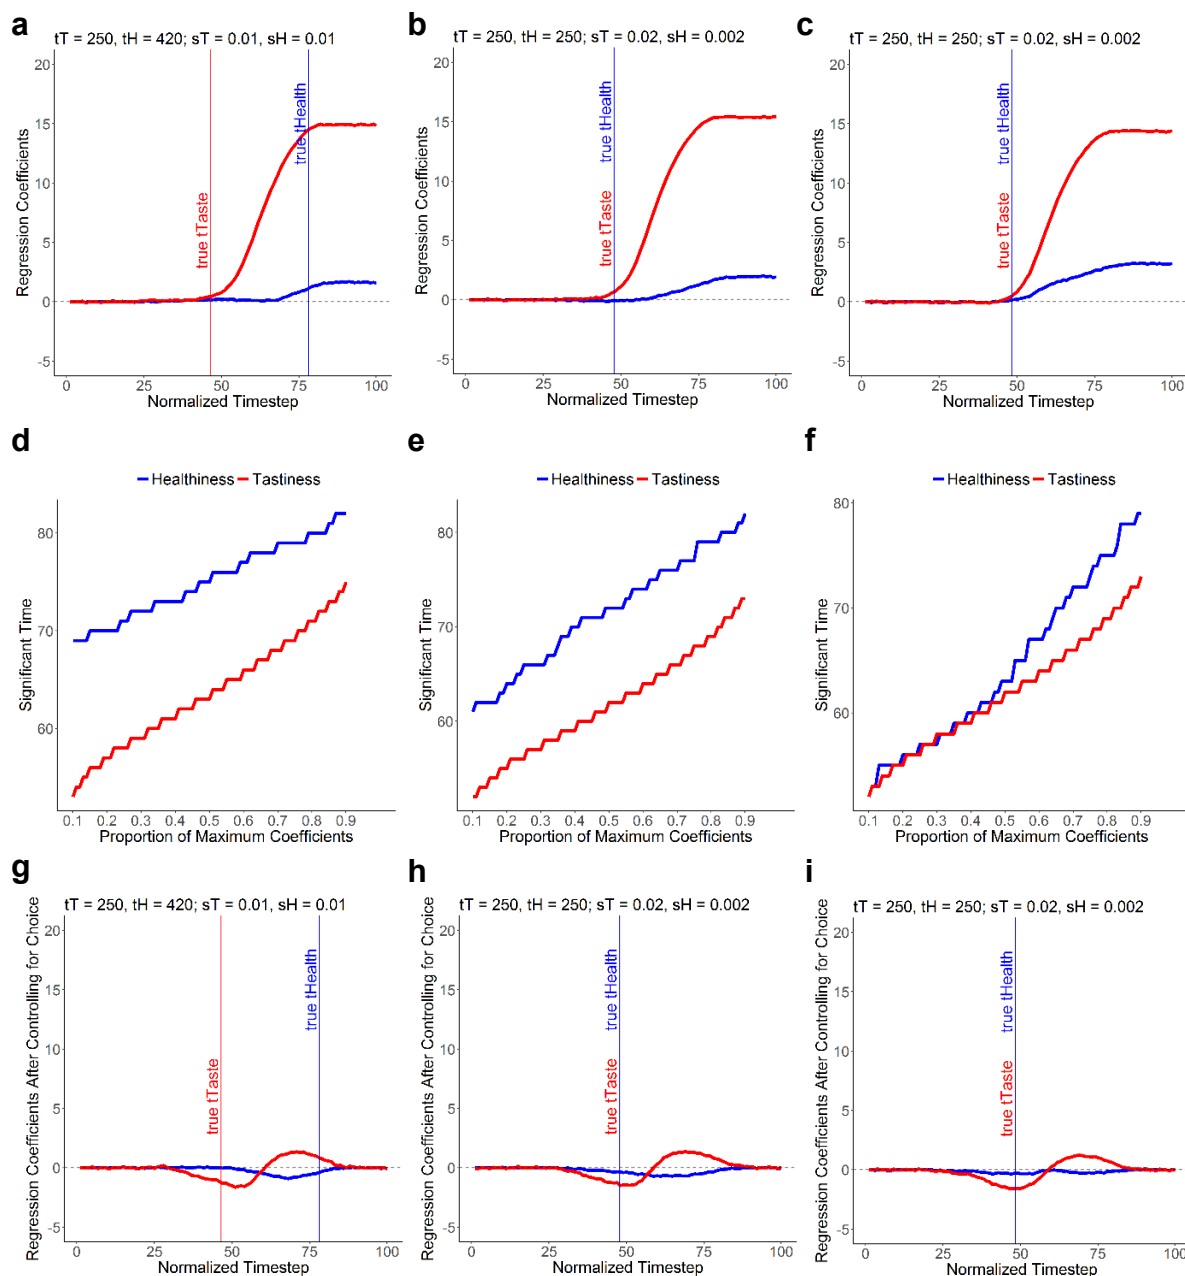


Figure 5.14 Simulations results in the realistic scenario (Top: results based on regression coefficient paths for the two critical conditions: (a) latency difference and equal slope, and (b & c) equal latency and slope difference for two instances; Middle: Results of the same conditions when using proportion paths; Bottom: Results when controlling for choice.

5.4.4 Discussion of Study 3

We have shown with the re-analyses and simulations that using Sullivan et al. (2015)'s method latency differences may appear even when a mechanism unrelated to processing speed underlined dietary self-control. Although the method works given a deterministic version of the moderate assumption of mouse-tracking, it is inadequate to differentiate different mechanisms under the realistic and noisy task environment that characterize the current mouse-tracking paradigm.

5.5 Study 4: food-choice experiment with practice trials and different orientations

Given all the findings in the previous three studies, it seemed to us that the moderate assumption was more plausible than the strong assumption for the mouse-tracking paradigm – there is no direct continuous mapping between decision-making and motor-control, but decision-making does partially influence the distances of upward movements. Yet it remained an interesting question whether the moderate assumption was also more plausible than the weak assumption, which states that trajectory curvatures are sole product of motor-control and that typical difficulty-deviation relationship is fully attributed to extreme direction reversals (i.e., change of mind). Results in Study 1 showed that even when trials with “change of mind” were excluded, small correlations between deviation parameters (e.g., MD, x-flip) and conflict measures (e.g., utility difference, utility of stronger option) still existed, even though the effect sizes were greatly attenuated. This seems to provide some support for the mechanism implied in the moderate assumption. On the other hand, the strong correlation between MD in the filler trials and MD in the food-choice trials might be in favor of a pure motor-control account of movement trajectory, trajectories in the filler trials indeed measure people's natural movements. However, although the filler trials did not involve food choices, decisions regarding on which side of the screen the only food item would appear were still required.

In Study 3, we designed a similar food-choice experiment as with Study 1, but included a neutral practice block to better measure people's natural dragging movements between two points on the screen. Decision-making of any form was not needed for the practice trials. If the upward movements, which result in various curvature sizes, are only shown in food-choice trials (with decision-making) but not in practice trials (without decision-making), the moderate assumption should gain some verisimilitude.

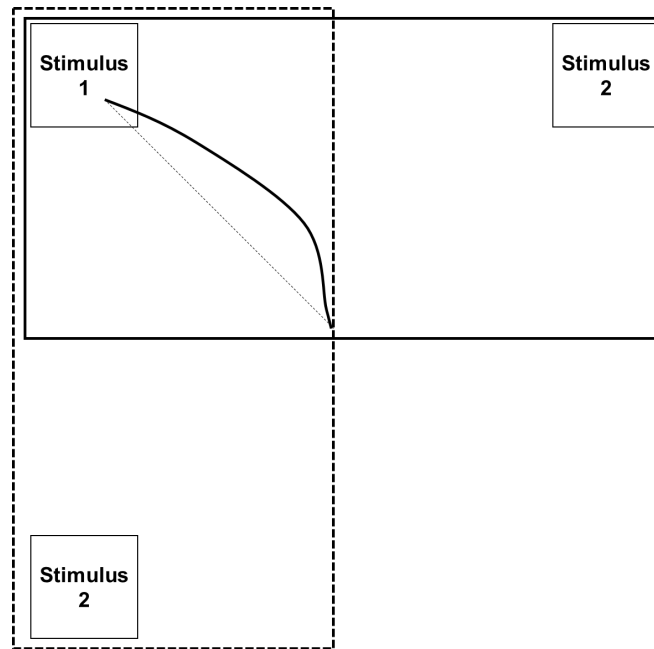


Figure 5.15 For the solid frame (normal spatial orientation), the trajectory is curved towards the midline between stimuli; For the dashed frame (a vertical orientation), the same trajectory is curved away from the midline.

In addition, as another way of testing the moderate assumption against the weak assumption, we designed the food-choice trials with different spatial orientations – participants do not only drag upwards (as in the traditional mouse-tracking paradigm), but also downwards, leftwards, and rightwards. If decision-making is the key for the vertical movements in between two options, then movements should be along the midline between two options in all orientation conditions. Conversely, if the upward movements in the traditional paradigm is mainly due to the motor-control constrain imposed by the particular spatial arrangement, then trajectory patterns could be dramatically different in the conditions with novel orientations. For example, in the traditional paradigm, if up-left and up-right movement trajectories are curved towards the midline due to a motor-control constrain but not decision-making, then in conditions with left-right or right-left orientations trajectories to the top responses should be curved away from the midline (see Figure 5.15).

5.5.1 Method

Participants & design

Forty-three students (20 males and 23 females) from Eindhoven University of Technology participated in the experiment as part of a course fulfillment. The age of the participants ranged between 19 and 25 (mean = 21.3, $SD = 1.36$). Their BMIs were between 17.6 and 31.2 (mean = 22.8, $SD = 2.46$). Of all the participants, four were following a diet, of which one followed a general healthy diet, two followed a gluten-free diet, and one followed a gluten-free and lactose-free diet. Seven participants indicated one or more food-related allergies,

including gluten intolerance, and allergies to shrimps, shellfish, fruits, milk, and pork meat. Four people indicated that they were vegetarians. All participants used their right hand to perform the mouse-tracking task.

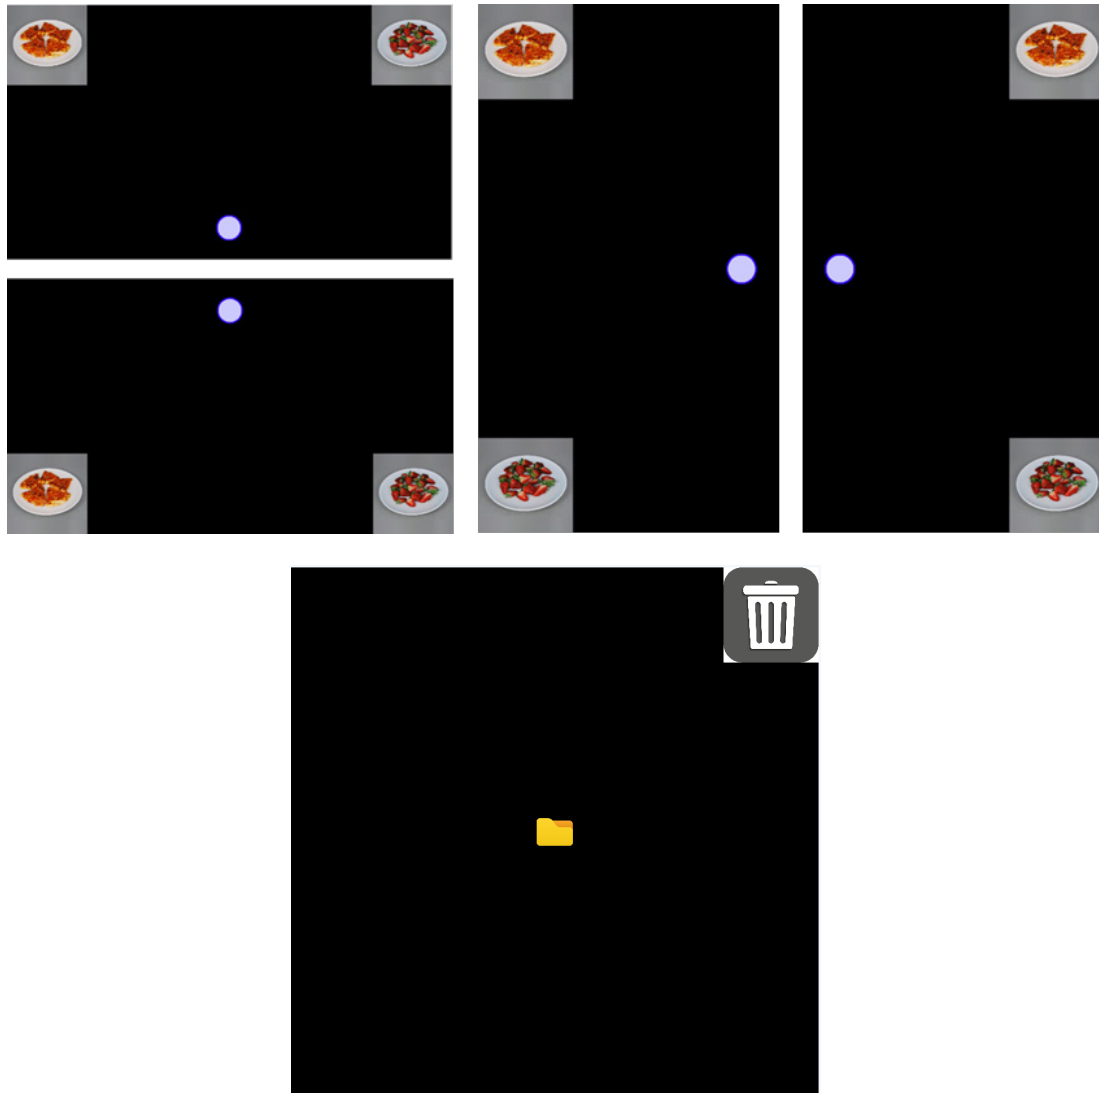


Figure 5.16 Top: Food-choice task with different spatial orientations (the horizontal conditions are scaled smaller for the ease of presentation); Bottom: One of the practice trials of dragging the folder to the recycle bin.

All participants were exposed to four different conditions in a within-subject design. In each condition, the orientation of the food stimuli relative to the starting position of the cursor changed (see Figure 5.16, top) – participants needed to either drag the cursor at the bottom upwards (the traditional *up* condition), at the top downwards (the *down* condition), from the right to the left (the *left* condition), or from the left to the right (the *right* condition). Ninety pairs of food stimuli were repeated twice, giving 180 trials, which were randomly distributed to the four orientation conditions. The order of the conditions was counter-balanced across participants. For the four conditions, all other spatial parameters except for the orientation

were identical (e.g., distance to travel, angles, and stimuli sizes).

Prior to the food-choice task, participants performed 40 practice trials, in which they were required to drag a folder icon at the center of the screen to the 4 corners (see Figure 5.16, bottom). For each direction, 10 trials were repeated consecutively. Unlike the filler trials used in Study 1, no decision had to be made at all. The supposed direction of dragging was explained in text before the start of the first trial for each direction, and later trials simply followed the same direction. The distance and angle of the 4 required movements in the practice were exactly matched to the distance and angle of movements in the corresponding orientation conditions in the food-choice task.

Food stimuli

The same 10 food stimuli were used as in Study 1.

Apparatus and measurements

The same mouse-tracking task and measurements were used as in Study 1. Because control modes did not seem to matter in Study 1, we only included the physical-mouse setup in this study.

Procedure

The procedure was very similar to Study 1. Participants first performed the practice trials, then the four conditions of food-choice trials in a random order, and lastly they rated the food items and completed the same general questionnaire. Before leaving the experiments, they were again asked to eat a food item randomly chosen from all the foods they chose at least once during the task, in order to increase the ecological validity of the experiment.

Data preprocessing

The raw trajectory data were preprocessed in the same way as with Study 1. Note that for trajectories in the non-conventional conditions (down, left, right), the trajectory data were rotated first, so that all spatial parameters computed would have the same meaning in all conditions. For example, a MD of larger than 0 always meant some curvature towards the midline between the response options. Moreover, the same trial-exclusion criteria were used. Firstly, trials in which the participants released the cursor prematurely were removed (6%). Secondly, 7 trials with technical faults, indicated by extremely large x and y-coordinates or negative area under curve (AUC), were removed (0.1%). Thirdly, trials with (log-transformed) reaction times larger than three standard deviations from the grand mean were removed (0.5%). Fourthly, trials in which the participants changed their minds more than twice were removed (0.8%). The final dataset for analysis consists of 7875 food-choice trials and 1585 filler trials from 43 participants.

5.5.2 Results²⁶

Comparison of trajectories in different orientation conditions and in practice

Figure 5.17 visualizes the averaged trajectories per person per condition (including practice). The pattern makes two things crystal clear. First, when there was no decision to be made in practice trials, participants' trajectories were very close to the straight lines connecting the origin and destinations (pink lines in Figure 5.17), and the deviations were much smaller than those in the other conditions. Second, for food-choice task with all orientations, trajectories tended to curve towards the corresponding midlines, showing upward movements at the early stage of a trial response.




Figure 5.17 Average trajectories per person per condition in the task space.

The same conclusions can be drawn by comparing the distributions of trajectory parameters in different conditions (see Figure 5.18). Compared with food-choice trials, practice trials produced much smaller MD, fewer x-flips, shorter response time, and faster maximum velocity. The fact that movements were close to straight lines was more evident in density plot for MD, in which the distribution of MD for practice centered exactly around zero. On the other hand, the distributions of parameters largely overlapped for the four orientation conditions. Without statistical tests, it was safe to say that if any differences existed between those conditions at all, the effect sizes would be even smaller than the small difference between the

²⁶ Data from Study 4 would also replicate the results in Study 2 (e.g., findings in Sullivan et al., 2015), but for brevity, I do not report them in this chapter.

touch-screen condition and the mouse condition used in Study 1.

Lastly, contrary to the results of the filler trials in Study 2, participants' average MD in practice trials did not correlate with their average MD in the food-choice trials for all four different spatial orientations (all $ps > .05$), nor was any correlation found for x-flip between the two types of trials (all $ps > .16$). However, participants did behave similarly in practice trials and in food-choice trials, in terms of how fast or slow they completed the trials. For all movement directions considered, weak to moderate correlations were observed for response time (Pearson's r between 0.25 and 0.48) and for maximum velocity (Pearson's r between 0.26 and 0.46).

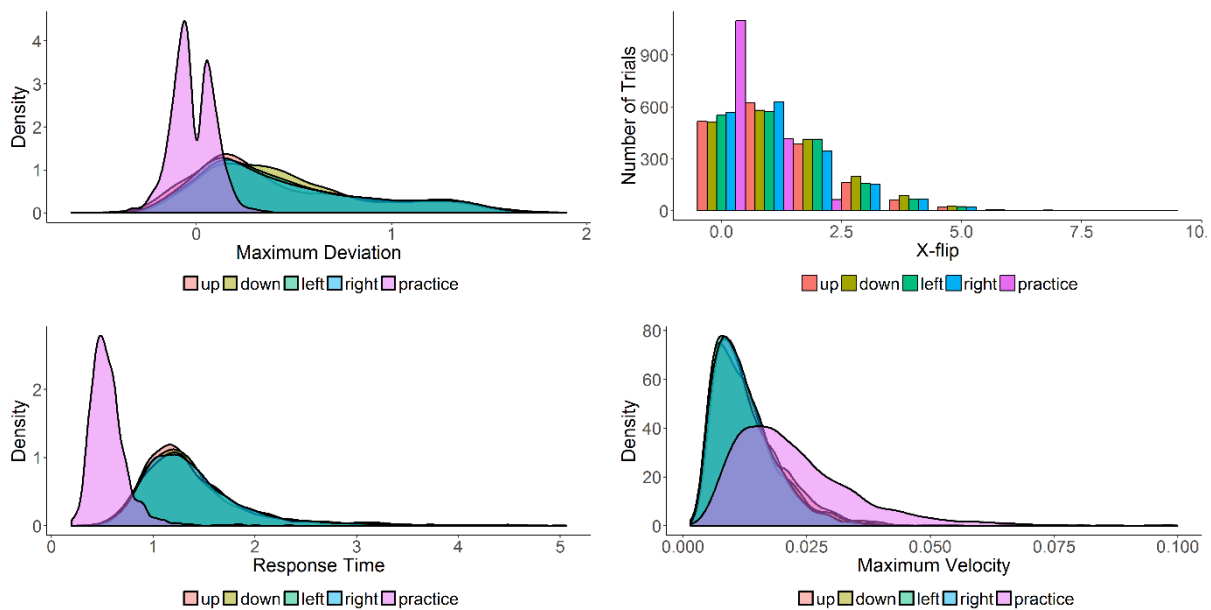


Figure 5.18 Density plots for the distributions of trajectory parameters in different conditions (histogram for x-flip).

Orientation and the sensitivity of trajectory parameters to conflict strength

We also examined whether spatial orientation affected the sensitivity of trajectory parameters to attribute-based measures of conflict strength. For this analysis, we used the model that showed the strongest effects in Study 1 – random-intercept model that predicted parameter values (MD, x-flip, response time, and maximum velocity) using utility difference, utility of stronger option, and their interaction. Table 5.4 compares the overall model fit (marginal R^2) in the four different conditions, and with or without the removal of trials that participants changed their minds (i.e., 19.2% of all trials). The overall pattern of the results did not change much across the orientation conditions, and the effect sizes generally replicated the results in Study 1. When trials became more difficult due to smaller utility difference and/or smaller utility of the stronger option, MD, x-flip, and response time all increased. Also consistent with Study 1, when trials with more than one choice commitment were removed, all parameters'

effect sizes were attenuated, among which MD was affected the most and response time the least. Finally, one might be tempted to conclude that the *right* condition produced the relatively largest sensitivity of parameters to decision conflict, but this result should be replicated with new data first.

Table 5.4 Model fit (marginal R^2) in different orientation conditions (with or without “change of mind (CoM)” trials).

	Up	Down	Left	Right
<i>MD</i>				
With CoM	0.056	0.036	0.059	0.101
Without CoM	0.003	0.004	0.001	0.012
<i>x-flip</i>				
With CoM	0.060	0.055	0.047	0.064
Without CoM	0.012	0.022	0.022	0.038
<i>Response time</i>				
With CoM	0.094	0.077	0.066	0.102
Without CoM	0.074	0.061	0.041	0.064
<i>Maximum velocity</i>				
With CoM	0.020	0.014	0.026	0.019
Without CoM	0.004	0.007	0.015	0.014

Note: The random-intercept models included each parameter as the outcome variable, and utility difference, utility of stronger option, and their interaction as predictors.

5.5.3 Discussion of Study 4

Results of Study 4 made two things clear. First, unlike trajectories in the filler trials in Study 2, trajectories in the practice trials were very close to straight lines, and their deviations did not correlate with the deviations of trajectories in food-choice trials. Second, trajectories in all spatial orientations curved towards the midline, suggesting some vertical movement along the midline at the early phases of responses, which were potentially influenced by the decision-making process. Both findings support the moderate assumption of the mouse-tracking paradigm that people move in between two response options while the decision-making process is still ongoing.

5.6 General discussion

In this extensive chapter of evaluating the mouse-tracking technique on self-control research, our investigation was a quite unexpected journey. Initially after reading the work by Sullivan et al. (2015) and Stillman et al. (2017), we were excited about the possibility of using the technique to uncover cognitive processes underlying health-related decision-making, particularly in real-world applications. We hoped to find the differences in trajectories between the trials with what we called “the two kinds of difficulty”: the trials with strong tradeoffs as in typical self-control dilemmas, and the trials with two options that were simply similar to each other. Our hypothesis was clearly rejected, and later we uncovered the methodological limitations with Sullivan et al. (2015)’s method on using mouse-tracking to reveal cognitive mechanisms of self-control. These results made us very skeptical about the whole mouse-tracking paradigm, and we suspected that the weak assumption was all there was: trajectory curvatures were completely under the influence of motor-control. Finally, Study 4 again proved us wrong – the moderate assumption turned out to be the most plausible.

5.6.1 Trajectory parameters as a function of different decision scenarios

A major contribution of this chapter is the systematic analysis on how common trajectory parameters are influenced by different decision scenarios. As discussed earlier, previous studies had looked at the same issue, for example, by comparing self-control success trials with self-control failure trials (Gillebaart et al., 2015; Ha et al., 2016), or by comparing trials involving self-control and trials with two health food options (Stillman et al., 2017). However, the focus of these studies was pretty scattered, and the hypotheses did not seem to follow clearly from theories. In addition, the trial types used for their comparisons were based on rough manipulations or categorizations (e.g., stimuli’s typical categories, success versus failure), rather than perceived characteristics of the trials. Our results suggest that trajectory parameters are more sensitive to different decision scenarios subjectively perceived by individuals (based on both perceived attribute values and subjective decision weights) than to the collectively defined food categories (healthy or unhealthy)²⁷. Unlike some other mouse-tracking applications, where stimuli characteristics and differences between stimuli could be univocal for everyone, food-choices are likely to be driven by personal perceptions of food. Therefore, to reason whether and how trajectories parameters might differ between decision scenarios defined at higher-level, the underlying subjective perceptions have to be taken into account.

²⁷ Although not reported in detail, we also examined the correlations of trajectory parameters to attribute differences with population-level decision weights or without any decision weights. In both cases, the correlations were weaker when compared with the analyses with individual decision weights.

It was found that both the utility difference between two options and the utility of the stronger option influenced trajectory parameters. This finding is quite intuitive: decisions should become easier and take less time when there is a large difference between two options in terms of their values, and also when one of the options (or both options) is very satisfying. For example, choosing between two very good job offers is probably easier than choosing between two that barely meet one's aspirations. This intuitive finding, if replicated in other decision domains as well, may have significant implications for some sequential sampling models of decision-making. This is because different variants of sequential sampling models assume different types of value signals or preference accumulators, which may have different predictions regarding the influences of utility difference and the utility of the stronger option. The drift diffusion model used in Study 3 and the multialternative decision field theory (Roe et al., 2001) use contrasts between options as value signals, so they seem to naturally predict an effect of utility difference, but not obviously an effect of the utility of the stronger option. Conversely, some variants, such as the associative accumulation model (Bhatia, 2013), employ independent value accumulators for individual options without any contrasting mechanism, so that decision time in these models seems to mainly depend on the value of the best option. Lastly, the multiattribute decision field theory proposed by Diederich includes an additional mechanism to distinguish approach and avoidance conflicts – deciding between two equally good options is easier and faster than deciding between two equally bad options (Diederich, 2003a; 2003b). Simulation studies can be done in the future to understand what these models predict and how the results compare to our empirical finding.

We did not find any differences in trajectory parameters between trials with the two types of difficulty, namely two options with a strong health-taste tradeoff, or two options that were equal in all aspects. Under the assumption that people shift their attentions between healthiness and tastiness, as in decision field theories (Roe et al., 2001), the null-result may cast some doubts on the strong assumption of the mouse-tracking paradigm, which states that the continuous attractions to the opposing options at the cognitive level are mapped directly to motor movements (Freeman et al., 2008; Spivey et al., 2005). However, if attribute values are integrated into one value signal, as in the drift diffusion model used in Study 3, then even the strong assumption may not predict the effect we wished to find. Given the theoretical and practical significance of self-control dilemma, as well as the phenomenological feelings of being pulled towards two extremes in situations with strong tradeoffs, it remains an important open question how such conflicts are resolved at the cognitive level, and whether any traces can be revealed using indirect measures.

Throughout the chapter, we have implicitly taken the position of calibrating the trajectory parameters as novel measures to some more refined measures of decision conflict, such as utility difference computed based on attribute ratings. This position was less of a theoretical commitment, but more of a convenient choice given that it was impossible to assert a ground

truth of conflict strength. For example, if one observes a much deviated trajectory and a long response time, but the attribute ratings of the decision-maker suggest an easy trial, what should one believe? Attribute ratings, as with any self-report, can be biased for various reasons and may not reflect the values processed in the subsequent decision moments. Some researchers have instead advocated the use of decision time as a measure of conflict strength (Diederich, 2003b). Moreover, it may be argued that decision conflicts are primarily subjective feelings and direct self-report measures are closer to the truth (for an example, see Gil-lebaart et al., 2015), but again how much such feelings are accessible to introspection is debatable. Finally, conflict strength might be reflected in actual choices as well, in a way that more choice reversals may occur if the same difficult decision is made repeatedly. If one's primarily goal is behavior prediction, the number of choice reversals may serve as a criterion variable for other direct or indirect measures. In our experiments, as the same food pairs were only repeated twice, this behavior-based approach was unfortunately unattainable.

5.6.2 Using mouse-tracking data to reveal specific cognitive mechanism of self-control

Through a conceptual replication, re-analyses on Sullivan et al. (2015)'s data, and model-based simulations, it has become clear that the temporal analysis method on mouse-tracking data cannot provide strong evidence to the processing latency difference hypothesis of self-control. Given that choice accounts for most of the attribute angle correlations, their method cannot distinguish different plausible mechanisms underlying self-control in realistic settings. We have shown in the simulation study that a slope difference as in the drift diffusion model is also a plausible cause of the weight difference at the choice-level, and it can appear as a latency difference using Sullivan et al. (2015)'s method. Moreover, according to decision field theories (e.g., Bhatia, 2013, Roe et al., 2001), the fact that people weight tastiness more than healthiness can also be explained by a higher sampling probability of tastiness than healthiness in the preference accumulation process.

Note that our criticism does not apply to previous works on the Simon task using the mouse-tracking paradigm, where a similar method of temporal analyses was used (Scherbaum et al., 2010; see also Scherbaum & Kieslich, 2018)²⁸. In fact, choice direction was used as a predictor in multiple regressions but the coefficients of other predictors were not attenuated. This was because, unlike attribute differences, those predictors of interests (e.g., location, congruence of previous trial) were orthogonal or independent to choice direction by experimental design.

²⁸ It should be noted that in those studies angle was defined as the angle to the y-axis for each difference vector between two time steps. Thus, it measured the instantaneous movement angle rather than the angle relative to the starting position as in Sullivan et al. (2015).

5.6.3 Theoretical assumptions underlying the mouse-tracking paradigm

The results of our experiments contribute to the growing empirical record of the relationship between trajectory deviations and decision difficulty (e.g., Dshemuchadse et al., 2013; Freeman et al., 2008; Koop, 2013; McKinstry et al., 2008; Spivey et al., 2005). In the introduction, we mentioned three possible theoretical assumptions that can account for this general effect. Although our studies cannot provide a definite answer to the verisimilitude of these theoretical assumptions, some insights are gained. Results of Study 4 imply that some form of decision-making is a necessity for trajectory curvatures (otherwise people move in straight lines as in the practice trials), and it is plausible as in the moderate assumption that while still making decisions people tend to move upwards in order to reduce the distances to both response options (as a strategy to improve response efficiency). Thus, the deviation-difficulty relationship is partially due to the (imperfect) mapping of decision duration and the distance of the upward movement. Apparently, trials in which people actually reverse their choices would result in larger deviations and they contribute to the effect as well, but these trials only count for a small portion of all food-choices (under 20% in our studies).

It was more difficult to find any evidence for the strong assumption that momentary preferential status or shift at the cognitive level are reflected in the movement trajectories, given the use of the common parameters. It remains interesting to see whether the direction of the upward movements during decision-making is affected by the relative attractiveness of the options, or any directional biases are purely motor-control noisy. For example, if an upward movement slightly leans towards the left, does it mean that although the commitment to the left option is not yet made, the movement is already biased because of the option's favorable status in preference accumulation? Testing this hypothesis would require more advanced method to disaggregate movement trajectories to distinct stages or segments (see e.g., Calzagni et al., 2017).

In general, when applying the mouse-tracking technique to self-control or any other research areas, researchers should carefully consider the theoretical assumptions they believe to underlie the mouse-tracking paradigm and to use theory-driven predictions whenever possible. This would require one to connect cognitive models of interests (e.g., a self-control mechanism based on processing latency difference) to a model of how the cognitive processes relate to motor-control processes. In this respect, although Sullivan et al. (2015)'s method is limited, their model-based approach should be encouraged. The field should be moving forward faster if researchers move beyond the vague verbal descriptions of how mouse-tracking is supposed to work, and not to be satisfied with merely demonstrating correlations between parameters and coarse task manipulations.

5.6.4 Implications for applying mouse-tracking techniques to digital interventions

An encouraging finding in this chapter is that the validity of the mouse-tracking paradigm does not seem to depend on control method or spatial orientation of the display. Thus, mouse-tracking techniques can be used on mobile devices, such as smartphones, and in portrait mode. There is a tentative indication that the unconventional setup of moving the cursor from the left to the right seems to produce the strongest relationship between trajectory parameters and attribute-based conflict strength. However, in general, the correlations between trajectory parameters and attribute-based conflict strength might be too small to be useful for the use cases discussed in the introduction. If one's goal is indeed to predict or measure attribute-based conflict strength (e.g., the utility difference between two food options experienced by a user), then a single trial or a few trials cannot guarantee much higher accuracy in prediction than random guesses. However, because attribute-based measures may not reflect the true underlying conflict strength, future research should also use behavior-based criterion variables in field intervention studies to see if mouse-tracking data can be useful for predicting actual behavior change (e.g., relapses).

For promoting the use of mouse-tracking, it is also important to demonstrate the added value of spatial parameters beyond response time. Although spatial parameters and response time do not correlate strongly, they seem to react in the same way to attribute-based conflict strength, and often response time is more sensitive than e.g., MD or x-flip. If only response time is needed, simpler response formats are no less applicable than the mouse-tracking setup, and they tend to be flexible in terms of interface design requirements in the digital systems. Nonetheless, there is at least one unique source of information that can only be obtained by mouse-tracking – choice reversals (change of mind) within each single trial (e.g., Szaszi et al., 2018). If detecting within-trial choice reversals is valuable for an application, the mouse-tracking paradigm is ready to be used in digital systems.

Chapter 6

Experience Sampling Method to Study the Variations of Self-Control Capacity in Daily Life

6.1 Introduction

As one of the hallmarks of human behavior, self-control has always been an important research topic in psychology and other behavioral sciences (Duckworth, 2011). Despite the lasting debate about the nature of self-control (e.g., Inzlicht & Berkman, 2015; Kurzban, 2012), self-control is generally described as inhibiting one's responses to immediate rewards (e.g., eating a delicious chocolate bar) for the sake of pursuing long-term goals, for example, to lose weight or to be healthier. (cf. Loewenstein, 2000). For decades, theorists have been puzzled by the ability of mankind to value abstract goals over concrete short-term rewards (Kanfer & Karoly, 1972), while at the same time people often fail to utilize this ability (Baumeister & Heatherton, 1996). There are many factors that may influence self-control outcomes, such as the nature and strength of temptation in a self-control task (e.g., Hur, Koo, & Hofmann, 2015; see also Chapter 5) and the personal value of a long-term goal (e.g., Saunders & Inzlicht, 2018). Among all the factors, self-control capacity as a person factor has been very central for the explanation in both folk theories (Bergen, 2011) and scientific models of self-control (Baumeister, Vohs, & Tice, 2007; Kotabe & Hofmann, 2015; Robinson, Schmeichel, & Inzlicht, 2010). Its central role is perhaps grounded in the casual observations as well as empirical evidence that the ability of controlling oneself differs greatly between individuals (de Ridder, Lensvelt-Mulders, Finkenauer, Stok, & Baumeister, 2012) as well as within individuals (e.g., Randles, Harlow, & Inzlicht, 2017). For developing digital interventions that can adapt to both individual differences and time-varying personal states, examining the inter-individual and intra-individual differences in self-control capacity is also of great importance.

²⁹ This chapter is based on Zhang, C., Smolders, K. C., Lakens, D., & IJsselstein, W. A. (2018). Two experience sampling studies examining the variation of self-control capacity and its relationship with core affect in daily life. *Journal of Research in Personality*, 74, 102-113.

Mischel and colleagues provided the first demonstration of the compelling effects of self-control capacity in the 1980s (Mischel, Shoda, & Peake, 1988; Mischel, Shoda, & Rodriguez, 1989). In a series of experiments, young children showed large differences in their ability to delay gratification (e.g., to receive more candies if they could withhold themselves from eating one candy for several minutes), and this individual difference was associated with their achievements much later in life. Although research on the origins of these differences has just begun (e.g., Hofmann, Gschwendner, Friese, Wiers, & Schmitt, 2008), the significance of inter-individual differences in self-control capacity have been suggested in a variety of domains (for a review, see de Ridder et al., 2012). For example, lower trait self-control is associated with lower academic achievements (Duckworth & Seligman, 2005) and poorer health (Crescioni et al., 2011; Moffitt et al., 2011).

Adding to the substantial correlational evidence, the experimental paradigm of ego-depletion (Muraven, Tice, & Baumeister, 1998) and its associated strength model (Baumeister et al., 2007) further established the special status of self-control capacity in self-control research. According to numerous ego-depletion studies (for a review, see Hagger et al., 2010), when people exert their self-control in an initial task, their performance in a second unrelated task that demands self-control worsened. This effect led to the view that self-control capacity is a resource that is domain-general and limited. However, recent meta-analyses failed to demonstrate convincing evidence for an ego-depletion effect, after correcting the literature for publication bias (Carter, Kofler, Forster, & McCullough, 2015; Etherton et al., 2018). More critically, a recent large multi-lab pre-registered replication study did not find support for a domain-general limited resource view of self-control (Hagger, et al., 2016).

Although the strength model seems too simplistic based on the empirical support, the concept of self-control capacity is still central to several more recent models of self-control (e.g., Hall & Fong, 2007; Hofmann et al., 2009; Kotabe & Hofmann, 2015; Robinson et al., 2010; but see Inzlicht & Schmeichel, 2012; Inzlicht, Schmeichel, & Macrae, 2013). Without direct evidence, it is perhaps prudent to be skeptical about the theoretical status of self-control capacity, especially because resource-like constructs have long been criticized as seemingly intuitive but lacking true exploratory power (see Navon, 1984). A specific question is whether and how self-control capacity is distinct from core affect in their causal roles of influencing self-control processes and outcomes. In the valence-arousal model (Russell & Barrett, 1999), core affect is defined as the change of a neurophysiological state underlying daily prototypical emotions (e.g., happiness, fear, anger, etc.) and can be expressed as subjective experiences on the valence (feeling good or bad) and arousal (feeling sleepy or activated) dimensions. The close relationship between core affect and self-control processes is both intuitive and supported by empirical research, as for example acute stress or tense arousal undermines self-control (e.g., Maier et al., 2015) and a certain level of arousal or alertness is required for the functioning of attention and cognitive control (e.g., Thomas et al., 2000). Moreover, some recent theories

have positioned affect at the core of cognitive control (Pessoa, 2009) and higher-level self-regulatory processes (Inzlicht, Bartholow, & Hirsh, 2015; Saunders & Inzlicht, 2018; also see Carver & Scheier, 1990). For example, in the appraisal framework by Saunders and Inzlicht (2018), conflict-elicited negative affect is proposed as a driving force for mobilizing self-control. Combining the empirical evidence and theoretical considerations, it is possible that the state variation of core affect is the major person-factor to explain self-control outcomes, which may at the same time give rise to the phenomenological feelings that capacity of self-control varies. One may even compare self-control capacity to emotional episodes (e.g., fear or anger), in the sense that they are psychologically constructed based on a momentary state of core affect and additional appraisal processes (Russell, 2003), rather than being causal entities. Given these open questions, it is valuable as a first step to accurately describe the inter-individual and intra-individual variations of self-control capacity, and to examine how these variations are associated with the variation of core affect (i.e., to build a nomological network, Cronbach & Meehl, 1955). We report here two experience sampling studies to fulfill these objectives with the hope to facilitate future theoretical work on the role of self-control capacity in self-control processes.

Firstly, we measure both trait self-control and self-reported state self-control capacity in a multi-session multi-day experience sampling protocol to differentiate the variation of self-control capacity at person-to-person, day-to-day, and moment-to-moment levels. We also aim to examine whether moment-to-moment variation can be explained by a time-of-day effect, and especially as a diurnal pattern, using a statistical technique called cosinor fitting (e.g., Hasler, Mehl, Bootzin, & Vazire, 2008; Murray et al., 2009). There are clear indications for diurnal patterns in people's core affect, such as arousal levels (e.g., Smolders, de Kort, & van den Berg, 2013; Stone et al., 2006; Wood & Magnello, 1992), and valence (Stone et al., 2006; Murray et al., 2009). A similar pattern in self-control capacity can be predicted based on its close link to the affective system, and this hypothesis was indirectly supported by a study in which students' performance in an online learning platform showed a diurnal pattern over the day (Randles et al., 2017). We use a self-report measure of state self-control capacity to examine its diurnal patterns more directly.

Secondly, we follow the valence-arousal model (Russell & Barrett, 1999) to measure four dimensions of core affect – valence, pure arousal, energetic arousal, and tense arousal (i.e., feeling of stress) – in the same experience sampling protocol. Thus, the relationships between self-control capacity and core affect can be quantified not only as correlations between trait self-control and dimensions of core affect, but also as disaggregated between-person and within-person correlations based on repeated state measures (see Curran & Bauer, 2011). Previous studies have demonstrated that trait self-control is associated with valence (Daly, Baumeister, Delaney, & MacLachlan, 2014; Galla & Wood, 2015; Hofmann, Luhmann, Fisher, Vohs, & Baumeister, 2014) and tense arousal (Bowlin & Baer, 2012; Galla & Wood, 2015;

Hofmann et al., 2014), using survey, daily diary, and experience sampling methods. These results suggest that people with higher trait self-control are also slightly happier in their daily lives (meta-analytic $r = 0.26$) and experience less psychological stress (meta-analytic $r = -0.28$; for forest plots with our studies added, see Figure 6.3 & 6.4 in the Discussion section³⁰). We replicate these results and extend these findings to include within-person correlations. Moreover, although arousal is clearly related to cognitive performance (Thomas et al., 2000) and its link to self-control capacity has been suggested (Randle et al., 2017), we provide a first direct description of the association.

The two research questions are also motivated by their relevance to digital lifestyle interventions. In theory, depending on a user's overall and temporary self-control capacity, the timing and types of interventions can be tailored. If state self-control capacity shows a strong and stable diurnal pattern, intervention messages for coaching a new healthy behavior can be sent during the hours when self-control capacity peaks for the user, as the person is likely to be more motivated to take effort to learn the new behavior. Similarly, if day-to-day variation in self-control capacity is large, preventive interventions may be used on the days when users has a more difficult time controlling themselves. Moreover, although experience sampling questions can be easily administrated in digital intervention systems to prob useful user information, answering these questions repeatedly can be burdening and irritating. Therefore, to minimize the number of questions, it is useful to know whether self-control capacity measures can provide different information than measures of core affect.

As the two experience sampling studies are very similar in their design and statistical analyses, we report the method and results of the two studies together in the following sections.

6.2 Method³¹

6.2.1 Participants

Study 1 was conducted within a larger experience sampling project in September 2015. In total, 172 Bachelor students at Eindhoven University of Technology participated in the data collection as partial fulfillment for a course. From this sample, 140 students provided permission to use their data for scientific research, so all analyses are based on this subset. The final sample consisted of 85 men and 55 women, and the median age was 19 (18-26, $SD = 2.43$). Study 2 was conducted in the same setting as Study 1 in September 2016. Out of the

³⁰ Meta-analyses and forest plots were done using the R package *metafor* (Viechtbauer, 2010).

³¹ Raw data, scripts, and other materials of the two experience sampling studies are available at Open Science Framework: <https://osf.io/hguz4/>. As our analyses focused on describing correlational patterns rather than testing specific hypotheses, we did not pre-register our data analyses.

163 students, 146 provided permission to use their data for scientific research. The final sample consisted of 83 men and 63 women, and the median age was 18 (17-25, $SD = 1.75$). Both studies were conducted in the third and fourth week of the academic year when students were not yet under the pressure of deadlines or exams.

The sample sizes were limited by the number of students enrolled in the course, which was deemed sufficient because analyses at measurement level (results in section 3.3, 3.5) were well-powered to detect even very small correlations with over 5000 observations. The intensively repeated measurements were also necessary to accurately measure diurnal patterns of the variables of interest. For the analyses at the person level (i.e., effects of trait self-control, section 3.4), sensitivity analyses in G*Power v3.192 revealed that our sample sizes in Study 1 and Study 2 were able to detect minimum effect sizes of $\rho = 0.230$ and $\rho = 0.234$ respectively, given an alpha level of 0.05 and 80% power. These effect sizes were just below the meta-analytic effect sizes of the correlations of trait self-control with affective valence and stress ($r = 0.26$ and $r = -0.28$).

6.2.2 Apparatus

A mobile experience sampling app was developed for the studies by the first author based on the open-source framework Experience Sampler³² (Thai & Page-Gould, 2018). The framework is a Cordova-based application template, supporting fast development of customized experience sampling apps for both Android and iOS platforms. In both studies, participants downloaded the app on their own smartphones to answer the experience sampling questionnaires.

6.2.3 Measurements

In both studies, trait self-control was measured by the brief version of the Trait Self-Control Scale (Tangney, Baumeister, & Boone, 2004). The scale is based on a comprehensive conceptualization of self-control and has been shown to have good internal reliability ($\alpha = 0.83$ and 0.85 in Tangney et al., 2004; $\alpha = 0.77$ and 0.83 in our studies), good test-retest reliability (0.87 in Tangney et al., 2004), and to correlate relatively well with behavioral indicators of self-control outcomes (de Ridder et al., 2012). Participants indicated to what extent they agreed with 13 statements on 5-point scales ranging from 1-*Completely disagree* to 5-*Completely agree*, such as “*I am good at resisting temptation*” and “*I wish I had more self-discipline*”.

³² A tutorial to get started with Experience Sampler by Sabrina Thai and Elizabeth Page-Gould can be found on <http://www.experiencesampler.com/>

We measured state self-control capacity using selected items from the State Self-Control Scale (Ciarocco, Twenge, Muraven, & Tice, 2015). With the intention to capture different aspects of state self-control, but not to overburden our participants, four items were selected, namely general willpower (*Right now, I feel my willpower is gone*), concentration (*Right now, it would take a lot of effort for me to concentrate on something*), urge-control (*Right now, I am having a hard time controlling my urges*), and motivation (*Right now, I am motivated to pursuit my (long-term) goals*³³). Participants responded to the items on 7-point scales, ranging from 1-*Not true* to 7-*Very true*. In a post-hoc content analysis on the State Self-Control Scale, our selection turned out to cover four out of the total six theoretical sub-constructs identified. Conceptually, the two omitted sub-constructs – mental fatigue and tense arousal – were captured when measuring the energetic and tense arousal dimensions of core affect. We decided to analyze the four items separately in order to explore potential sub-constructs in state self-control capacity and their variations, and to assess the reliability and validity of a potential short composite scale. According to a recent review on construct validity (Strauss & Smith, 2009), focusing on cohesive unidimensional rather than complex constructs should help to describe nomological networks more precisely.

The measurements of core affect followed the valence-arousal model of core affect (Russell & Barrett, 1999). In Study 1, each dimension of core affect was measured by a single item to minimize the burden on participants. Specifically, affective valence was measured on one 7-point bipolar scale, ranging from 1-*Very bad* to 7-*Very good*. Energetic arousal was measured on one 7-point unipolar scale, ranging from 1-*Not at all* to 7-*Very energetic*. Tense arousal or stress was measured on one 7-point scale, ranging from 1-*Not at all stressful* to 7-*Very stressful*. Sleepiness was measured by the 9-point Karolinska Sleepiness Scale (Åkerstedt & Gillberg, 1990) which included the following labels: 1-*Extremely sleepy*, 3-*Sleepy, but no difficulty remaining awake*, 5-*Neither alert nor sleepy*, 7-*Alert*, and 9-*Extremely alert*. In Study 2, two changes were made to the measurements. First, in order to increase reliability of the measures (see Schimmack, 2003), two items were used for each dimension of core affect, except that the standard 1-item Karolinska Sleepiness Scale was unchanged. Second, bipolar scales were used for all variables as recommended by Russell and Carroll (1999). As a result, two 7-point scales each were used for affective valence (*Unpleasant-Pleasant*, and *Sad-*

³³ This item was adapted as the original item was simply “I am motivated”, and the phrase “long-term” was only added in Study 2. The reason was to emphasize long-term goals, and to match a related item in the Total Trait Self-Control Scale (Tangney et al., 2004).

Happy), and for energetic arousal (*Inactive-Active*, and *Depleted-Energetic*). For the dimension of tense arousal or stress, two 7-point bipolar items (*Relaxed-Nervous*, and *Calm-Tense*) were added to the unipolar item used in Study 1³⁴.

6.2.4 Procedure

Both studies were introduced to the participants during one of their lectures in an introduction to psychology course. They were instructed to download the experience sampling app on their smartphone and to use it for one week. One day prior to the start of the study, participants indicated in the app when they normally wake up and go to sleep on weekdays and weekends. Between the self-reported wake-up time and sleep time, they were prompted eight times a day to answer the experience sampling questionnaires. The notifications were triggered by a semi-random algorithm to ensure that the adjacent two sessions were always at least one hour apart. In each experience sampling questionnaire, the state questions described above were presented in the following order on separate pages: affective valence, stress, energetic arousal, sleepiness, and state self-control (Study 1), or affective valence, energetic arousal, tense arousal, sleepiness, stress, and state self-control (Study 2). To prevent a drop in response rate, we sent two motivational messages to the participants during the study (on the 3rd and 6th day since the beginning of the study). Trait self-control was measured in a separate online questionnaire either before (Study 1) or after the sampling period (Study 2).

6.2.5 Data analysis

Given the three-level crossed design of our studies (all participants completed surveys in all seven days of a week, and in all 8 sessions from waking-up to sleep), we first built variance component models (multilevel null models) to decompose variance in each variable to difference sources. For variables with single-item measures, three-level models were built, resulting in variance from person (P), day-of-week (D), session number (S), and their interactions (see Equation 1). For variables with multiple-item measures, four-level models were built, resulting in variance from person (P), day-of-week (D), session number (S), item (I), and their interactions (see Equation 2). After variance decompositions, we followed generalizability theory (Cronbach, Gleser, Nanda, & Rajaratnam, 1972; Webb, Shavelson, & Haertel, 2006) to compute between-person and within-person reliability coefficients for the measurements used in the studies. Specifically, generalizability coefficients (E_{ρ}^2) as measures of between-

³⁴ Initially, stress was considered as a different construct from tense arousal, but the correlation between the two measures turned out to be very high and their relationships with other constructs were almost identical. Thus, we report stress and tense arousal as the same construct. Our current position is that even though the conceptualization of stress may be broader (e.g., to include stressors and behavioral responses), feelings of stress overlaps with the tense arousal dimension of core affect.

person reliability were computed (see Equation 3, based on Webb et al., 2006), which indicated the reliability of obtaining relative ranks of participants in terms of the measured variables when observed scores were averaged over all measurement facets (e.g., days, sessions, and items). Coefficients for within-person reliability of change (R_c) were computed to estimate the precision of measuring systematic changes in participants' self-control capacity and core affect over the 56 observations (see Equation 4, based on Cranford et al., 2015; Shrout & Lane, 2011). R_c could only be calculated for multiple-item measures, for which variance components of the person by day by session interaction (σ_{p*d*s}) could be separated from measurement error. Alternatively, we also computed averages of the observation-specific Cronbach's alpha coefficients for multiple-item measures (α_{ave} in Table 6.1 and 6.3).

$$Y_{ijk} = \beta_0 + P_i + D_j + S_k + (PD)_{ij} + (PS)_{ik} + (DS)_{jk} + \varepsilon_{ijk} \quad (1)$$

$$Y_{ijkl} = \beta_0 + P_i + D_j + S_k + I_l + (PD)_{ij} + (PS)_{ik} + (PI)_{il} + (DS)_{jk} + (DI)_{jl} + (SI)_{kl} + (PDS)_{ijk} + (PDI)_{ijl} + (PSI)_{ikl} + (DSI)_{jkl} + \varepsilon_{ijk} \quad (2)$$

$$\frac{\sigma_{person}^2}{\sigma_{person}^2 + \frac{\sigma_{p*d}^2}{n_d} + \frac{\sigma_{p*s}^2}{n_s} + \frac{\sigma_{p*i}^2}{n_i} + \frac{\sigma_{p*d*s}^2}{n_d n_s} + \frac{\sigma_{p*d*i}^2}{n_d n_i} + \frac{\sigma_{p*s*i}^2}{n_s n_i} + \frac{\sigma_{residual}^2}{n_d n_s n_i}} \quad (3)$$

$$\frac{\sigma_{p*d*s}^2}{\sigma_{p*d*s}^2 + \frac{\sigma_{residual}^2}{n_i}} \quad (4)$$

After obtaining variance components and reliability coefficients, we moved on to build four types of multilevel models to answer more substantial research questions. First, for each variable, a simplified multilevel null model was built to serve as the baseline model, with only person and person by day-of-week interaction as grouping variables in the model (model 1, see Equation 5)³⁵.

$$Y_{ijk} = \beta_0 + P_i + (PD)_{ij} + \varepsilon_{ijk} \quad (5)$$

Secondly, a cosinor fitting procedure (e.g., Hasler et al., 2008; Murray et al., 2009) was used to model the diurnal pattern of state self-control capacity and core affect. The basic idea was to estimate how much variance of a variable can be explained by fitting sinusoids to the data. Technically, for each variable, a random-intercept model (model 2) was fitted by adding a sine component ($\sin T_{ijk}$) and a cosine component ($\cos T_{ijk}$) to the null model defined in Equation 5. The sine and cosine components were computed from the time variable (T_{ijk} in hours after midnight, e.g., 10.5 for 10:30) as in Equation 7 and 8.

³⁵ Simplification was made for two reasons: (1) variance components of day-of-week ($\sigma_{day-of-week}$) and day-of-week by session (σ_{d*s}) were very small (see Section 3.2); (2) time variables $\sin T_{ijk}$ and $\cos T_{ijk}$ would correlate highly with session number and person by session as grouping variables.

$$Y_{ijk} = \beta_0 + P_i + (PD)_{ij} + \beta_1 \sin T_{ijk} + \beta_2 \cos T_{ijk} + \varepsilon_{ijk} \quad (6)$$

$$\sin T_{ijk} = \sin(T_{ijk} \times 2\pi / 24) \quad (7)$$

$$\cos T_{ijk} = \cos(T_{ijk} \times 2\pi / 24) \quad (8)$$

The improvement of model fit from model 1 to model 2 provides a measure of the strength of the diurnal pattern of modeled variables. The amplitude and phase angle of each fitted sinusoid can be calculated from β_1 and β_2 using Equation 9 and 10. Amplitude measures the maximum deviation of a variable from its mean so that a larger amplitude implies a stronger diurnal pattern, or in other words, larger cyclical fluctuations around the mean. Phase angle indicates at what time of day a variable shows its maximum deviation from the mean (peaks and valleys of the sinusoid).

$$Amp = \sqrt{\beta_1^2 + \beta_2^2} \quad (9)$$

$$Phase\ shift = \begin{cases} \tan^{-1}(\beta_1/\beta_2) + \pi, & IF\ \beta_2 < 0 \\ \tan^{-1}(\beta_1/\beta_2), & IF\ \beta_2 > 0\ AND\ \beta_1 > 0 \\ \tan^{-1}(\beta_1/\beta_2) + 2\pi, & IF\ \beta_2 > 0\ AND\ \beta_1 < 0 \end{cases} \quad (10)$$

Thirdly, for each variable, a random-slope model (model 3) was fitted by allowing the coefficients of $\sin T_{ijk}$ and $\cos T_{ijk}$ to vary for different participants and for different days within participants, in order to explore variations in amplitude and phase angle at the person-to-person and day-to-day levels (see Equation 11).

$$Y_{ijk} = \beta_0 + P_{0i} + (PD)_{0ij} + \beta_1 \sin T_{ijk} + \beta_2 \cos T_{ijk} + P_{1i} \sin T_{ijk} + (PD)_{1ij} \sin T_{ijk} \\ + P_{2i} \cos T_{ijk} + (PD)_{2ij} \cos T_{ijk} + \varepsilon_{ijk} \quad (11)$$

Finally, for examining the relationship between self-control capacity and dimensions of core affect, analyses were done using three different methods. To estimate the association between trait self-control and core affect, trait self-control was added as a person-level predictor in model 4. The effects were measured as β_3 (see Equation 12). Between-person and within-person correlations were estimated based on the data of state self-control capacity and core affect. The between-person component of each variable was estimated as β_0 in the corresponding random-slope models (Equation 11), and then Pearson correlation coefficients were computed based on the estimated person means. The within-person component of each variable was estimated as the residuals (ε_{ijk}) in model 3 (Equation 11), and then Pearson correlation coefficients were computed based on the estimated within-person residuals. The within-person correlations estimated represented correlations between the momentary fluctuations of pairs of variables, after removing stable time-of-day effects.

$$Y_{ijk} = \beta_0 + P_{0i} + (PD)_{0ij} + \beta_1 \sin T_{ijk} + \beta_2 \cos T_{ijk} + P_{1i} \sin T_{ijk} + (PD)_{1ij} \sin T_{ijk} + P_{2i} \cos T_{ijk} + (PD)_{2ij} \cos T_{ijk} + \beta_3 TSC_i + \varepsilon_{ijk} \quad (12)$$

All models were fitted using the package *lme4* version 1.1-12 (Bates, Maechler, Bolker, & Walker, 2015) in R programming environment 3.22³⁶ (R Development Core Team, 2015).

6.3 Results

6.3.1 Response rate

In Study 1, 4987 observations were collected from 140 participants. The mean response rate was 63.6% (median = 69.6%, $SD = 23.8\%$). In Study 2, 5599 observations were collected from the 146 participants. The mean response rate was 68.5% (median = 75.0%, $SD = 22.0\%$)³⁷. Given that there was no monetary reward for the participants, the respondent rates were slightly lower than the ones in some other experience sampling studies (77% in a meta-analysis, see Hofmann & Patel, 2015), but they were still comparable to some studies with payment (e.g., Wilt, Funkhouser, & Revelle, 2011). As we planned to estimate diurnal patterns in model 3 and model 4, a sufficient number of observations for different periods of the day were required. Therefore, participants with less than 5 observations in either the morning (06:00 - 12:00), afternoon (12:00 - 18:00), or evening (later than 18:00) were excluded, leaving us with 125 participants in Study 1 and 132 participants in Study 2 for the following analyses³⁸.

6.3.2 Descriptives, variance components, and reliability coefficients

Table 6.1 to Table 6.3 summarize descriptive results of the data, including sample means, variance components, and derived reliability coefficients. According to the sample means, on average, participants were in positive mood, experienced low tense arousal or stress, and reported to have relatively high state self-control capacity (particularly for the *willpower* and *urge control* items). In general, inter-individual difference (σ_{person}^2) accounted for a sizable portion of total variance for all variables (10-40%), and its contribution was larger for three components of state self-control capacity (willpower, urge control, and motivation) and tense arousal (25-40%) than for energetic arousal, sleepiness, and concentration (10-20%). In contrast, there was no overall day-to-day variance ($\sigma_{day-of-week}^2$), but person-specific day-to-day

³⁶ We also fitted a Bayesian version of all models by using the package “brms” version 1.10.0 (Bürkner, 2017), with order-1 autoregressive structure specified within the outcome variables. As results of the two approaches were almost identical, only results from the “lme4” package are reported. Code for the Bayesian models can be found in the scripts at the Open Science Framework.

³⁷ For each scheduled notification, only responses that were made within half an hour were included to avoid largely overlapping observations in time.

³⁸ Including the excluded participants did not change the results to any significant extent.

6.3 Results

variation (σ_{p*d}^2) did explain a small portion of total variance for each variable (6-12%).

This trend of was slightly more evident for core affect than for state self-control capacity. Moreover, overall and person-specific session-to-session variation accounted for about 10% of the total variance for sleepiness, around 3-5% for energetic arousal, concentration, and motivation, but not for other variables. Finally, when within-person within-day momentary variations (σ_{p*d*s}^2) were estimated for variables with multiple-item measures, they greatly reduced residual variance and accounted for about 15-40% of the total variance. Results also indicated that state self-control capacity had less momentary fluctuations than core affect.

In terms of reliability, between-person reliability (E_{ρ^2}) was good for most measures, with generalizability coefficient ranging from 0.75 to 0.95. For determining the relative rank of participants' scores on core affect and state self-control capacity, single-item measures seem to be good enough in typical experience sampling designs, since large number of observations are averaged (56 in our case). However, within-person reliability coefficients of change (R_c) were lower, ranging from acceptable for core affect (0.69 to 0.77) to insufficient for state self-control capacity (below 0.7). Although the number of measurement items should be increased whenever possible to increase within-person reliability, our results of within-person correlations were unlikely to suffer from this issue because measurement errors of a pair of variables were independent in theory. Our analyses of reliability did question the validity of combining state self-control capacity items to a total score (see Table 6.3), as evident by the low average alpha coefficients (α_{ave}), low reliability of change, and the large percentages of total variance accounted by overall and person-specific item-to-item variations (σ_{item}^2 and σ_{p*i}^2).

Table 6.1 Means, variance components, and reliability coefficients for core affect.

	Valence	%	EA	%	TA	%	Sleepiness	%
σ_{person}^2	0.31	20.7	0.28	15.0	0.61	30.6	0.32	11.0
	0.25	17.6	0.24	12.1	0.50	24.9	0.58	18.0
$\sigma_{day-of-week}^2$	0.02	1.2	0.02	0.9	0.03	1.6	0.02	0.6
	0.01	0.4	0.02	0.8	0.01	0.5	0.03	0.8
$\sigma_{session}^2$	0.03	1.9	0.08	4.1	0.02	0.8	0.32	10.9
	0.01	1.0	0.06	3.2	0.01	0.6	0.25	7.9
σ_{item}^2	—	—	—	—	—	—	—	—
	0.00	0	0.00	0.2	0.01	0.5	—	—
σ_{p*d}^2	0.18	11.9	0.19	10.4	0.23	11.7	0.31	10.4
	0.11	7.7	0.14	7.4	0.16	7.7	0.27	8.5

(To be continued)

Chapter 6 - Experience Sampling Method to Study the Variations of Self-Control Capacity

	Valence	%	EA	%	TA	%	Sleepiness	%
σ_{p*s}^2	0.02	1.2	0.07	4.0	0.02	0.9	0.20	6.9
	0.02	1.2	0.10	5.0	0.03	1.7	0.29	8.9
σ_{p*i}^2	–	–	–	–	–	–	0.03	0.9
	0.03	2.2	0.03	1.8	0.05	2.4	0.04	1.3
σ_{d*s}^2	0.01	0.5	0.02	1.2	0.00	0.2	–	–
	0.01	0.6	0.02	0.9	0.01	0.4	–	–
σ_{d*i}^2	–	–	–	–	–	–	–	–
	0.00	0	0.00	0	0.00	0.2	–	–
σ_{s*i}^2	–	–	–	–	–	–	–	–
	0.00	0	0.00	0	0.00	0	–	–
σ_{p*d*s}^2	–	–	–	–	–	–	–	–
	0.62	43.3	0.69	35.6	0.58	29.0	–	–
σ_{p*d*i}^2	–	–	–	–	–	–	–	–
	0.00	0.2	0.03	1.3	0.05	2.4	–	–
σ_{p*s*i}^2	–	–	–	–	–	–	–	–
	0.00	0	0.00	0	0.01	0.3	–	–
σ_{d*s*i}^2	–	–	–	–	–	–	–	–
	0.00	0	0.00	0	0.00	0	–	–
$\sigma_{residual}^2$	0.93	62.6	1.19	64.4	1.08	54.2	1.76	59.4
	0.37	25.8	0.61	31.6	0.59	29.4	1.76	54.6
σ_{total}^2	1.48	100	1.85	100	1.99	100	2.95	100
	1.44	100	1.94	100	2.01	100	3.22	100
α_{ave}	–	–	–	–	–	–	–	–
	0.82	–	0.77	–	0.84	–	–	–
E_{ρ^2}	0.87	–	0.83	–	0.92	–	0.76	–
	0.84	–	0.77	–	0.89	–	0.85	–
R_c	–	–	–	–	–	–	–	–
	0.77	–	0.69	–	0.75	–	–	–
mean	5.04	–	4.10	–	2.67	–	4.34	–
	4.81	–	4.12	–	2.96	–	4.71	–

Note: EA = energetic arousal; TA = tense arousal; In interaction terms, p = person, d = day-of-week, s = session, i = item. For ease of presentation, the results of Study 1 and Study 2 are shown at the top and bottom of each cell respectively.

6.3 Results

Table 6.2 Means, variance components, and reliability coefficients for items of state self-control capacity.

	SSC_con	%	SSC_wil	%	SSC_urg	%	SSC_mot	%
σ_{person}^2	0.54	20.1	0.90	37.5	0.88	39.0	0.62	26.4
	0.51	16.1	0.70	27.2	0.78	31.2	0.76	32.1
$\sigma_{day-of-week}^2$	0.01	0.4	0.01	0.2	0.01	0.3	0.01	0.3
	0.00	0.1	0.00	0	0.01	0.5	0.00	0.2
$\sigma_{session}^2$	0.11	4.1	0.02	0.9	0.01	0.6	0.06	2.4
	0.18	5.6	0.04	1.5	0.03	1.0	0.07	3.0
σ_{p*d}^2	0.19	7.2	0.15	6.2	0.14	6.3	0.19	7.9
	0.21	6.6	0.18	7.0	0.21	8.4	0.16	6.8
σ_{p*s}^2	0.11	4.3	0.07	2.9	0.05	2.2	0.07	3.1
	0.19	5.9	0.07	2.6	0.06	2.3	0.11	4.5
σ_{d*s}^2	0.01	0.3	0.00	0.1	0.00	0	0.01	0.4
	0.02	0.7	0.01	0.4	0.00	0.1	0.00	0.1
$\sigma_{residual}^2$	1.70	63.5	1.25	52.2	1.16	51.6	1.41	59.6
	2.08	65.0	1.57	61.2	1.42	56.4	1.26	53.2
σ_{total}^2	2.68	100	2.40	100	2.24	100	2.24	100
	3.20	100	2.57	100	2.51	100	2.51	100
E_{ρ^2}	0.88		0.95		0.95		0.91	
	0.85		0.92		0.93		0.93	
mean	4.40		5.29		5.35		3.98	
	4.28		5.06		5.20		3.95	

Note: SSC_con = state self-control, concentration; ssc_wil = state self-control, willpower; SSC_urg = state self-control, urge control; SSC_mot = state self-control, motivation; In interaction terms, p = person, d = day-of-week, s = session, i = item. For ease of presentation, the results of Study 1 and Study 2 are shown at the top and bottom of each cell respectively.

Table 6.3 Means, variance components, and reliability coefficients for 4-item and 2-item state self-control scales.

	4-item state self-control	%	2- item state self-control (willpower, urge control)	%
σ_{person}^2	0.28 / 0.31	9.8 / 10.2	0.65 / 0.61	28.1 / 24.0
$\sigma_{day-of-week}^2$	0.01 / 0.00	0.3 / 0	0.01 / 0.00	0.3 / 0.1
$\sigma_{session}^2$	0.04 / 0.07	1.3 / 2.2	0.02 / 0.03	0.8 / 1.3
σ_{item}^2	0.43 / 0.37	15.0 / 12.1	0.00 / 0.01	0.1 / 0.4
σ_{p*d}^2	0.09 / 0.10	3.1 / 3.3	0.09 / 0.14	3.8 / 5.5
σ_{p*s}^2	0.06 / 0.05	2.1 / 1.6	0.04 / 0.00	1.8 / 0
σ_{p*i}^2	0.46 / 0.38	16.0 / 12.6	0.23 / 0.13	10.0 / 5.1
σ_{d*s}^2	0.00 / 0.00	0.1 / 0.2	0.00 / 0.00	0 / 0.1
σ_{d*i}^2	0.00 / 0.00	0 / 0.1	0.00 / 0.00	0 / 0.1
σ_{s*i}^2	0.01 / 0.01	0.4 / 0.4	0.00 / 0.00	0 / 0
σ_{p*d*s}^2	0.43 / 0.52	15.0 / 17.1	0.40 / 0.51	17.1 / 20.0
σ_{p*d*i}^2	0.08 / 0.09	2.8 / 3.1	0.05 / 0.06	2.4 / 2.2
σ_{p*s*i}^2	0.02 / 0.06	0.6 / 1.9	0.02 / 0.07	0.7 / 2.8
σ_{d*s*i}^2	0.00 / 0.01	0.1 / 0.2	0.00 / 0.00	0 / 0.2
$\sigma_{residual}^2$	0.95 / 1.06	33.4 / 35.1	0.81 / 0.98	34.8 / 38.2
σ_{total}^2	2.85 / 3.04	100 / 100	2.32 / 2.56	100 / 100
E_{ρ^2}	0.65 / 0.70		0.81 / 0.85	
R_c	0.64 / 0.66		0.50 / 0.51	
α_{ave}	0.69 / 0.70		0.69 / 0.66	
mean	4.76 / 4.62		5.32 / 5.14	

Note: In interaction terms, $p = person$, $d = day-of-week$, $s = session$, $i = item$. For ease of presentation, the results of Study 1 and Study 2 are separated by “/”.

6.3.3 Modeling diurnal patterns in state self-control and affective states

Adding the time variables ($\sin T_{ijk}$ and $\cos T_{ijk}$) into the models (from model 1 to model 2) significantly improved model fit for all state measures ($p < .0001$ in all Chi-square tests), indicating some degree of diurnal patterns in the temporal variations of state self-control capacity and core affect. Table 6.4 summarizes the results of the random-intercept models (model 2) and the random-slope models (model 3). An increase in pseudo R^2 indicates an improved model fit by adding the time variables as predictors (model 2), or by further allowing the coefficients for the time variables to vary between participants and between days within each participant (model 3)³⁹. Among the four state self-control components, concentration had a stronger diurnal pattern than the other three components, and according to the phase angle estimates, concentration level on average peaked at around 13:30 - 14:00 for the student participants. It was also interesting to observe that the peak of urge control happened in the late morning (around 10:30), while all the other components of state self-control capacity peaked in the early afternoon (13:00 - 14:30). However, as the diurnal patterns for urge control and willpower were very weak, the estimated phase angles should be interpreted with caution. Among the dimensions of core affect, sleepiness had the strongest diurnal pattern, followed by energetic arousal, while valence and tense arousal had very weak patterns. According to the phase angle estimates, participants were happiest in the evening (cf. Kahneman, Krueger, Schkade, Schwarz, & Stone, 2004), most stressed in the early afternoon (around 13:00), most vital at around 15:00, and sleepest at around 3:00.

³⁹ Δ pseudo- R^2 was calculated as the increase in conditional R^2_{LMM} (see Johnson, 2014; Nakagawa & Schielzeth, 2013) from the null models (model 1) to the random-intercept models (model 2), and to the random-slope models (model 3). The calculations were done using the piecewiseSEM package in R (Lefcheck, 2015).

Table 6.4 Results of the random-intercept and random-slope cosinor models.

State Variable	<i>Model 2</i>			<i>Model 3</i>		
	Model fit change (Δ pseudo- R^2)	Amplitude	Phase angles	Model fit change (Δ pseudo- R^2)	Amplitude	Phase angles
Valence	1.5%	0.18	18:52	11.8%	0.18	18:50
	1.3%	0.16	20:41	9.9%	0.15	20:34
EA	4.3%	0.36	15:03	16.6%	0.37	15:04
	4.0%	0.37	15:05	16.6%	0.37	15:11
TA	1.1%	0.16	13:55	8.0%	0.16	13:56
	0.8%	0.14	13:06	9.1%	0.13	13:14
Sleepiness	11.6%	0.62	3:17	27.2%	0.63	3:16
	7.8%	0.53	3:03	23.3%	0.54	3:06
SSC_con	5.2%	0.37	14:14	14.4%	0.36	14:09
	5.7%	0.36	13:21	19.1%	0.37	13:38
SSC_wil	1.3%	0.17	13:11	7.9%	0.17	13:01
	1.5%	0.16	11:50	9.1%	0.17	12:06
SSC_urg	0.7%	0.11	11:05	7.3%	0.10	10:50
	1.1%	0.14	11:17	7.4%	0.13	11:09
SSC_mot	3.1%	0.28	14:26	11.9%	0.27	14:25
	3.7%	0.27	13:10	11.5%	0.27	13:12

Note: For ease of presentation, the results of Study 1 and Study 2 are shown at the top and bottom of each cell respectively.

6.3.4 Effects of trait self-control on affective states and state self-control capacity

The effects of trait self-control on affective states were estimated as the parameter β_3 (standardized regression coefficient) in model 4. As expected, trait self-control seemed to positively correlate with valence across both studies (see Figure 6.1), although the effects were quite small and nonsignificant in the individual studies (Study 1: $\beta = 0.07$, 95% CI = [-0.01, 0.16], $p = .097$, $r = 0.13$ ⁴⁰; Study 2: $\beta = 0.08$, 95% CI = [-0.01, 0.16], $p = .090$, $r = 0.11$). Moreover, trait self-control was correlated to tense arousal in the studies (see Figure 6.1), meaning that students with higher trait self-control experienced, on average, less tension or stress during the sampling week (Study 1: $\beta = -0.14$, 95% CI = [-0.24, -0.05], $p = .005$, $r = -0.24$; Study 2: $\beta = -0.15$, 95% CI = [-0.26, -0.04], $p = .009$, $r = -0.22$). However, in both studies, trait self-control did not significantly correlate with energetic arousal (Study 1: $\beta = 0.02$, 95% CI = [-0.05, 0.10], $p = 0.57$; Study 2: $\beta = 0.05$, 95% CI = [-0.03, 0.13], $p = .24$) nor sleepiness (Study 1: $\beta = -0.04$, 95% CI = [-0.11, 0.02], $p = .20$; Study 2: $\beta = -0.05$, 95% CI = [-0.13, 0.03], $p = .24$).

We also examined the relationship between trait self-control and state self-control capacity in the same way (see Figure 6.2). Results revealed reliable positive correlations between trait self-control and three components of state self-control capacity, including concentration (Study 1: $\beta = 0.15$, 95% CI = [0.08, 0.23], $p < .001$, $r = 0.28$; Study 2: $\beta = 0.16$, 95% CI = [0.08, 0.23], $p < .001$, $r = 0.31$), willpower (Study 1: $\beta = 0.18$, 95% CI = [0.08, 0.28], $p < .001$, $r = 0.28$; Study 2: $\beta = 0.25$, 95% CI = [0.16, 0.34], $p < .001$, $r = 0.41$), urge control (Study 1: $\beta = 0.25$, 95% CI = [0.15, 0.35], $p < .001$, $r = 0.39$; Study 2: $\beta = 0.26$, 95% CI = [0.17, 0.35], $p < .001$, $r = 0.48$). For motivation, the results were inconsistent across the two studies – trait self-control had a small positive effect on motivation in Study 2 ($\beta = 0.12$, 95% CI = [0.02, 0.23], $p = .027$, $r = 0.19$), but not in Study 1 ($\beta = 0.005$, 95% CI = [-0.09, 0.10], $p = .92$, $r = 0.01$). The inconsistency might be due to the change in the phrasing of the item between the two studies – emphasizing long-term goals better matched the conceptualization of self-control in the Brief Trait Self-Control Scale (see 6.2.3 for details).

⁴⁰ The r henceforth, as an effect-size measure, was calculated as the correlational strength between trait self-control and the between-person components of the state variables. It provides a way to compare our data with the effect sizes in previous studies with different data structures.

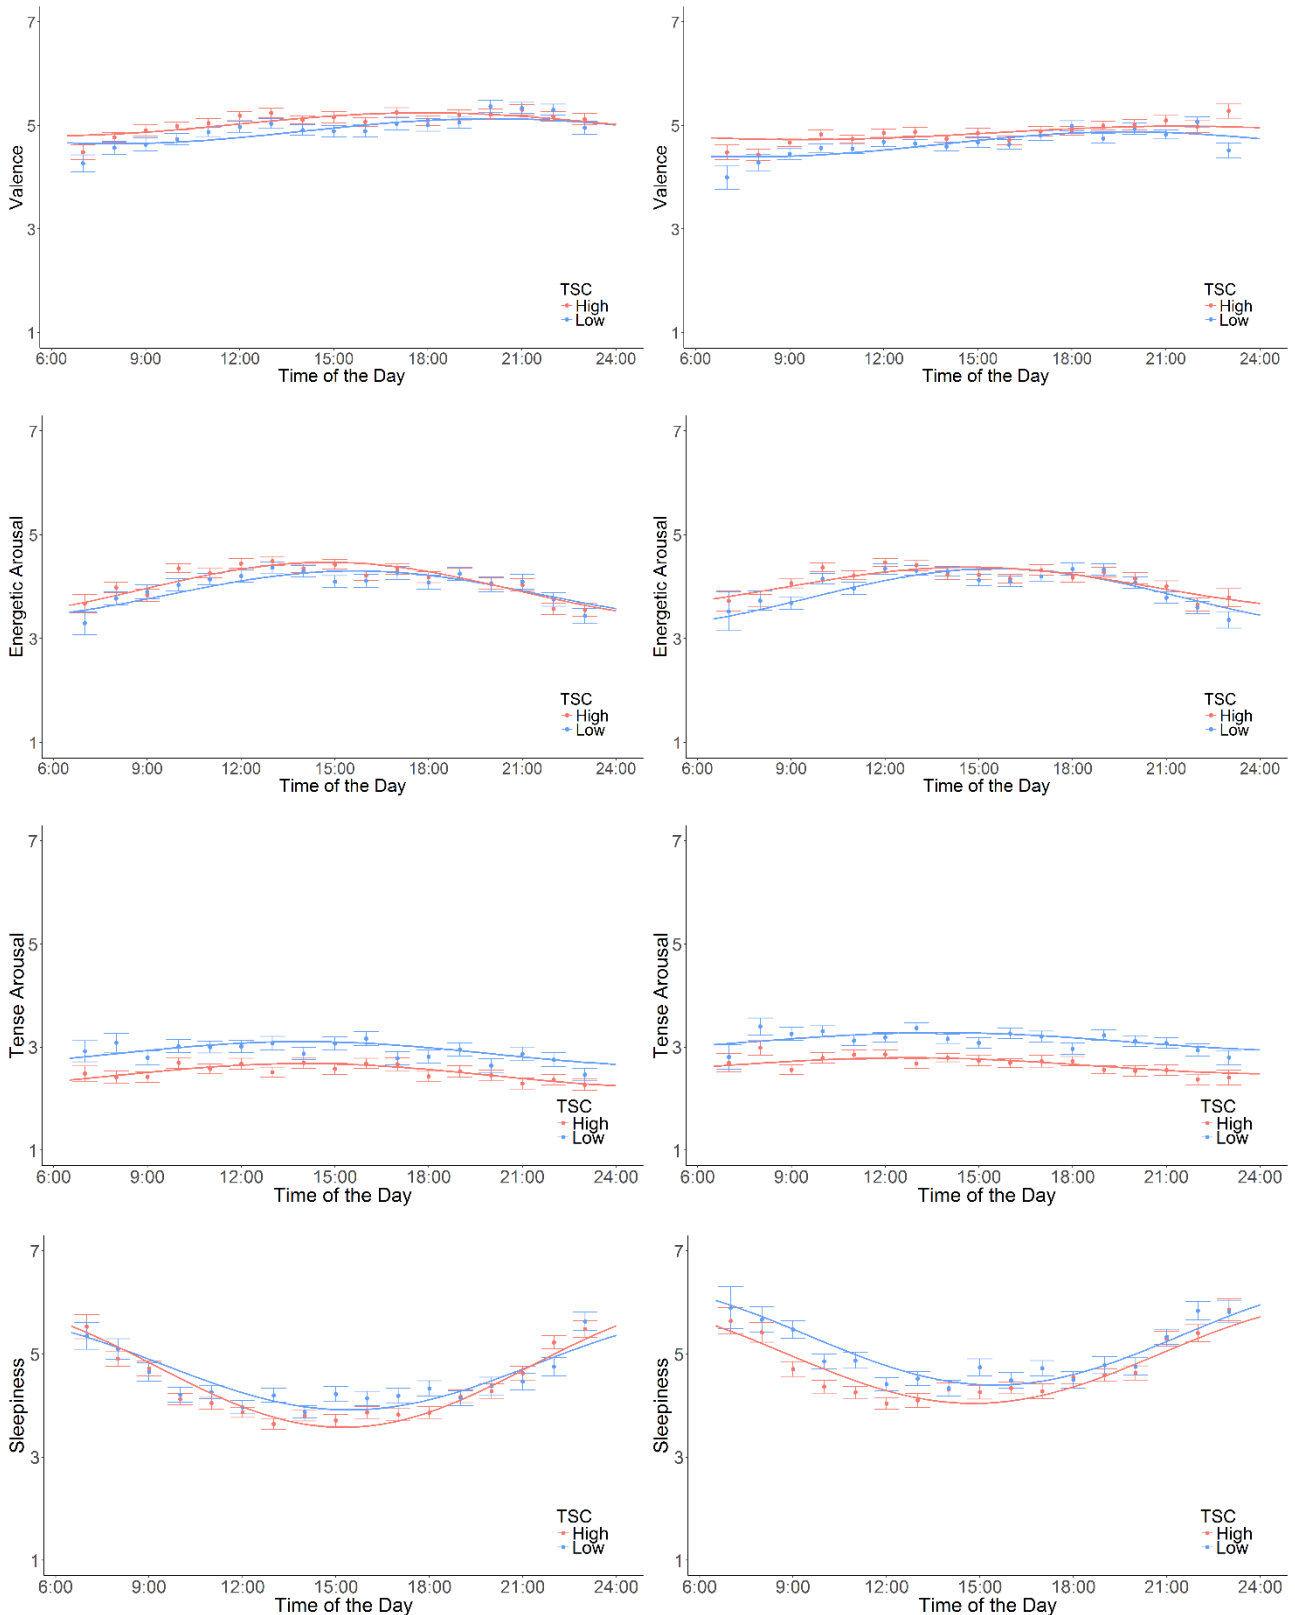


Figure 6.1 Modeled diurnal patterns of core affect dimensions for high and low trait self-control capacity groups. Mean values with error bars (one SE) for different hours of the day (between 7 and 23) are plotted together with the fitted sinusoidals.

6.3 Results

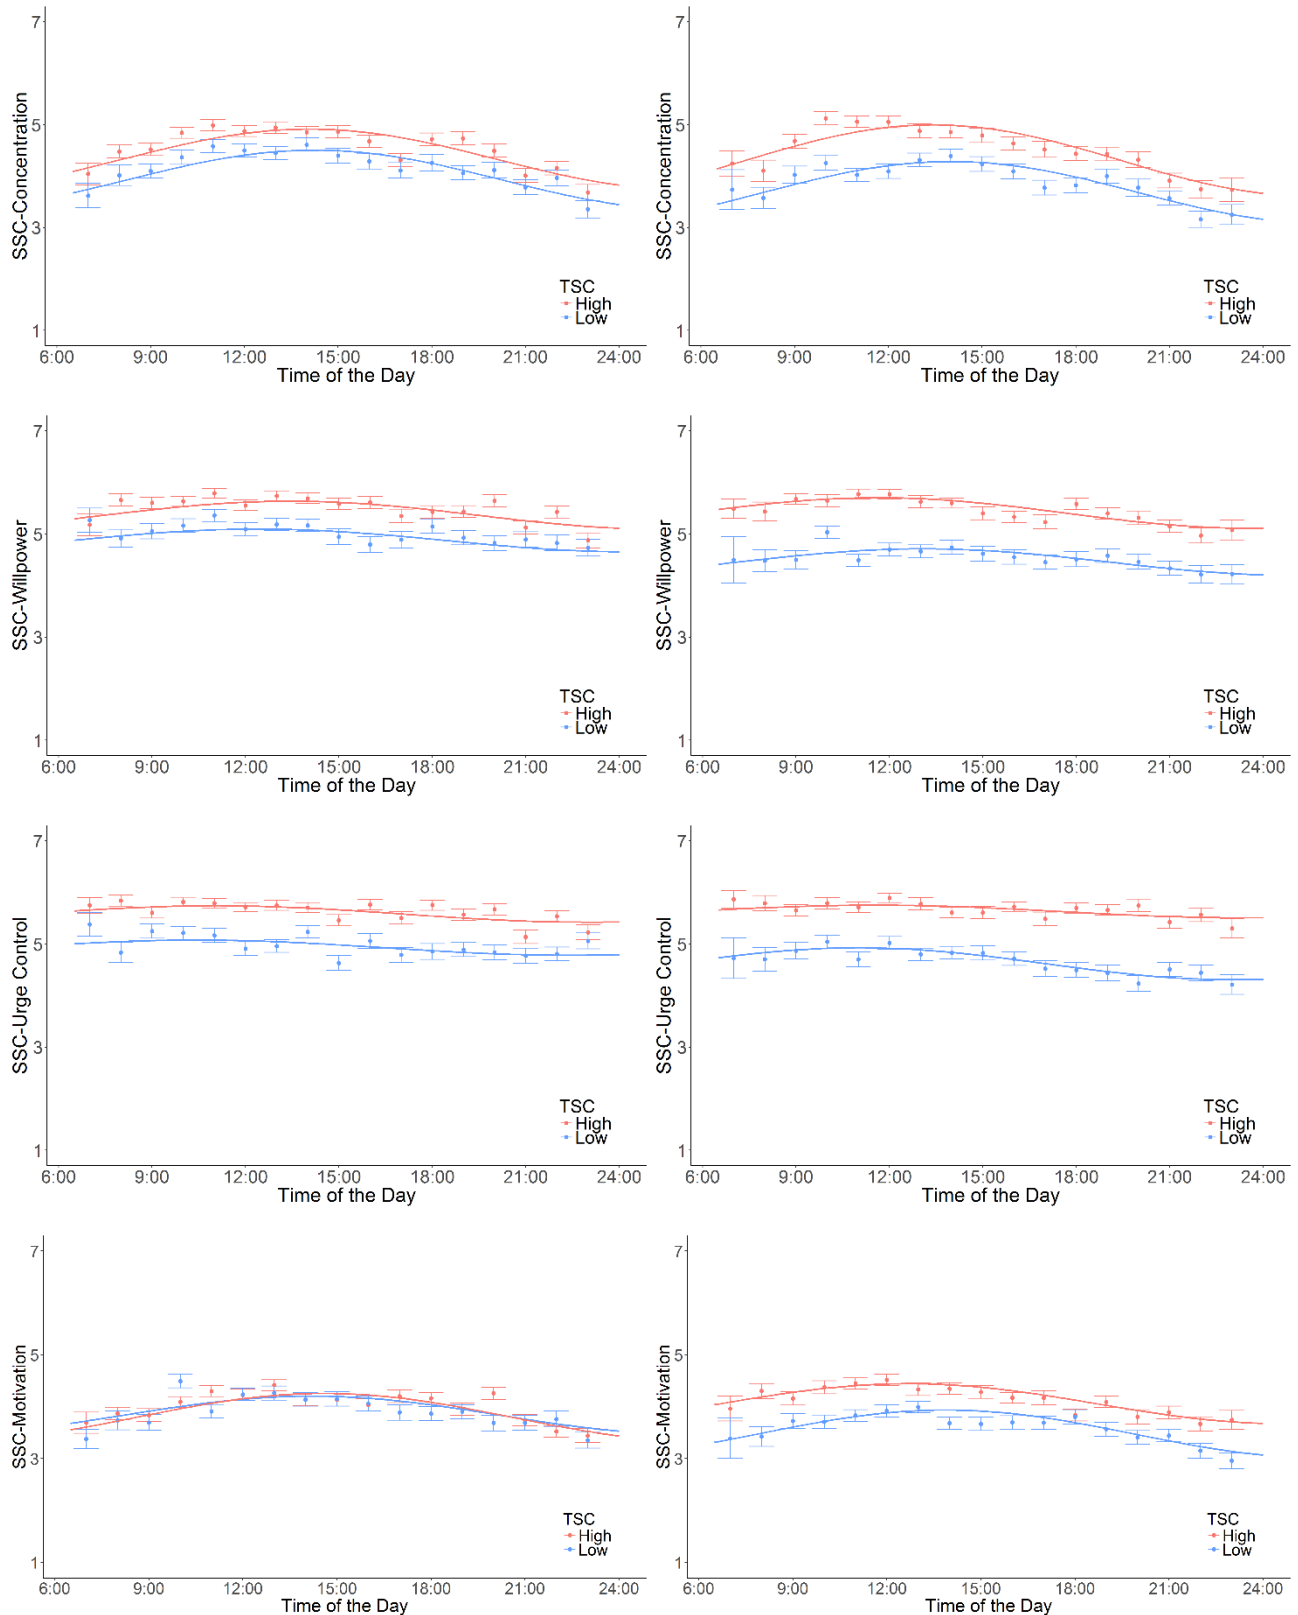


Figure 6.2 Modeled diurnal patterns of components of state self-control capacity for high and low trait self-control capacity groups. Mean values with error bars (one SE) for different hours of the day (between 7 and 23) are plotted together with the fitted sinusoids.

6.3.5 Between and within-person correlations between state self-control and affective states

Table 6.5 and Table 6.6 summarizes between-person and within-person correlations between pairs of state measures used in the two studies respectively. For understanding the relationship between two variables, a between-person correlation reflects the accumulated association over time at person-level, indicating whether a person scoring high on one variable tends to score high or low on the other variable. In contrast, a within-person correlation indicates whether two variables tend to deviate from personal means at the same time. Overall, results showed a clear pattern that between-person correlations were generally stronger than within-person correlations.

In the cluster of core affect (top-left of the tables), the within-person correlations in general supported the valence-arousal model of core affect (Russell & Barrett, 1999): Weak or no correlation between the orthogonal dimensions (valence-arousal, energetic arousal-tense arousal), and moderate correlations in the expected directions between the adjacent dimensions (valence-energetic arousal, valence-tense arousal, and energetic arousal-sleepiness). However, inconsistent with valence-arousal model, a negative correlation between tense arousal and sleepiness (or positive correlation between tense arousal and pure arousal) was not found.

In the state self-control cluster (bottom-right), it was interesting to observe that the motivation component only correlated very weakly with the other three highly correlated components at the between-person level, but it did correlate strongly with them at the within-person level. In other words, participants' average levels of motivation were largely independent from the other components, but they tended to have stronger motivation at the time when they had more willpower, could concentrate well, and control their urges.

The relationship between state self-control capacity and core affect is shown in the left-bottom of the two tables. At the between-person level (Table 6.5), participants with higher average self-control capacity were also happier, more energetic, less stressed, and less sleepy on average. These results were in line with the correlations of trait self-control and these dimensions of core affect. However, the motivation component of self-control correlated strongly with energetic arousal and sleepiness, weakly with valence, but was independent from tense arousal. When within-person correlations were examined (Table 6.6), it was intriguing to observe that state self-control's correlations with valence shrank, and the correlation with stress disappeared almost completely. In contrast, its relationship with energetic arousal and sleepiness remained at the same level, with even a small increase for the motivation component.

6.3 Results

Table 6.5 Between-person correlation matrix of state measures.

	Valence	EA	TA	Sleepiness	SSC_con	SSC_wil	SSC_urg	SSC_mot
Valence	1							
EA	0.64 0.68	1						
TA	-0.56 -0.46	-0.33 -0.28	1					
Sleepi- ness	-0.51 -0.43	-0.66 -0.52	0.33 0.27	1				
SSC_con	0.31 0.32	0.27 0.36	-0.46 -0.36	-0.43 -0.48	1			
SSC_wil	0.50 0.24	0.34 0.25	-0.55 -0.35	-0.38 -0.23	0.65 0.60	1		
SSC_urg	0.32 0.17	0.16 0.12	-0.47 -0.47	-0.30 -0.19	0.58 0.55	0.71 0.76	1	
SSC_mot	0.13 0.23	0.35 0.30	-0.04 0.02	-0.36 -0.17	0.14 0.24	0.08 0.24	-0.04 0.09	1

Note: For ease of presentation, the results of Study 1 and Study 2 are shown at the top and bottom of each cell respectively.

Table 6.6 Within-person correlation matrix of state measures.

	Valence	EA	TA	Sleepiness	SSC_con	SSC_wil	SSC_urg	SSC_mot
Valence	1							
EA	0.43 0.44	1						
TA	-0.28 -0.36	-0.06 -0.03	1					
Sleepi- ness	-0.34 -0.28	-0.56 -0.54	0.01 -0.02	1				
SSC_con	0.23 0.19	0.37 0.36	-0.05 -0.01	-0.41 -0.38	1			
SSC_wil	0.36 0.31	0.32 0.39	-0.16 -0.15	-0.33 -0.33	0.41 0.45	1		
SSC_urg	0.16 0.14	0.14 0.18	-0.08 -0.11	-0.18 -0.19	0.30 0.30	0.35 0.34	1	
SSC_mot	0.27 0.23	0.36 0.33	-0.02 -0.04	-0.35 -0.28	0.33 0.34	0.33 0.32	0.25 0.23	1

Note: For ease of presentation, the results of Study 1 and Study 2 are shown at the top and bottom of each cell respectively.

6.4. Discussion

In two experience sampling studies, we studied the variation of self-control capacity among university students in their daily lives. Using a multilevel modeling approach and the cosinor fitting method, we were able to quantify the variation of self-control capacity at different levels of interests, and to model its temporal variation as diurnal patterns. Moreover, we systematically examined the relationship between self-control capacity and core affect, in order to facilitate better theorization of self-control capacity as a construct. With two large group of participants and time-intensive repeated measures, results were quite consistent across the two studies and are likely to be robust, at least for similar student samples in similar contexts. We discuss the implications of the results for self-control research and for potential applications to digital lifestyle interventions.

6.4.1 Variations of self-control capacity

The first set of findings simply characterizes the variability of self-control capacity. Compared with valence, pure arousal (sleepiness), and energetic arousal, self-control capacity seems to differ more strongly between participants (approximately one-third of the total variance is due to between-subject factors). Although numerous studies have suggested a stable trait difference in the ability of self-control (de Ridder et al., 2012), our results provide the first quantitative information about the percentage of the inter-individual difference in comparison with other psychological states. Even though the absolute percentage estimation could have been inflated by response style differences between participants, the relative ranking suggests that inter-individual difference is more stable for self-control capacity than for core affect. In contrast, intra-individual day-to-day variation of self-control capacity is quite small in our data, despite common folk psychological beliefs (e.g., people often talk about having days in which they are able to control themselves better or worse, see Bergen, 2011). In fact, among all variables measured, state self-control components had the smallest day-to-day variations (all less than 9% of the total variance). Similarly, as estimated using the 4-item composite scale, self-control capacity also appears to have smaller intra-individual variation from moment to moment (around 15%), when compared to the momentary variation in core affect (30-40%). For the concentration and motivation components of self-control capacity, a small portion of their variations can be modeled as diurnal patterns. A typical student's ability to concentrate, and their motivation to pursue important goals, seems to increase after awakening, reach a peak at around 13:00 in the afternoon, and then decrease till sleep time. This mirrors the time-of-day pattern of learning performance in Randle et al. (2017), where performance also peaked early in the afternoon. Nonetheless, these time-of-day effects are only small parts of their temporal variations, which may otherwise be influenced by contextual factors (e.g., their physical and social environment, ongoing activities, etc.). Future experience sampling studies should attempt to measure contextual factors in people's daily lives.

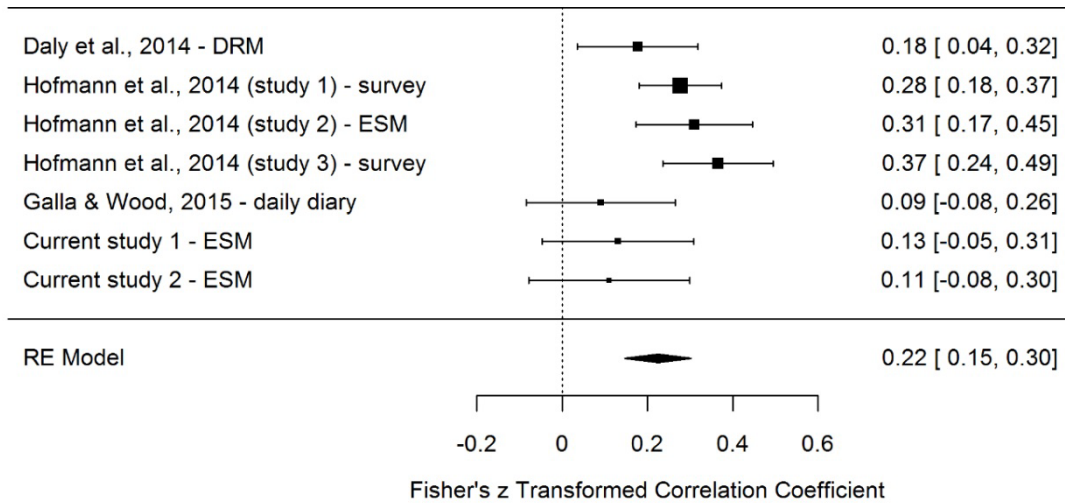


Figure 6.3 Forest plot of the correlation between trait self-control and affective valence. Note: Forest plot of the correlation between trait self-control and affective valence. The meta-analytic effect size and 95% CIs were estimated using a random-effect model. (DRM = day reconstruction method; ESM = experience sampling method).




Figure 6.4 Forest plot of the correlation between trait self-control and stress (tense arousal). Note: Forest plot of the correlation between trait self-control and stress (tense arousal). The meta-analytic effect size and 95% CIs were estimated using a random-effect model. (ESM = experience sampling method).

6.4.2 Relationship between self-control capacity and core affect

A second contribution of our studies is the systematic examination of the relationship between self-control and core affect. When self-control is measured at trait level, our results corroborate previous studies on the associations of trait self-control with affective valence and tense arousal (Bowlin & Baer, 2012; Daly et al., 2014; Galla & Wood, 2015; Hofmann et al., 2014). Although the effect on valence was estimated to be small and nonsignificant, we can reject a null-effect if we combine the data across the studies meta-analytically (see Figure 6.3). One reason for the smaller effect sizes in our studies ($r = 0.14$ and $r = 0.11$) compared with the ones in earlier studies (*meta-analytic* $r = 0.26$) might be that we used bipolar measures of valence while other studies mostly used unipolar scales selected from PANAS (Watson, Clark, & Tellegen, 1988). The bipolar items we used (e.g., happy-sad, pleasant-unpleasant) measure only the valence dimension, whereas the unipolar scales of positive affect sometimes also include items related to tense arousal, such as *calm*, which might have biased effect sizes in those studies. Regardless of the different conceptualizations and measurements of affective valence, when evaluated based on all available data, we can conclude that a small positive association between trait self-control and valence is reliable and robust in different populations. For tense arousal, our data provide the first clear evidence using an experience sampling design that people with higher trait self-control are likely to experience less psychological stress in daily lives, complementing evidence obtained using a one-time survey (Bowlin & Baer, 2012) and a daily diary method (Galla & Wood, 2015). The effect sizes of correlation between trait self-control and tense arousal are similar to those observed in earlier studies (see Figure 6.4).

In addition, by measuring state self-control capacity multiple times for multiple days, we could model the relationship between self-control capacity and core affect as both between-person and within-person correlations. The estimated between-person correlations between self-control capacity and core affect are in line with the results at the trait level, but with larger effect sizes (e.g., an increase of r from around 0.2 to 0.4 for sleepiness). The larger effect sizes may be explained by the fact that state measures of both self-control capacity and core affect are constrained by participants' experience in the week of the studies, while the trait self-control scale measures people's retrospective beliefs about their abilities based on broader past experience. However, as state measures of all variables were contingent to each other in each experience sampling session and were in similar formats, the between-person correlations might have been inflated by individual differences in response styles (e.g., some participants tended to rate all scales more extremely). When the between-person components and structural time-of-day effects are accounted for, the residual correlations at a within-person level reflect the genuine intra-individual relationship between self-control capacity and core affect. Intra-individually, it is clear that a person's self-control capacity changes together with their level of alertness and vitality. When someone is alert and energetic, they are likely better

able to concentrate, to control urges, and to pursue long-term goals. Surprisingly, results suggest no correlation between self-control capacity and tense arousal at the within-person level, even though inter-individually the correlation seems to be robust (see Figure 6.4). Despite the evidence that acute stress elicited experimentally can impede self-control process (Maier et al., 2015), the natural variation of stress and self-control capacity may not be strongly related in people's daily lives. This conclusion should be treated with caution, given the uncontrolled factors in our experience sampling design and the limited measurement reliability.

Another intriguing finding, for which we currently see no clear theoretical explanation, was that the motivation component of state self-control capacity did not correlate with concentration, willpower, and urge-control at between-person level, but did correlate strongly with these components intra-individually. According to self-control models from a dual-system perspective (Hofmann et al., 2009), it might be that the motivation item captured the dynamics of control motivation in the exertion cluster, whereas the willpower and urge-control items captured desire strength in the activation cluster. Moreover, our finding that motivation is distinct from willpower and urge-control seems to complement recent work that implies an overarching control role of motivation in self-control processes (Milyavskaya, Inzlicht, Hope, & Koestner, 2015). Given the seemingly deviating result for motivation, future work should particularly explore the role of motivation in self-control processes.

6.4.3 Implications for the conceptualization and measurement of self-control capacity

Our studies have important implications for the conceptualization of state self-control capacity, and how this construct should be measured. By positioning self-control capacity in a nomological network with dimensions of core affect, it becomes clear that the state variation of self-control capacity cannot be fully explained by the variation in core affect. This distinction is consistent with results from the ego-depletion paradigm in that changes in self-control capacity are generally not mirrored by changes in positive and negative affect (e.g., Muraven et al., 1998), and the distinction also extends to feelings of arousal and vitality. Although our results are far from resolving the theoretical debate about the causal role of self-control capacity, they do suggest that self-control capacity may provide additional explanatory power, above and beyond variation in core affect. This tentative conclusion could be strengthened if future studies using behavioral measures of self-control outcomes would conceptually replicate these results. For such studies to be successful, other potential determinants of self-control outcomes need to be carefully controlled, including task factors (e.g., attractiveness of immediate rewards) as well as control motivation as a different person factor (cf. Kotabe & Hofmann, 2015).

On the other hand, our results cast further doubt on the view that self-control capacity is a unified finite resource (see also Randle et al., 2017). At the very least, results suggest that the

concentration and motivation components of state self-control capacity are two sub-constructs that are distinct from a third sub-construct as measured by the willpower and urge-control items. With the temporal richness of experience sampling data, the distinctions are made based not only on the pattern of between-person correlations, but also on the pattern of within-person correlations and the temporal characteristics of the measured items. For future research, instead of treating state self-control capacity as a unified construct, a more fruitful approach might be to examine concrete sub-constructs of self-control capacity and to study how each of these sub-constructs influences self-control outcomes. For measuring the sub-constructs, we would recommend using 2-3 items for each construct based on our reliability analyses, in order to achieve sufficient within-person reliability. Following the network perspective of psychometrics, these sub-constructs can be treated as real psychological entities that are causally linked (e.g., bad concentration may make controlling urges more difficult), rather than as measurement items that are causally determined by one single construct called self-control capacity, as in the traditional latent variable approach (for network perspectives on other psychological constructs, see e.g., Costantini et al., 2015; Schmittmann et al., 2011).

6.4.4 Implications for digital lifestyle interventions

The results of decomposing variations in self-control capacity have some implications for digital lifestyle interventions. As our results corroborated previous findings of the large individual difference in self-control capacity, digital systems may personalize intervention techniques to users with high or low self-control capacity to increase intervention effectiveness. For within-person variations, it is less whether tailoring the timing of interventions to strong or weak “self-control moments” is beneficial, given the very small day-to-day variation and diurnal patterns of the state self-control measures. Future research is required to confirm the hypothetical associations between self-reported state self-control capacity and meaningful lifestyle behavioral measures (e.g., giving in to temptations, actual engagement in new healthy behaviors), and to test if tailoring to the small within-person variations is of any usefulness. At the very least, the relative independence of state self-control capacity from core affect justifies more research on applying the construct in digital lifestyle interventions.

Chapter 7

Epilogue

On one of those days busy writing my thesis, my attention was caught by a tweet from a psychology professor I followed: “Put away your Fitbit. You won’t walk any less and your wrist won’t smell funky”. It was a retweet commenting on a mobile health news titled “consumer fitness apps show nonsignificant behavior improvements”. The media coverage was based on a brand-new meta-analysis published in the *Journal of Medical Internet Research*, showing no evidence based on available randomized controlled trials for the benefits of using health apps to increase physical activity levels (Romeo et al., 2019). I saved the link and thought that the study was a nice addition to the long existing list of negative findings I referred to in the introduction chapter.

Given the apparent lack of evidence for the usefulness of digital intervention systems, it seems puzzling that there has been and still is an optimism towards the technology and that millions of self-tracking devices are sold. Of course, self-tracking devices and mobile apps work as pedometers (e.g., Case, Burwick, Volpp, & Patel, 2015), but their marketing has always been linked with promotions of better lifestyle and health, a benefit that is not fully delivered yet. There seems to be a wishful thinking by both producers and consumers that fills the gap between the steps being tracked and actual changes of behaviors. To draw on the distinction I made between physical technologies and behavioral technologies, a clear boundary between what works and what doesn’t work can always be applied to the former, and selling something that doesn’t work as advertised would be considered as unethical. In the infamous fraud of Theranos, the fake promises were eventually debunked cleanly. This is not so much true for behavioral technologies. Taking mindfulness or cognitive behavioral therapy as examples, despite the continuous doubts and criticisms over their efficacy, they continue to enjoy market success to some degree. The same applies to digital solutions for behavior change. Perhaps because the relatively weaker status of psychology compared to physical sciences, it is more difficult to evaluate these technologies rigorously. One might also suspect whether behavior change is so challenging that people are desperate to celebrate any new approaches that might work. Either way it feels critical to scrutinize the question why digital interventions would be a good solution for behavior change in theory by examining the problem in a broader context.

As an intro to this thesis, I quoted two paragraphs from Yuval Noah Harari's excellent book, *Sapiens: A Brief History of Humankind*, which contrasts the drastic differences between the environment people live in now and the one lived by Paleolithic humans. The same contrast can be used to explain why some unhealthy behaviors are so difficult to change from an evolutionary perspective (Tybur, Bryan, & Hooper, 2012). One popular example is the explanation of why most people prefer sweets and fast food (see Birch, 1999; Tybur & Griskevicius, 2013). In the stone-age when food were scarce, a preference for food with a lot of sugar and fat can be adaptive, as these nutrients provide more energy for surviving⁴¹. In today's society with abundant food and less physical work, such a selected genetic trait becomes maladaptive, or at least harmful for long-term health. More generally, it can be conjectured that humans are not evolved to remain healthy later in their lives, but to survive and have enough energy to reproduce at the right time. Our brains are not hardwired to consider long-term health as very important.

A lesson can also be learned from a narrower stretch of human history. While ruminating on why health behavior change is so hard, I started to suspect that the obsession with healthy lifestyle, such as low-carb diet or daily workout, is a very recent social phenomenon. A book on the public's health perception and awareness throughout different eras of human existence would be a fantastic book to read, but in its absence there are at least a few points to be considered.

First, never before has worrying about unhealthy lifestyles been so constantly salient at the level of individuals, since the concept of health is traditionally more of a social responsibility, and economical status largely constrains the quality of life one can lead. The general public has to take hold of what they can have, but they were not in the luxury position to choose between living healthily or unhealthily. Indeed, health promotion (the idea that individuals and groups should take care of their health) as a terminology was only coined in 1974 by the Canadian government (see the 1974 Lalonde report) and a decade later became recognized by the WHO.

Even at the level of society, there were many more pressing issues facing humanity than preventing chronic diseases in the past. Other threats, such as infectious disease, newborn mortality, poverty, and war, were more salient. Chronic diseases only became the main cause of mortality in the mid-20th century. The interests in health promotion and lifestyle intervention have perhaps also been fueled by scientific progress in fields of behavioral medicine and health psychology, which are only a few decades old.

⁴¹ Despite its popularity and plausibility, as with many other evolutionary accounts of behavior, I do not think it has been thoroughly tested.

Third, whereas religions have had a large impact on individuals' lives in the past, it is doubtful whether living a healthy life has ever been at the core of people's belief systems. It is nonetheless true that some religious practices played a positive role in promoting healthy behaviors, e.g., through prescribing a certain diet or prohibiting excessive drinking, though the focus was always on combating one's desires and doing what's right in a moral or godly sense. When the religious obligations are removed for many people today, it remains questionable whether the abstract concept of health can play the same role.

Fourth, beyond a mere consideration of health, having a healthy lifestyle is also a signaling of social status today. For example, going to gym regularly shows that one has time as well as a virtue of self-control and persistence. Restaurants that prepare natural, fresh, and healthy food are usually more expensive than their fast-food competitors. Such a perception has not always been the case. When foods are not abundant and heavy labor is a main form of physical exercise, being fat and pale were symbols of wealth and high social class. The old perception might be culturally universal, considering the corpulent and fleshy depiction of female beauty in both Western and Eastern traditional paintings. This is still an aesthetic ideal even in many areas of present-day Africa.

Overall, it is safe to say that compared with hardwired evolutionary preferences, the values attached to healthy lifestyles are more transient and perhaps more fragile. For the society as a whole, the newly established emphasis on self-driven lifestyle improvements certainly has great values for battling chronic diseases and for alleviating the burden on the healthcare system. However, as with many other social norms, although healthy living has become an important part of how individuals express themselves today, it may not be internalized enough to shape actual behaviors. As two sides of the self-control problem, evolutionary values and modern health concerns are not balanced in their strengths.

Having considered some reasons why lifestyle behavior change is so hard, it is time to realize that behavior change in general can also be very easy. Large-scale social changes and the associated changes of individual behaviors have frequently occurred throughout history. Often driven by technology progressions, changes in these scenarios are usually due to the availability of superior means to achieve certain goals, such as private cars for transportation and internet for acquiring information. Other times changes are simply forced by social institutions, in the sense that the individuals involved do not have a choice. One example is the food rationing by the UK government during the Second World War. Due to the special economic and social status, millions of British citizens adapted to the regulated diet overnight. Hard restrictions can also come from personal conditions, such as serious medical conditions that force one to avoid certain food or drinks. Instead of changing for a better long-term health, the immediate negative consequences or even life threats can quickly reshape behavior. Finally, a lesser restrictive but still powerful factor is the manipulation of monetary incentives,

such as the use of taxation to regulate consumer behavior. If not for other constraints, rewarding healthy behaviors financially would likely to be effective. In fact, a mobile app called SweatCoin, which I discovered in a behavior change conference, allows its users to convert steps to products and services (Derlyatka, Fomenko, Eck, Khmelev, & Elliott, 2019).

Of course, behavior change being difficult or easy is a human perception, but the underlying driving forces of human behaviors are the same. The situations where people find lifestyle behavior so hard to change are also the ones where people want the changes so much. The strong wanting for a healthy lifestyle comes from the Zeitgeist, which is also fueled by media exposure, social conversations, and even the digital gadgets around, but a different kind and often strong wanting for things that are more fundamentally pleasing. I believe that it is in those situations where we expect a form of behavioral technology on the basis of digital technologies to really help out for the motivated individuals. In many other situations the job is better to be left to society-level interventions, such as changes in incentive structure, policies, and commercial environment.

So why would digital intervention systems work for promoting healthy lifestyles? What is it special about digital technologies that would make what was not possible before possible? It is not that they can change the evolutionary forces, nor can they function in the same way as policies, regulations, and restrictive personal conditions. A more likely candidate is the pervasiveness and ubiquitousness of the technology. The fact that everyone carries digital devices with them means that if the technology works, it can be scaled up so easily and be a much more cost-effective solution than traditional intervention programs. However, the technology has to work first before the scaling up becomes relevant. The fact that people carry digital devices everywhere and at any time means that specific interventions can always be initiated, and machines can always stay alert and objective. But higher quantity does not necessarily mean higher quality. Compared with human coaches, the digital machines lack the persuasive power from a social presence⁴², as well as some high-level human capacities, such as empathy.

Machine intelligence, as being hyped to revolutionize various applications, is another factor that should contribute to lifestyle interventions. However, there are some points of caution to be made. In many artificial intelligence applications of today, machines are made to perform as well as humans (e.g., image recognition) or to beat humans in areas where humans are also doing very well (e.g., chess and Go). In contrast, how lifestyle behavior change or human psychology in general works is not even comprehensible by ourselves, so making machines to do the same task is of a very different nature. Nonetheless, machines are better than

⁴² This is not to say that machines cannot be perceived as social actors at all (see Reeves & Nass, 1996), but certainly to a much lesser degree than humans.

humans at extracting complex relationships from large amount of data and potentially making more accurate predictions.

This has naturally lead us to a feature of digital intervention systems that I regard as crucial for their (future) success in promoting lifestyle behavior change – the large amount of time-intensive data collected in people’s natural living environments. These data contain information about how people behave, feel, and even think, and how their bodies function. This data abundance has never happened in history before. If digital interventions have a chance to solve the health behavior change problem, the breakthrough has to do with data.

From data abundance, there are two routes to successful digital lifestyle interventions. The first route is the one of self-tracking, also known as quantified-self (see Lupton, 2016). Self-tracking can serve other purposes than behavior change, for example, to track one’s health status (e.g., trend of heart rate, blood pressure), or to assist planning (e.g., on which days of the week to do exercises), and at these tasks the technology works well. However, for behavior change, it relies on the following assumption: the reason why people do not behave healthy enough is that they do not know enough, so the new kind of data provided by the digital systems can lead to discoveries of knowledge and insights about oneself (Kersten-van Dijk et al., 2017; Li et al., 2011). As much as this route is promising and research continues to improve it, there is one caveat. It is not guaranteed that the digital data can lead to self-discovery. From an evolutionary perspective, because the complex and time-intensive data have only existed for a decade, our brains may not be adapted to deal with such data optimally, even if the data indeed contain valuable information. In addition, machines are much faster and more accurate in processing complex data than humans, and it has been shown that when making complex decisions, data-driven models often perform better than human experts (e.g., Ayres, 2008; Dawes, 1979; Silver, 2012).

The second route is the one I argued throughout this thesis. This route relies on the data abundance on one hand but also theory-based models about human behavior on the other, in order to exploit the dynamics of daily behavioral processes, including the cognitive processes of self-control and habit formation. Instead of letting the users learn about themselves from data, in this route the digital systems take the responsibility of learning, prediction, anticipation, and eventually guidance for actions. The future of this approach depends critically on the progress of behavioral and psychological sciences, and also on transforming theories into computational models that can be implemented in digital systems. This route, I call the psychological computing approach to lifestyle behavior change.

7.1 Contributions and insights of the thesis to the psychological computing approach

So what are the contributions and insights of this thesis, if one considers it as the first few steps towards the psychological computing approach? In the first half of the thesis (Chapter 2 to Chapter 4), we have realized a particular paradigm within the psychological computing approach, starting from a theoretical framework (Chapter 2), to a computational model of some processes in the framework (Chapter 3), and finally to an application of behavior prediction (Chapter 3). Here I shall mainly discuss the contributions of the research results and outputs from our paradigm to the psychological computing approach, and what we learned about the strengths and limitations of the approach. It should be noted that the psychological computing approach is much broader than the particular paradigm in this thesis. For example, the control-engineering model of self-efficacy (Riley et al., 2015) and the COMBI model (Klein et al., 2014) discussed in Chapter 1 can also be considered as examples of the same general approach, but with different focuses. Based on the stage model of behavior change, the COMBI model tailors its interventions on daily decisions according to the current stages of users and the identified behavioral determinants that limit their progressions (i.e., bottlenecks). Unlike our approach, because the identifications of the bottlenecks are based on periodical self-reports, the COMBI model does not compute rapid changes in cognitive and psychological states that underline repeated daily decisions. The control-engineering model of self-efficacy clearly focuses on a specific cognitive mechanism in behavior change and it uses a very different modeling language.

The adaptive decision-making framework is one of the most valuable outputs of the thesis, as it offers a much needed integration of traditional and more cutting-edge theories that focuses specifically on lifestyle behaviors and digital interventions. It also serves as a point of contact between psychologists, intervention designers, and modelers. We chose to transform some processes of the framework to simulation models using a sequential sampling approach, but the framework is open to computational models of other processes or using other modeling methods (e.g., differential equations, probabilistic graphical model). It can be the starting point of many more paradigms within the psychological computing approach.

A recurrent theme in this thesis is the use of sequential sampling models. It was used for explaining habit-goal integration (Chapter 3) and also for modeling dietary decisions during mouse-tracking tasks (Chapter 5). The beauty of the modeling approach is that, when being validated in one domain (e.g., instrumental learning), it can be extended, for example, by adding an option generation component (Chapter 4), to simulate results in a much broader behavioral space and to make new testable predictions. Although the sequential sampling approach has had success in many sub-fields of psychology, we are the first to introduce it to

behavior change research. As a result, there are many open questions regarding how the approach, besides its explanatory power, can be practically useful for digital interventions. It may be questioned, for example, what's the use of the low-level cognitive process of preference accumulation, if one is only interested in behavior-level prediction or intervention. Without any further proof, we consider a unique strength of sequential sampling models to be the disentanglement of three decision determinants, the habit-related starting position, the goal-related drift rate, and the decision threshold. With recent progresses in fitting decision field theory to actual data (e.g., Berkowitsch, Scheibehenne, & Rieskamp, 2014), it may become possible to estimate the values of these determinants from data about a user's repeated decisions, and then to find the suitable interventions accordingly.

In Chapter 4, we applied the habit formation part of the model in Chapter 3 in a realistic task of digital intervention systems – behavior prediction. The results provide some support for the value of using theory-based models for behavior prediction, although its advantage over simple data-driven models (i.e., based on past behavior) was not clear-cut. More generally, our work can be viewed as a concrete example how the psychological computing approach can combine theory-driven computational models and machine learning for applications, which resembles the use of domain knowledge in machine learning in many engineering problems. Essentially, what a theory-based equation does is defining a-priori relationships among some cognitive or behavioral variables, and allowing computations of some variables from others. When incorporated into machine learning models, the equation can supply unobservable features based on raw observed features or constrain the combination of features, in order to increase the efficiency of learning (e.g., less trial-and-error, or parameter tuning). Besides behavior prediction, it might also be valuable for systems to simply have knowledge about unobservable states of the users, although this has to be tested in future research.

The remaining part of the thesis (Chapter 5 & 6) turned to a specific issue with the psychological computing approach – how digital systems collect information about users' psychological and cognitive states that cannot be observed directly. Two methods, mouse-tracking and experience sampling, were explored because they can be relatively easily implemented in mobile computing devices. Our work of evaluating the mouse-tracking method highlights the challenges of using task-based indirect measures for practical applications. These tasks, including mouse-tracking, are usually used in laboratory settings where researchers are interested in showing differences between manipulated extreme conditions. Even small effects can have significant theoretical implications. However, applying them in behavior change interventions imposes a very different set of requirements, such as measurement reliability and sensitivity to the changes in the variables of interests.

7.2 Challenges and future research

There is still a long way to go from here, but the results and lessons I learned from the last 4 years already point to some future directions. First of all, while I considered data abundance as the main reason why digital interventions are promising for lifestyle behavior change, more data are needed in terms of variety and quality. When the concept of “big data” is talked about today, the bulk of behavioral data is generated through online behaviors, such as Facebook posts, Twitter conversations, and logged events in using mobile apps. Offline behavioral data that can be tracked by digital devices in the field are still relatively rare, and mostly are restricted to physical activities measured by accelerometers and motion sensors. The review paper by Gardner (2015) I cited in the last chapter perhaps illustrates the point nicely and somewhat ironically: in the tables summarizing all the reviewed studies, there was a dedicated column that codes the behavior measure in each study as self-report or objective, but in over 90 studies included only 7 employed sensor-based measures in the field⁴³. Except in one paper where an electronic monitoring pill bottle was used to measure medication adherence (Alison Phillips, Leventhal, & Leventhal, 2013), all others used accelerometers to measure physical activities. Another example for the lack of variety in sensor-based measures are the goal-compliance monitoring mechanisms used in GameBus, an excellent gamification app to promote healthy lifestyle developed by colleagues in our neighboring research group (see Shahrestani et al., 2017). GameBus users compete socially by complying with personalized lifestyle goals (e.g., daily steps, drinking water, sleep on time etc.), but all measures of compliance are self-reported except for physical activities, which can be measured by accelerometers in smartphones. Despite the current status, I do believe the future is bright in this regard, as sensor technologies and activity recognition algorithms continue to improve and the emerging trend of “internet of things” may also help.

Improving data quality can be more challenging. Often the large volume of data collected by self-tracking devices are not very useful from a psychological perspective. Data from digital systems can have very high temporal and spatial resolution, but often very poor psychological resolution. Considering data of daily step count as an example, numbers of steps can be monitored per hour or even minute, and associated location data can be available, but what meaningful behavioral units constitute the steps are often unknown. Ten-thousand steps a day could come from many different behaviors that serve very different goals, for example, commuting to work by foot, mandatory walking to meetings and toilets, and intentional walking exercises after lunch to improve one’s fitness. Improved activity recognition algorithms may help to segregate continuous movement data to isolated exercise events (walking, cycling,

⁴³ In three studies about snacking choice, the behavior measures were coded as objective, but those measures were one-time snacking choices made in the laboratory experiments.

etc.), but even if several walking episodes are detected, inferring a person's goals and intentions behind the episodes remains difficult. If a digital intervention aims at creating a habit of doing a walking exercise, or a study focuses on understanding goal-setting in behavior change, coding measured behaviors based on their cognitive driving forces is required. My impression is that for non-psychologists, the importance of the goals and intentions (or the lack thereof) behind behaviors is not fully appreciated. This is understandable because the aggregated behavioral patterns (e.g., daily step count) are indeed what counts for people's health, but to promote change behavioral data need higher psychological resolution.

Second, more computational models are highly needed, especially those that can be implemented and tested in digital interventions. It is important to point out that models do not have to be perfect to be useful (see Smaldino, 2017), and models of different levels of complexity and realism are needed. This need is also recognized by computer scientists and AI researchers, and collaborations between them and behavioral scientists can really help moving the field forward. Recently, I started a project with a machine-learning researcher from University of Amsterdam, who was interested in using reinforcement learning to optimize the notification function in digital intervention systems for promoting physical exercise. Considering a digital system as a learning agent, its rules for sending notifications under some constrained budget can be gradually optimized by observing the users' actual behavioral responses to the notifications as feedback (e.g., compliance as reward). What is also important in the project is to design a simple simulation model of human behavior – how users make exercise decisions depending on both external factors (e.g., time of day, weather) and whether notifications are received at the decision moments. Without gathering any empirical data, such a simulation model can be used to train the system using reinforcement learning, in order to reduce training time substantially. When real data are available, the differences between the learning outcomes in the simulated environment and the real-world can also tell us how to improve the computational model.

A crucial part of the adaptive decision framework that was not addressed in any depth in my PhD work are the reflection-level processes, including goal-setting, planning, and self-reflection, as well as their interactions with the action-level processes. My first personal advice for behavior change is always to set up a concrete and measurable goal and to keep the goal and goal-related progress always accessible. For example, in order to meet my goal of drinking at least 1 liter of pure water a day, my girlfriend bought me a nicely-designed bottle with exactly 1-liter of capacity and now I always see my progress in front of me on my working desk. My choice of not focusing on the reflection-level processes was due to research interests and chances, but one of the more logical reasons was that these flexible higher-level cognitive functions are more difficult to model computationally (i.e., lack of literature compared with self-control and habit formation) and to measure in the field. For example, we actually included a protocol of measuring self-reflective moments in an experience sampling study to

study when people reflected and what they reflected about in relation to their goals and what the consequences of these reflections had. It turned out that participants found it particularly hard to recall and report the details of such episodes (e.g., timing, frequency, and content), clearly more difficult than reporting momentary affects, thoughts, and feelings. Methodological innovations are required in the future to answer these very basic questions about the reflection-level processes.

Third, future research should explore different methods for measuring the cognitive and affective states included in computational models. The measured states can contribute not only to behavior predictions and the selection of intervention techniques, but also to the testing of the core mechanisms and other assumptions in the models. Advances in measurements are of crucial importance to any fields in behavioral sciences, but for research on behavior change and digital interventions, measurement developments should pay special attention to the applicability of new methods to digital intervention systems to be used in noisy real-world environments. Chapter 5 & 6 have shown how challenging this endeavor can be. The experience sampling method has proved to provide high-quality data and the method continues to be more accessible to researchers, but measures in most studies are still based on self-reports. Although self-reports of affective variables are easily administered and the data are usually accurate and reliable, in real digital intervention trials answering those questions does require a lot of effort from users and it may interfere with their natural behaviors. A promising alternative would be to estimate affective states from physiological data, such as heart rate, skin conductance, as in the approach of affective computing. Future research should also evaluate the possibility of transforming other indirect measures of cognitive states from laboratory settings to real-world settings, such as eye-tracking and neuroimaging. A lesson learned in our experience with mouse-tracking is that because ground truth for a cognitive measure is often lacking, computational models about how the associated cognitive processes work can contribute a lot to measurement developments.

7.3 Ethical considerations

The psychological computing approach may eventually lead us to a future where digital systems know more about us than we do ourselves. At one point, new ethical issues will emerge, so perhaps it is desirable to discuss some significant ethical implications of the technology even in this early stage of its development. A very important issue will of course be privacy. Threats to privacy in the digital age and how to take measures to relieve them have been discussed frequently in scholarly articles, books, and online media, but what's particularly interesting is how the systems envisioned here may redefine what private and personal information is. The systems may not only know about users' past and current behaviors, but also their future behaviors to some degree of accuracy based on the model predictions. Even in

2012, some fears for machine predictions were spread when a *Forbes* article entered the spotlight, describing how predictive models used by the supermarket Target could know a teenager was pregnant before her father did. Although it was later confirmed as a hypothetical story, one day digital intervention systems may actually predict, for example, what users would eat for lunch the next day, or how likely a particular user is going to relapse to a sedentary lifestyle. In addition to predictions, such systems may also be able to infer with confidence about users' cognitive and affective states that they themselves know very little about. It is unclear how users should decide when and when not to share this type of information if they do not know what the information entails (e.g., what does it mean if a mobile app knows my habit strength of eating crisps?). Perhaps, if such measurements are accurate enough, they should be considered as a form of "privacy of mind" as comparable to the classical notion of "privacy of body" (Westin, 1967).

Trust is another ethics-related issue that has to be addressed in the future. Obviously, users of digital intervention systems put themselves at some risks if they opt to let the machines figure out what's the best for them to change unhealthy lifestyles, in order to prevent or even manage chronic diseases. Among the situational factors that determine trust (e.g., Hurley, 2006), benevolence might be less of an issue assuming there is a good will behind the developments of these systems. But for the reasons that it is difficult to verify the accuracy of predictions and inferences by users themselves, their opinions about the competence of the systems may differ greatly. Like many other behavioral technologies, effectiveness of the systems may only be guaranteed at the level of groups rather than individuals, which complicates the process of developing trust by observing trustworthy behaviors from the systems. Also because users may have a hard time understanding how the systems work (e.g., its inner computational model), judging the integrity or predictability of the systems will not be easy. The systems may need to be able to explain its recommendations to the users to enhance trust.

Finally, compared with the route of self-tracking, the route of psychological computing may be criticized for undermining users' sense of autonomy in lifestyle behavior change (see Kamphorst & Kalis, 2015). I would defend the later approach for two reasons. First, at the action-level as in the adaptive decision framework, digital lifestyle interventions in most cases support decisions by providing and highlighting options but leave the final choice to the users⁴⁴. After decision-making, the actual actions (motor-control) to execute the behaviors are totally at the disposal of humans. Thus, the level of machine autonomy would mostly remain at a low level (e.g., level 3 in Sheridan & Verplank, 1978). Second, at the reflection-level, with the assistance from the digital applications on the how of goal-setting and planning, users would

⁴⁴ In this sense, nudging or other influences outside of users' awareness are exception, and there is indeed ongoing debate on whether nudging restricts people's freedom of choice or not (see Marchiori, Adriaanse, & de Ridder, 2017).

still be allowed to choose freely what specific behaviors in their lifestyles they would want to change, and this should enforce a sense of meaning. If one wishes to carefully examine this ethical issue, the adaptive decision framework proposed in this thesis could also be useful for differentiating the impacts of different intervention techniques on human autonomy based on the nature of the behavioral and cognitive processes being targeted.

7.4 Conclusion

For researchers interested in both psychology and technology, lifestyle behavior change is one of the most challenging yet rewarding problems to be tackled. The psychological computing approach proposed, as it relies on computational modeling of cognitive and behavioral processes, is applicable particularly to the challenges in lifestyle behaviors – the complexities in long-term habit formation and the self-control problem in the on-the-fly daily decisions. This thesis has taken some concrete steps towards the proposed approach through a collection of theoretical, computational, and empirical studies. With limitations and open questions remained, I wish it to stimulate similar future works and to reinforce a belief that psychological science and digital technologies must advance in symbiosis.

Bibliography

- Aarts, H., & Dijksterhuis, A. (2000). Habits as knowledge structures: Automaticity in goal-directed behavior. *Journal of Personality and Social Psychology, 78*, 53-63.
- Aarts, H., Verplanken, B., & Van Knippenberg, A. (1997). Habit and information use in travel mode choices. *Acta Psychologica, 96*, 1-14.
- Abraham, C., & Michie, S. (2008). A taxonomy of behavior change techniques used in interventions. *Health Psychology, 27*, 379-387.
- Adams, A. T., Costa, J., Jung, M. F., & Choudhury, T. (2015). Mindless computing: designing technologies to subtly influence behavior. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (pp. 719-730). ACM.
- Adams, C. D. (1982). Variations in the sensitivity of instrumental responding to reinforcer devaluation. *The Quarterly Journal of Experimental Psychology, 34*, 77-98.
- Adriaanse, M. A., Vinkers, C. D., De Ridder, D. T., Hox, J. J., & De Wit, J. B. (2011). Do implementation intentions help to eat a healthy diet? A systematic review and meta-analysis of the empirical evidence. *Appetite, 56*, 183-193.
- Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological Bulletin, 82*, 463-496.
- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes, 50*, 179-211.
- Ajzen, I. (2002). Residual effects of past on later behavior: Habituation and reasoned action perspectives. *Personality and Social Psychology Review, 6*, 107-122.
- Ajzen, I., & Madden, T. J. (1986). Prediction of goal-directed behavior: Attitudes, intentions, and perceived behavioral control. *Journal of Experimental Social Psychology, 22*, 453-474.
- Akaishi, R., Umeda, K., Nagase, A., & Sakai, K. (2014). Autonomous mechanism of internal choice estimate underlies decision inertia. *Neuron, 81*, 195-206.
- Åkerstedt, T., & Gillberg, M. (1990). Subjective and objective sleepiness in the active individual. *International Journal of Neuroscience, 52*, 29-37.
- Al-Durra, M., Torio, M. B., & Cafazzo, J. A. (2015). The use of behavior change theory in Internet-based asthma self-management interventions: a systematic review. *Journal of Medical Internet Research, 17*(4): e89.
- Allison, P. D. (1990). Change scores as dependent variables in regression analysis. *Sociological Methodology, 20*, 93-114.

- Alison Phillips, L., Leventhal, H., & Leventhal, E. A. (2013). Assessing theoretical predictors of long-term medication adherence: Patients' treatment-related beliefs, experiential feedback and habit development. *Psychology & Health, 28*, 1135-1151.
- Allom, V., Mullan, B., & Sebastian, J. (2013). Closing the intention-behaviour gap for sun-screen use and sun protection behaviours. *Psychology & Health, 28*, 477-494.
- Anderson, J. R. (1996). ACT: A simple theory of complex cognition. *American Psychologist, 51*, 355-365.
- Armstrong, A. W., Watson, A. J., Makredes, M., Frangos, J. E., Kimball, A. B., & Kvedar, J. C. (2009). Text-message reminders to improve sunscreen use: a randomized, controlled trial using electronic monitoring. *Archives of Dermatology, 145*, 1230-1236.
- Ayres, I. (2008). *Super crunchers: How anything can be predicted*. Hachette UK.
- Bagozzi, R. P. (1981). Attitudes, intentions, and behavior: A test of some key hypotheses. *Journal of Personality and Social Psychology, 41*, 607-627.
- Bandura, A. (1971). *Social learning theory*. Morristown.
- Bandura, A. (1982). Self-efficacy mechanism in human agency. *American Psychologist, 37*, 122-147.
- Bandura, A. (1989). Human agency in social cognitive theory. *American Psychologist, 44*, 1175-1184.
- Bandura, A., & McDonald, F. J. (1963). Influence of social reinforcement and the behavior of models in shaping children's moral judgment. *The Journal of Abnormal and Social Psychology, 67*, 274-281.
- Bandura, A., & Wood, R. (1989). Effect of perceived controllability and performance standards on self-regulation of complex decision making. *Journal of Personality and Social Psychology, 56*, 805-814.
- Bates, D., Maechler, M., Bolker B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*, 1-48.
- Bauer, U. E., Briss, P. A., Goodman, R. A., & Bowman, B. A. (2014). Prevention of chronic disease in the 21st century: elimination of the leading preventable causes of premature death and disability in the USA. *The Lancet, 384*, 45-52.
- Baumeister, R. F., & Heatherton, T. F. (1996). Self-regulation failure: An overview. *Psychological Inquiry, 7*, 1-15.
- Baumeister, R. F., Vohs, K. D., & Tice, D. M. (2007). The strength model of self-control. *Current Directions in Psychological Science, 16*, 351-355.

- Beck, L., & Ajzen, I. (1991). Predicting dishonest actions using the theory of planned behavior. *Journal of Research in Personality, 25*, 285-301.
- Bennett, J. E., Stevens, G. A., Mathers, C. D., Bonita, R., Rehm, J., Kruk, M. E., ... & Beagley, J. (2018). NCD Countdown 2030: worldwide trends in non-communicable disease mortality and progress towards Sustainable Development Goal target 3.4. *The Lancet, 392*, 1072-1088.
- Bergen, A. E. (2011). *Self-control in context: A mixed methods investigation of self-control and self-control failure*. PhD dissertation. The University of Guelph.
- Berkman, E. T., Hutcherson, C. A., Livingston, J. L., Kahn, L. E., & Inzlicht, M. (2017). Self-control as value-based choice. *Current Directions in Psychological Science, 26*, 422-428.
- Berkovsky, S., Coombe, M., Freyne, J., Bhandari, D., & Baghaei, N. (2010). Physical activity motivating games: virtual rewards for real activity. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 243-252). ACM.
- Berkowitsch, N. A., Scheibehenne, B., & Rieskamp, J. (2014). Rigorously testing multialternative decision field theory against random utility models. *Journal of Experimental Psychology: General, 143*, 1331-1348.
- Bhatia, S. (2013). Associations and the accumulation of preference. *Psychological Review, 120*, 522-543.
- Birch, L. L. (1999). Development of food preferences. *Annual Review of Nutrition, 19*, 41-62.
- Blechert, J., Meule, A., Busch, N. A., & Ohla, K. (2014). Food-pics: an image database for experimental research on eating and appetite. *Frontiers in Psychology, 5*:617.
- Bode, S., Sewell, D. K., Lilburn, S., Forte, J. D., Smith, P. L., & Stahl, J. (2012). Predicting perceptual decision biases from early brain activity. *Journal of Neuroscience, 32*, 12488-12498.
- Borsboom, D., Mellenbergh, G. J., & Van Heerden, J. (2003). The theoretical status of latent variables. *Psychological Review, 110*, 203-219.
- Bowlin, S. L., & Baer, R. A. (2012). Relationships between mindfulness, self-control, and psychological functioning. *Personality and Individual Differences, 52*, 411-415.
- Bussemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review, 100*, 432-459.
- Buc Calderon, C., Verguts, T., & Gevers, W. (2015). Losing the boundary: Cognition biases action well after action selection. *Journal of Experimental Psychology: General, 144*, 737-743.

- Bürkner, P. C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, *80*, 1-28.
- Calcagni, A., Lombardi, L., & Sulpizio, S. (2017). Analyzing spatial data from mouse tracker methodology: An entropic approach. *Behavior Research Methods*, *49*, 2012-2030.
- Carter, E. C., Kofler, L. M., Forster, D. E., & McCullough, M. E. (2015). A series of meta-analytic tests of the depletion effect: Self-control does not seem to rely on a limited resource. *Journal of Experimental Psychology: General*, *144*, 796-815.
- Carver, C. S., & Scheier, M. F. (1982). Control theory: A useful conceptual framework for personality–social, clinical, and health psychology. *Psychological Bulletin*, *92*, 111-135.
- Carver, C. S., & Scheier, M. F. (1990). Origins and functions of positive and negative affect: A control-process view. *Psychological Review*, *97*, 19-35.
- 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*, 319-333.
- Case, M. A., Burwick, H. A., Volpp, K. G., & Patel, M. S. (2015). Accuracy of smartphone applications and wearable devices for tracking physical activity data. *JAMA*, *313*, 625-626.
- Chajut, E., & Algom, D. (2003). Selective attention improves under stress: implications for theories of social cognition. *Journal of Personality and Social Psychology*, *85*, 231-248.
- Chapman, G. B. (1996). Temporal discounting and utility for health and money. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*, 771-791.
- Charbonnier, L., van Meer, F., van der Laan, L. N., Viergever, M. A., & Smeets, P. A. (2016). Standardized food images: a photographing protocol and image database. *Appetite*, *96*, 166-173.
- Choe, E. K., Lee, N. B., Lee, B., Pratt, W., & Kientz, J. A. (2014). Understanding quantified-selfers' practices in collecting and exploring personal data. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems* (pp. 1143-1152). ACM.
- Ciarocco, N. J., Twenge, J. M., Muraven, M., & Tice, D. M. (2015). The State Self-Control Capacity Scale: Reliability, validity, and correlations with physical and psychological stress. Manuscript under revision.
- Conroy, D. E., Maher, J. P., Elavsky, S., Hyde, A. L., & Doerksen, S. E. (2013). Sedentary behavior as a daily process regulated by habits and intentions. *Health Psychology*, *32*, 1149-1157.
- Conroy, D. E., Yang, C. H., & Maher, J. P. (2014). Behavior change techniques in top-ranked mobile apps for physical activity. *American Journal of Preventive Medicine*, *46*, 649-652.

- Costantini, G., Epskamp, S., Borsboom, D., Perugini, M., Mõttus, R., Waldorp, L. J., & Cramer, A. O. (2015). State of the aRt personality research: A tutorial on network analysis of personality data in R. *Journal of Research in Personality, 54*, 13-29.
- Cranford, J. A., Shrout, P. E., Iida, M., Rafaeli, E., Yip, T., & Bolger, N. (2006). A procedure for evaluating sensitivity to within-person change: Can mood measures in diary studies detect change reliably?. *Personality and Social Psychology Bulletin, 32*, 917-929.
- Crescioni, A. W., Ehrlinger, J., Alquist, J. L., Conlon, K. E., Baumeister, R. F., Schatschneider, C., & Dutton, G. R. (2011). High trait self-control predicts positive health behaviors and success in weight loss. *Journal of Health Psychology, 16*, 750-759.
- Cronbach, I. J., Gleser, G. C., Nanda, H., & Rajaratnam, N. (1972). *The dependability of behavioral measurements: Theory of generalizability for scores and profiles*. New York: John Wiley.
- Cronbach, L. J., & Meehl, P. E. (1955). Construct validity in psychological tests. *Psychological Bulletin, 52*, 281-302.
- Crutzen, R., & Peters, G. J. Y. (2018). Evolutionary learning processes as the foundation for behaviour change. *Health Psychology Review, 12*, 43-57.
- Cugelman, B. (2013). Gamification: what it is and why it matters to digital health behavior change developers. *JMIR Serious Games, 1*(1):e3.
- Curran, P. J., & Bauer, D. J. (2011). The disaggregation of within-person and between-person effects in longitudinal models of change. *Annual Review of Psychology, 62*, 583-619.
- Daly, M., Baumeister, R. F., Delaney, L., & MacLachlan, M. (2014). Self-control and its relation to emotions and psychobiology: Evidence from a Day Reconstruction Method study. *Journal of Behavioral Medicine, 37*, 81-93.
- Davis, R., Campbell, R., Hildon, Z., Hobbs, L., & Michie, S. (2015). Theories of behaviour and behaviour change across the social and behavioural sciences: a scoping review. *Health Psychology Review, 9*, 323-344.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience, 8*, 1704-1711.
- Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist, 34*, 571-582.
- de Bruijn, G. J., Gardner, B., van Osch, L., & Sniehotta, F. F. (2014). Predicting automaticity in exercise behaviour: the role of perceived behavioural control, affect, intention, action planning, and behaviour. *International journal of behavioral medicine, 21*, 767-774.

- de Bruijn, G. J., Kremers, S. P., De Vet, E., De Nooijer, J., Van Mechelen, W., & Brug, J. (2007). Does habit strength moderate the intention–behaviour relationship in the Theory of Planned Behaviour? The case of fruit consumption. *Psychology and Health, 22*, 899-916.
- de Ridder, D. T., Lensvelt-Mulders, G., Finkenauer, C., Stok, F. M., & Baumeister, R. F. (2012). Taking stock of self-control: A meta-analysis of how trait self-control relates to a wide range of behaviors. *Personality and Social Psychology Review, 16*, 76-99.
- de Vries, H. (2017). An integrated approach for understanding health behavior; the I-change model as an example. *Psychology and Behavioral Science International Journal, 2*(2), [555585], 10-19080.
- Derlyatka, A., Fomenko, O., Eck, F., Khmelev, E., & Elliott, M. T. (2019). Bright spots, physical activity investments that work: Sweatcoin: a steps generated virtual currency for sustained physical activity behaviour change. *British Journal of Sports Medicine*, Published Online First: 04 January 2019. doi: 10.1136/bjsports-2018-099739.
- DeRusso, A., Fan, D., Gupta, J., Shelest, O., Costa, R. M., & Yin, H. H. (2010). Instrumental uncertainty as a determinant of behavior under interval schedules of reinforcement. *Frontiers in Integrative Neuroscience, 4*:17.
- DeSmet, A., Shegog, R., Van Ryckeghem, D., Crombez, G., & De Bourdeaudhuij, I. (2015). A systematic review and meta-analysis of interventions for sexual health promotion involving serious digital games. *Games for Health Journal, 4*, 78-90.
- Deterding, S., Dixon, D., Khaled, R., & Nacke, L. (2011). From game design elements to gamefulness: defining gamification. In *Proceedings of the 15th international academic MindTrek conference: Envisioning future media environments* (pp. 9-15). ACM.
- Dickinson, A. (1985). Actions and habits: the development of behavioural autonomy. *Philosophical Transactions of the Royal Society B, 308*, 67-78.
- Dickinson, A., Nicholas, D. J., & Adams, C. D. (1983). The effect of the instrumental training contingency on susceptibility to reinforcer devaluation. *The Quarterly Journal of Experimental Psychology Section B, 35*, 35-51.
- Diederich, A. (2003a). MDFT account of decision making under time pressure. *Psychonomic Bulletin & Review, 10*, 157-166.
- Diederich, A. (2003b). Decision making under conflict: Decision time as a measure of conflict strength. *Psychonomic Bulletin & Review, 10*, 167-176.
- Diefenbach, S., & Müssig, A. (2019). Counterproductive effects of gamification: An analysis on the example of the gamified task manager Habitica. *International Journal of Human-Computer Studies, 127*, 190-210.

- Dillen, H., Phillips, J. G., & Meehan, J. W. (2005). Kinematic analysis of cursor trajectories controlled with a touchpad. *International Journal of Human-Computer Interaction, 19*, 223-239.
- Dogangün, A., Schwarz, M., Kloppenborg, K., & Le, R. (2017). An Approach to Improve Physical Activity by Generating Individual Implementation Intentions. In *Adjunct Publication of the 25th Conference on User Modeling, Adaptation and Personalization* (pp. 370-375). ACM.
- Doherty, A., Jackson, D., Hammerla, N., Plötz, T., Olivier, P., Granat, M. H., ... & Preece, S. J. (2017). Large scale population assessment of physical activity using wrist worn accelerometers: the UK biobank study. *PloS one, 12*, e0169649.
- Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron, 80*, 312-325.
- Dshemuchadse, M., Scherbaum, S., & Goschke, T. (2013). How decisions emerge: Action dynamics in intertemporal decision making. *Journal of Experimental Psychology: General, 142*, 93-100.
- Duckworth, A. L. (2011). The significance of self-control. *Proceedings of the National Academy of Sciences, 108*, 2639-2640.
- Duckworth, A. L., & Seligman, M. E. (2005). Self-discipline outdoes IQ in predicting academic performance of adolescents. *Psychological Science, 16*, 939-944.
- Dunton, G. F., & Atienza, A. A. (2009). The need for time-intensive information in healthful eating and physical activity research: a timely topic. *Journal of the American Dietetic Association, 109*, 30-35.
- Etherton, J. L., Osborne, R., Stephenson, K., Grace, M., Jones, C., & De Nadai, A. (2018). Bayesian analysis of multimethod ego-depletion studies favours the null hypothesis. *British Journal of Social Psychology, 57*, 368-385
- Evans, J. S. B. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology, 59*, 255-278.
- Ezzati, M., & Riboli, E. (2013). Behavioral and dietary risk factors for noncommunicable diseases. *New England Journal of Medicine, 369*, 954-964.
- Farrell, S., & Lewandowsky, S. (2010). Computational models as aids to better reasoning in psychology. *Current Directions in Psychological Science, 19*, 329-335.
- Fishbein, M., & Ajzen, I. (1975). *Belief, attitude, intention, and behavior: An introduction to theory and research*. Reading, MA: Addison-Wesley.
- Fleig, L., Pomp, S., Parschau, L., Barz, M., Lange, D., Schwarzer, R., & Lippke, S. (2013). From intentions via planning and behavior to physical exercise habits. *Psychology of Sport and Exercise, 14*, 632-639.

- Fontanesi, L., Gluth, S., Spektor, M. S., & Rieskamp, J. (2019). A reinforcement learning diffusion decision model for value-based decisions. *Psychonomic Bulletin & Review*. First online publication.
- Forstmann, B. U., Ratcliff, R., & Wagenmakers, E. J. (2016). Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annual Review of Psychology*, *67*, 614-666.
- Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., & Badre, D. (2015). fMRI and EEG predictors of dynamic decision parameters during human reinforcement learning. *Journal of Neuroscience*, *35*, 485-494.
- Frederick, S., Loewenstein, G., & O'donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, *40*, 351-401.
- Free, C., Phillips, G., Galli, L., Watson, L., Felix, L., Edwards, P., ... & Haines, A. (2013). The effectiveness of mobile-health technology-based health behaviour change or disease management interventions for health care consumers: a systematic review. *PLoS Medicine*, *10*(1):e1001362.
- Freeman, J. B. (2018). Doing psychological science by hand. *Current Directions in Psychological Science*, *27*, 315-323.
- Freeman, J. B., & Ambady, N. (2009). Motions of the hand expose the partial and parallel activation of stereotypes. *Psychological science*, *20*, 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. *Journal of Experimental Psychology: General*, *137*, 673-690.
- Freeman, J. B., & Ambady, N. (2010). MouseTracker: Software for studying real-time mental processing using a computer mouse-tracking method. *Behavior Research Methods*, *42*, 226-241.
- Galla, B. M., & Wood, J. J. (2015). Trait Self-Control Predicts Adolescents' Exposure and Reactivity to Daily Stressful Events. *Journal of Personality*, *83*, 69-83.
- Gainforth, H. L., West, R., & Michie, S. (2015). Assessing connections between behavior change theories using network analysis. *Annals of Behavioral Medicine*, *49*, 754-761.
- Gardner, B. (2012). Habit as automaticity, not frequency. *European Health Psychologist*, *14*, 32-36.
- Gardner, B. (2015). A review and analysis of the use of 'habit' in understanding, predicting and influencing health-related behaviour. *Health Psychology Review*, *9*, 277-295.

- Gardner, B., Abraham, C., Lally, P., & de Bruijn, G. J. (2012). Towards parsimony in habit measurement: Testing the convergent and predictive validity of an automaticity subscale of the Self-Report Habit Index. *International Journal of Behavioral Nutrition and Physical Activity*, 9:102.
- Gardner, B., & Lally, P. (2013). Does intrinsic motivation strengthen physical activity habit? Modeling relationships between self-determination, past behaviour, and habit strength. *Journal of Behavioral Medicine*, 36, 488-497.
- Gardner, B., Sheals, K., Wardle, J., & McGowan, L. (2014). Putting habit into practice, and practice into habit: a process evaluation and exploration of the acceptability of a habit-based dietary behaviour change intervention. *International Journal of Behavioral Nutrition and Physical Activity*, 11:135.
- Gershman, S. J., & Daw, N. D. (2017). Reinforcement learning and episodic memory in humans and animals: an integrative framework. *Annual Review of Psychology*, 68, 101-128.
- Gillebaart, M., Schneider, I. K., & De Ridder, D. T. (2016). Effects of trait self-control on response conflict about healthy and unhealthy food. *Journal of Personality*, 84, 789-798.
- Glymour, M. M., Weuve, J., Berkman, L. F., Kawachi, I., & Robins, J. M. (2005). When is baseline adjustment useful in analyses of change? An example with education and cognitive change. *American Journal of Epidemiology*, 162, 267-278.
- Gollwitzer, P. M. (1999). Implementation intentions: strong effects of simple plans. *American Psychologist*, 54, 493-503.
- Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin*, 130, 769-792.
- Guo, L. (2016). Quantified-Self 2.0: Using context-aware services for promoting gradual behaviour change. *arXiv preprint arXiv:1610.00460*.
- Ha, O. R., Bruce, A. S., Pruitt, S. W., Cherry, J. B. C., Smith, T. R., Burkart, D., ... & Lim, S. L. (2016). Healthy eating decisions require efficient dietary self-control in children: a mouse-tracking food decision study. *Appetite*, 105, 575-581.
- Hagger, M. S., Chatzisarantis, N. L. D., Alberts, H., Anggono, C. O., Batailler, C., Birt, A. R., ... Zwienerberg, M. (2016). A multilab preregistered replication of the ego-depletion effect. *Perspectives on Psychological Science*, 11, 546-573.
- Hagger, M. S., Wood, C., Stiff, C., & Chatzisarantis, N. L. (2010). Ego depletion and the strength model of self-control: a meta-analysis. *Psychological Bulletin*, 136, 495-525.
- Hall, P. A., & Fong, G. T. (2007). Temporal self-regulation theory: A model for individual health behavior. *Health Psychology Review*, 1, 6-52.

- Hare, T. A., Camerer, C. F., & Rangel, A. (2009). Self-control in decision-making involves modulation of the vmPFC valuation system. *Science*, *324*, 646-648.
- Harari, Y. N. (2015). *Sapiens: a brief history of humankind*. New York: Harper.
- Harris, A., Hare, T., & Rangel, A. (2013). Temporally dissociable mechanisms of self-control: early attentional filtering versus late value modulation. *Journal of Neuroscience*, *33*, 18917-18931.
- Hasler, B. P., Mehl, M. R., Bootzin, R. R., & Vazire, S. (2008). Preliminary evidence of diurnal rhythms in everyday behaviors associated with positive affect. *Journal of Research in Personality*, *42*, 1537-1546.
- Hastie, R. (2001). Problems for judgment and decision making. *Annual Review of Psychology*, *52*, 653-683.
- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. New York: Wiley.
- Helman, E., Stolier, R. M., & Freeman, J. B. (2015). Advanced mouse-tracking analytic techniques for enhancing psychological science. *Group Processes & Intergroup Relations*, *18*, 384-401.
- Hekler, E. B., Michie, S., Pavel, M., Rivera, D. E., Collins, L. M., Jimison, H. B., ... & Spruijt-Metz, D. (2016). Advancing models and theories for digital behavior change interventions. *American Journal of Preventive Medicine*, *51*, 825-832.
- Hermsen, S., Frost, J., Renes, R. J., & Kerkhof, P. (2016). Using feedback through digital technology to disrupt and change habitual behavior: a critical review of current literature. *Computers in Human Behavior*, *57*, 61-74.
- Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology*, *4*, 11-26.
- Hintzman, D. L. (1991). Why are formal models useful in psychology? In W. E. Hockley & S. Lewandowsky (Eds.), *Relating theory and data: Essays on human memory in honor of Bennet B. Murdock* (pp. 175-194). Hillsdale, NJ: Erlbaum.
- Heron, K. E., & Smyth, J. M. (2010). Ecological momentary interventions: incorporating mobile technology into psychosocial and health behaviour treatments. *British Journal of Health Psychology*, *15*, 1-39.
- Hofmann, W., Friese, M., & Strack, F. (2009). Impulse and self-control from a dual-systems perspective. *Perspectives on Psychological Science*, *4*, 162-176.
- Hofmann, W., Friese, M., & Wiers, R. W. (2008). Impulsive versus reflective influences on health behavior: A theoretical framework and empirical review. *Health Psychology Review*, *2*, 111-137.

- Hofmann, W., Gschwendner, T., Friese, M., Wiers, R. W., & Schmitt, M. (2008). Working memory capacity and self-regulatory behavior: Toward an individual differences perspective on behavior determination by automatic versus controlled processes. *Journal of Personality and Social Psychology*, *95*, 962-77.
- Hofmann, W., Luhmann, M., Fisher, R. R., Vohs, K. D., & Baumeister, R. F. (2014). Yes, but are they happy? Effects of trait self-control on affective well-being and life satisfaction. *Journal of Personality*, *82*, 265-277.
- Hofmann, W., & Patel, P. V. (2015). SurveySignal: A convenient solution for experience sampling research using participants' own smartphones. *Social Science Computer Review*, *33*, 235-253.
- Hull, C. L. (1943). *Principles of behavior: an introduction to behavior theory*. Oxford, England: Appleton-Century.
- Hur, J. D., Koo, M., & Hofmann, W. (2015). When temptations come alive: How anthropomorphism undermines self-control. *Journal of Consumer Research*, *42*, 340-358.
- Hurley, R. F. (2006). The decision to trust. *Harvard business review*, *84*, 55-62.
- IJsselsteijn, W., De Kort, Y., Midden, C., Eggen, B., & Van Den Hoven, E. (2006, May). Persuasive technology for human well-being: setting the scene. In *International conference on persuasive technology* (pp. 1-5). Springer, Berlin, Heidelberg.
- Intille, S. S., Kukla, C., Farzanfar, R., & Bakr, W. (2003). Just-in-time technology to encourage incremental, dietary behavior change. In *AMIA Annual Symposium Proceedings* (Vol. 2003, p. 874). American Medical Informatics Association.
- Inzlicht, M., Bartholow, B. D., & Hirsh, J. B. (2015). Emotional foundations of cognitive control. *Trends in Cognitive Sciences*, *19*, 126-132.
- Inzlicht, M., & Berkman, E. (2015). Six questions for the resource model of control (and some answers). *Social and Personality Psychology Compass*, *9*, 511-524.
- Inzlicht, M., & Schmeichel, B. J. (2012). What is ego depletion? Toward a mechanistic revision of the resource model of self-control. *Perspectives on Psychological Science*, *7*, 450-463.
- Inzlicht, M., Schmeichel, B. J., & Macrae, C. N. (2014). Why self-control seems (but may not be) limited. *Trends in Cognitive Sciences*, *18*, 127-133.
- Jaimes, L. G., Calderon, J., Lopez, J., & Raij, A. (2015). Trends in mobile cyber-physical systems for health just-in time interventions. In *SoutheastCon 2015* (pp. 1-6). IEEE.
- Jaimes, L. G., Llofriú, M., & Raij, A. (2014). A stress-free life: just-in-time interventions for stress via real-time forecasting and intervention adaptation. In *Proceedings of the 9th International Conference on Body Area Networks* (pp. 197-203). ICST.

- Jakicic, J. M., Davis, K. K., Rogers, R. J., King, W. C., Marcus, M. D., Helsel, D., ... & Belle, S. H. (2016). Effect of wearable technology combined with a lifestyle intervention on long-term weight loss: the IDEA randomized clinical trial. *JAMA*, *316*, 1161-1171.
- Janz, N. K., & Becker, M. H. (1984). The health belief model: A decade later. *Health Education Quarterly*, *11*, 1-47.
- Ji, M. F., & Wood, W. (2007). Purchase and consumption habits: Not necessarily what you intend. *Journal of Consumer Psychology*, *17*, 261-276.
- Johnson, E. J., Shu, S. B., Dellaert, B. G., Fox, C., Goldstein, D. G., Häubl, G., ... & Wansink, B. (2012). Beyond nudges: Tools of a choice architecture. *Marketing Letters*, *23*, 487-504.
- Johnson, P. C. (2014). Extension of Nakagawa & Schielzeth's R^2_{GLMM} to random slopes models. *Methods in Ecology and Evolution*, *5*, 944-946.
- Kable, J. W., & Glimcher, P. W. (2009). The neurobiology of decision: consensus and controversy. *Neuron*, *63*, 733-745.
- Kahneman, D. (2003). A perspective on judgment and choice: mapping bounded rationality. *American Psychologist*, *58*, 697-720.
- Kahneman, D., Krueger, A. B., Schkade, D. A., Schwarz, N., & Stone, A. A. (2004). A survey method for characterizing daily life experience: The day reconstruction method. *Science*, *306*, 1776-1780.
- Kamphorst, B., & Kalis, A. (2015). Why option generation matters for the design of autonomous e-coaching systems. *AI & SOCIETY*, *30*, 77-88.
- Kamphorst, B. A., Klein, M. C., & Van Wissen, A. (2014). Autonomous e-coaching in the wild: empirical validation of a model-based reasoning system. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems* (pp. 725-732). International Foundation for Autonomous Agents and Multiagent Systems.
- Kanfer, F. H., & Karoly, P. (1972). Self-control: A behavioristic excursion into the lion's den. *Behavior Therapy*, *3*, 398-416.
- Karoly, P. (1993). Mechanisms of self-regulation: A systems view. *Annual Review of Psychology*, *44*, 23-52.
- Kaplan, R. M., & Stone, A. A. (2013). Bringing the laboratory and clinic to the community: mobile technologies for health promotion and disease prevention. *Annual Review of Psychology*, *64*, 471-498.
- Keramati, M., Dezfouli, A., & Piray, P. (2011). Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS Computational Biology*, *7*, e1002055.

- Kersten-van Dijk, E. T., Westerink, J. H., Beute, F., & IJsselsteijn, W. A. (2017). Personal informatics, self-insight, and behavior change: A critical review of current literature. *Human-Computer Interaction, 32*, 268-296.
- Killcross, S., & Coutureau, E. (2003). Coordination of actions and habits in the medial prefrontal cortex of rats. *Cerebral Cortex, 13*, 400-408.
- Klasnja, P., & Pratt, W. (2012). Healthcare in the pocket: mapping the space of mobile-phone health interventions. *Journal of Biomedical Informatics, 45*, 184-198.
- Klein, M. C., Mogles, N., Treur, J., & Van Wissen, A. (2011). A computational model of habit learning to enable ambient support for lifestyle change. In *International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems* (pp. 130-142). Springer, Berlin, Heidelberg.
- Klein, M., Mogles, N., & Van Wissen, A. (2011). Why won't you do what's good for you? Using intelligent support for behavior change. In *International Workshop on Human Behavior Understanding* (pp. 104-115). Springer, Berlin, Heidelberg.
- Klein, M., Mogles, N., & van Wissen, A. (2014). Intelligent mobile support for therapy adherence and behavior change. *Journal of Biomedical Informatics, 51*, 137-151.
- Kohl, L. F., Crutzen, R., & de Vries, N. K. (2013). Online prevention aimed at lifestyle behaviors: a systematic review of reviews. *Journal of Medical Internet Research, 15*(7):e146.
- Koop, G. J. (2013). An assessment of the temporal dynamics of moral decisions. *Judgment and Decision Making, 8*, 527-539.
- Kosaki, Y., & Dickinson, A. (2010). Choice and contingency in the development of behavioral autonomy during instrumental conditioning. *Journal of Experimental Psychology: Animal Behavior Processes, 36*, 334-342.
- Kotabe, H. P., & Hofmann, W. (2015). On integrating the components of self-control. *Perspectives on Psychological Science, 10*, 618-638.
- Kurzban, R. (2012). *Why everyone (else) is a hypocrite: Evolution and the modular mind*. Princeton University Press.
- Lally, P., Van Jaarsveld, C. H., Potts, H. W., & Wardle, J. (2010). How are habits formed: Modelling habit formation in the real world. *European Journal of Social Psychology, 40*, 998-1009.
- Landis, D., Triandis, H. C., & Adamopoulos, J. (1978). Habit and behavioral intentions as predictors of social behavior. *The Journal of Social Psychology, 106*, 227-237.
- Ledger, D., & McCaffrey, D. (2014). Inside wearables: How the science of human behavior change offers the secret to long-term engagement. *Endeavour Partners, 200*, 1-17.

- Lefcheck, J. S. (2016). piecewiseSEM: Piecewise structural equation modelling in r for ecology, evolution, and systematics. *Methods in Ecology and Evolution*, 7, 573-579.
- Lee, M. K., Kiesler, S., & Forlizzi, J. (2011). Mining behavioral economics to design persuasive technology for healthy choices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 325-334). ACM.
- Lee, S. W., Shimojo, S., & O'Doherty, J. P. (2014). Neural computations underlying arbitration between model-based and model-free learning. *Neuron*, 81, 687-699.
- Lehto, T., & Oinas-Kukkonen, H. (2011). Persuasive features in web-based alcohol and smoking interventions: a systematic review of the literature. *Journal of Medical Internet Research*, 13(3): e46.
- Li, I., Dey, A. K., & Forlizzi, J. (2011). Understanding my data, myself: supporting self-reflection with ubicomp technologies. In *Proceedings of the 13th international conference on Ubiquitous computing* (pp. 405-414). ACM.
- Liljeholm, M., Dunne, S., & O'Doherty, J. P. (2015). Differentiating neural systems mediating the acquisition vs. expression of goal-directed and habitual behavioral control. *European Journal of Neuroscience*, 41, 1358-1371.
- Lim, S. L., Penrod, M. T., Ha, O. R., Bruce, J. M., & Bruce, A. S. (2018). Calorie labeling promotes dietary self-control by shifting the temporal dynamics of health-and taste-attribute integration in overweight individuals. *Psychological Science*, 29, 447-462.
- Locke, E. A., & Latham, G. P. (2002). Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *American Psychologist*, 57, 705-717.
- Loewenstein, G. (2000). Willpower: A decision-theorist's perspective. *Law and Philosophy*, 19, 51-76.
- Lord, F. M. (1967). A paradox in the interpretation of group comparisons. *Psychological Bulletin*, 68, 304-305.
- Luszczynska, A., Sobczyk, A., & Abraham, C. (2007). Planning to lose weight: Randomized controlled trial of an implementation intention prompt to enhance weight reduction among overweight and obese women. *Health Psychology*, 26, 507-512.
- Maier, S. U., Raja Beharelle, A., Polania, R., Ruff, C. C., & Hare, T. A. (2018). Dissociable mechanisms govern when and how strongly reward attributes affect decisions. *bioRxiv*, 434860.
- Maier, S. U., Makwana, A. B., & Hare, T. A. (2015). Acute stress impairs self-control in goal-directed choice by altering multiple functional connections within the brain's decision circuits. *Neuron*, 87, 621-631.

- Marchiori, D. R., Adriaanse, M. A., & De Ridder, D. T. (2017). Unresolved questions in nudging research: Putting the psychology back in nudging. *Social and Personality Psychology Compass*, *11*, e12297.
- Marcolino, M. S., Oliveira, J. A. Q., D'Agostino, M., Ribeiro, A. L., Alkmim, M. B. M., & Novillo-Ortiz, D. (2018). The impact of mHealth interventions: systematic review of systematic reviews. *JMIR mHealth and uHealth*, *6*, e23.
- Marewski, J. N., & Olsson, H. (2009). Beyond the null ritual: Formal modeling of psychological processes. *Zeitschrift für Psychologie*, *217*, 49-60.
- Marien, H., Custers, R., & Aarts, H. (2019). Studying Human Habits in Societal Context: Examining Support for a Basic Stimulus–Response Mechanism. *Current Directions in Psychological Science*. Advance online publication.
- Maris, E. (1998). Covariance adjustment versus gain scores—revisited. *Psychological Methods*, *3*, 309-327.
- Mateo, G. F., Granado-Font, E., Ferré-Grau, C., & Montaña-Carreras, X. (2015). Mobile phone apps to promote weight loss and increase physical activity: a systematic review and meta-analysis. *Journal of Medical Internet Research*, *17*(11): e253.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, *88*, 375-407.
- McDaniel, M. A., & Einstein, G. O. (2000). Strategic and automatic processes in prospective memory retrieval: A multiprocess framework. *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition*, *14*, S127-S144.
- Mercer, K., Li, M., Giangregorio, L., Burns, C., & Grindrod, K. (2016). Behavior change techniques present in wearable activity trackers: a critical analysis. *JMIR mHealth and uHealth*, *4*(2): e40.
- McKinstry, C., Dale, R., & Spivey, M. J. (2008). Action dynamics reveal parallel competition in decision making. *Psychological Science*, *19*, 22-24.
- Michie, S., Abraham, C., Whittington, C., McAteer, J., & Gupta, S. (2009). Effective techniques in healthy eating and physical activity interventions: a meta-regression. *Health Psychology*, *28*, 690-701.
- Michie, S., Richardson, M., Johnston, M., Abraham, C., Francis, J., Hardeman, W., ... & Wood, C. E. (2013). The behavior change technique taxonomy (v1) of 93 hierarchically clustered techniques: building an international consensus for the reporting of behavior change interventions. *Annals of Behavioral Medicine*, *46*, 81-95.

- Michie, S., & West, R. (2013). Behaviour change theory and evidence: a presentation to Government. *Health Psychology Review*, 7, 1-22.
- Michie, S. F., West, R., Campbell, R., Brown, J., & Gainforth, H. (2014). *ABC of behaviour change theories*. Silverback Publishing.
- Michie, S., West, R., & Spring, B. (2013). Moving from theory to practice and back in social and health psychology. *Health Psychology*, 32, 581-585.
- Miller, K. J., Ludvig, E. A., Pezzulo, G., & Shenhav, A. (2018). Realigning Models of Habitual and Goal-Directed Decision-Making. In *Goal-Directed Decision Making: Computations and Neural Circuits* (pp. 407-428). Academic Press.
- Miller, K. J., Shenhav, A., & Ludvig, E. A. (2019). Habits without values. *Psychological Review*. Advance online publication.
- Milyavskaya, M., Inzlicht, M., Hope, N., & Koestner, R. (2015). Saying “no” to temptation: Want-to motivation improves self-regulation by reducing temptation rather than by increasing self-control. *Journal of Personality and Social Psychology*, 109, 677-93.
- Mischel, W., Shoda, Y., & Peake, P. K. (1988). The nature of adolescent competencies predicted by preschool delay of gratification. *Journal of Personality and Social Psychology*, 54, 687- 696.
- Mischel, W., Shoda, Y., & Rodriguez, M. L. (1989). Delay of gratification in children. *Science*, 244, 933-938.
- Mittal, B. (1988). Achieving Higher Seat Belt Usage: The Role of Habit in Bridging the Attitude-Behavior Gap 1. *Journal of Applied Social Psychology*, 18, 993-1016.
- Moffitt, T. E., Arseneault, L., Belsky, D., Dickson, N., Hancox, R. J., Harrington, H., ... & Sears, M. R. (2011). A gradient of childhood self-control predicts health, wealth, and public safety. *Proceedings of the National Academy of Sciences*, 108, 2693-2698.
- Molenaar, P. C. (2004). A manifesto on psychology as idiographic science: Bringing the person back into scientific psychology, this time forever. *Measurement*, 2, 201-218.
- Motoki, K., Saito, T., Nouchi, R., Kawashima, R., & Sugiura, M. (2018). Tastiness but not healthfulness captures automatic visual attention: Preliminary evidence from an eye-tracking study. *Food Quality and Preference*, 64, 148-153.
- 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. *Journal of Neuroscience*, 32, 2335-2343
- Munson, S. A., & Consolvo, S. (2012). Exploring goal-setting, rewards, self-monitoring, and sharing to motivate physical activity. In *6th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*. IEEE.

- Muraven, M., Tice, D. M., & Baumeister, R. F. (1998). Self-control as a limited resource: Regulatory depletion patterns. *Journal of Personality and Social Psychology*, *74*, 774-789.
- Murtagh, S., Rowe, D. A., Elliott, M. A., McMinn, D., & Nelson, N. M. (2012). Predicting active school travel: the role of planned behavior and habit strength. *International Journal of Behavioral Nutrition and Physical Activity*, *9*:65.
- Murray, G., Nicholas, C. L., Kleiman, J., Dwyer, R., Carrington, M. J., Allen, N. B., & Trinder, J. (2009). Nature's clocks and human mood: The circadian system modulates reward motivation. *Emotion*, *9*, 705-716.
- Nacke, L. E., & Deterding, C. S. (2017). Editorial: The maturing of gamification research. *Computers in Human Behaviour*, *450-454*.
- Nahum-Shani, I., Hekler, E. B., & Spruijt-Metz, D. (2015). Building health behavior models to guide the development of just-in-time adaptive interventions: A pragmatic framework. *Health Psychology*, *34*(S), 1209-1219.
- Nakagawa, S., & Schielzeth, H. (2013). A general and simple method for obtaining R² from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, *4*, 133-142.
- Navon, D. (1984). Resources – A theoretical soup stone?. *Psychological Review*, *91*, 216-234.
- Neal, D. T., Wood, W., & Quinn, J. M. (2006). Habits—A repeat performance. *Current Directions in Psychological Science*, *15*, 198-202.
- Nilsen, W. J., & Pavel, M. (2013). Moving behavioral theories into the 21st century: technological advancements for improving quality of life. *IEEE Pulse*, *4*, 25-28.
- Noar, S. M., & Head, K. J. (2014). Mind the gap: bringing our theories in line with the empirical data—a response to commentaries. *Health Psychology Review*, *8*, 65-69.
- Norman D.A., Shallice T. (1986) Attention to Action. In: Davidson R.J., Schwartz G.E., Shapiro D. (Eds.) *Consciousness and Self-Regulation* (pp. 1-18). Springer, Boston, MA
- Norman, P. (2011). The theory of planned behavior and binge drinking among undergraduate students: Assessing the impact of habit strength. *Addictive Behaviors*, *36*, 502-507.
- Nour, M., Chen, J., & Allman-Farinelli, M. (2016). Efficacy and external validity of electronic and mobile phone-based interventions promoting vegetable intake in young adults: systematic review and meta-analysis. *Journal of Medical Internet Research*, *18*(4): e58.
- O'doherty, J. P., Cockburn, J., & Pauli, W. M. (2017). Learning, reward, and decision making. *Annual Review of Psychology*, *68*, 73-100.
- op den Akker, H. (2014). *Smart tailoring of real-time physical activity coaching systems*. PhD dissertation. Twente University Press.

- Oppenheimer, D. M., & Kelso, E. (2015). Information processing as a paradigm for decision making. *Annual Review of Psychology*, *66*, 277-294.
- Ouellette, J. A., & Wood, W. (1998). Habit and intention in everyday life: The multiple processes by which past behavior predicts future behavior. *Psychological Bulletin*, *124*, 54-74.
- Patel, M. S., Asch, D. A., & Volpp, K. G. (2015). Wearable devices as facilitators, not drivers, of health behavior change. *JAMA*, *313*, 459-460.
- Pauli, W. M., Cockburn, J., Pool, E. R., Pérez, O. D., & O'Doherty, J. P. (2018). Computational approaches to habits in a model-free world. *Current Opinion in Behavioral Sciences*, *20*, 104-109.
- Payne, H. E., Lister, C., West, J. H., & Bernhardt, J. M. (2015). Behavioral functionality of mobile apps in health interventions: a systematic review of the literature. *JMIR mHealth and uHealth*, *3*(1):e20.
- Pedersen, M. L., Frank, M. J., & Biele, G. (2017). The drift diffusion model as the choice rule in reinforcement learning. *Psychonomic Bulletin & Review*, *24*, 1234-1251.
- Pessiglione, M., Czernecki, V., Pillon, B., Dubois, B., Schüpbach, M., Agid, Y., & Tremblay, L. (2005). An effect of dopamine depletion on decision-making: the temporal coupling of deliberation and execution. *Journal of Cognitive Neuroscience*, *17*, 1886-1896.
- Pessoa, L. (2009). How do emotion and motivation direct executive control?. *Trends in Cognitive Sciences*, *13*, 160-166.
- Phillips, J. G., & Triggs, T. J. (2001). Characteristics of cursor trajectories controlled by the computer mouse. *Ergonomics*, *44*, 527-536.
- Picard, R. W. (2000). *Affective computing*. MIT press.
- Pina, L. R., Ramirez, E., & Griswold, W. G. (2012). Fitbit+: A behavior-based intervention system to reduce sedentary behavior. In *6th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth) and Workshops* (pp. 175-178). IEEE.
- Pinder, C., Vermeulen, J., Wicaksono, A., Beale, R., & Hendley, R. J. (2016). If this, then habit: exploring context-aware implementation intentions on smartphones. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct* (pp. 690-697). ACM.
- Platt, J. R. (1964). Strong inference. *Science*, *146*, 347-353.
- Podsakoff, P. M., MacKenzie, S. B., Lee, J. Y., & Podsakoff, N. P. (2003). Common method biases in behavioral research: A critical review of the literature and recommended remedies. *Journal of Applied Psychology*, *88*, 879-903.

- Postman, L. (1947). The history and present status of the law of effect. *Psychological Bulletin*, 44, 489-563.
- Powers, W. T. (1973). Feedback: Beyond Behaviorism: Stimulus-response laws are wholly predictable within a control-system model of behavioral organization. *Science*, 179, 351-356.
- Prestwich, A., Sniehotta, F. F., Whittington, C., Dombrowski, S. U., Rogers, L., & Michie, S. (2014). Does theory influence the effectiveness of health behavior interventions? Meta-analysis. *Health Psychology*, 33, 465-471.
- Prestwich, A., Webb, T. L., & Conner, M. (2015). Using theory to develop and test interventions to promote changes in health behaviour: evidence, issues, and recommendations. *Current Opinion in Psychology*, 5, 1-5.
- Prochaska, J. O., & DiClemente, C. C. (1982). Transtheoretical therapy: toward a more integrative model of change. *Psychotherapy: Theory, Research & Practice*, 19, 276-288
- Psarra, I. (2016). *A bounded rationality model of short and long-term dynamics of activity-travel behavior*. PhD dissertation. Eindhoven University of Technology.
- R Development Core Team (2008). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL <http://www.R-project.org>.
- Radha, M., Willemsen, M. C., Boerhof, M., & IJsselsteijn, W. A. (2016). Lifestyle recommendations for hypertension through Rasch-based feasibility modeling. In *Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization* (pp. 239-247). ACM.
- Rahman, T., Czerwinski, M., Gilad-Bachrach, R., & Johns, P. (2016). Predicting about-to-eat moments for just-in-time eating intervention. In *Proceedings of the 6th International Conference on Digital Health Conference* (pp. 141-150).
- Randles, D., Harlow, I., & Inzlicht, M. (2017). A pre-registered naturalistic observation of within domain mental fatigue and domain-general depletion of self-control. *PLoS One*, 12, e0182980.
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9, 545-556.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59-108.
- Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, 9, 347-356.
- Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods* (Vol. 1). Sage.

- Riley, W. T., Martin, C. A., Rivera, D. E., Hekler, E. B., Adams, M. A., Buman, M. P., ... & King, A. C. (2015). Development of a dynamic computational model of social cognitive theory. *Translational Behavioral Medicine, 6*, 483-495.
- Riley, W. T., Rivera, D. E., Atienza, A. A., Nilsen, W., Allison, S. M., & Mermelstein, R. (2011). Health behavior models in the age of mobile interventions: are our theories up to the task?. *Translational Behavioral Medicine, 1*, 53-71.
- Robinson, M. D., Schmeichel, B. J., & Inzlicht, M. (2010). A cognitive control perspective of self-control strength and its depletion. *Social and Personality Psychology Compass, 4*, 189-200.
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review, 108*, 370-392.
- Romeo, A., Edney, S., Plotnikoff, R., Curtis, R., Ryan, J., Sanders, I., ... & Maher, C. (2019). Can Smartphone Apps Increase Physical Activity? Systematic Review and Meta-Analysis. *Journal of Medical Internet Research, 21*, e12053.
- Rothman, A. J., Gollwitzer, P. M., Grant, A. M., Neal, D. T., Sheeran, P., & Wood, W. (2015). Hale and hearty policies: How psychological science can create and maintain healthy habits. *Perspectives on Psychological Science, 10*, 701-705.
- Russell, J. A. (2003). Core affect and the psychological construction of emotion. *Psychological Review, 110*, 145-172.
- Russell, J. A., & Barrett, L. F. (1999). Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant. *Journal of Personality and Social Psychology, 76*, 805-819.
- Russell, J. A., & Carroll, J. M. (1999). On the bipolarity of positive and negative affect. *Psychological Bulletin, 125*, 3-30.
- Russo, D. J., Van Roy, B., Kazerouni, A., Osband, I., & Wen, Z. (2018). A tutorial on Thompson sampling. *Foundations and Trends® in Machine Learning, 11*, 1-96.
- Rutjes, H., Willemsen, M. C., & IJsselsteijn, W. A. (2019, April). Beyond Behavior: The Coach's Perspective on Technology in Health Coaching. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (p. 670). ACM.
- Saranummi, N., Spruijt-Metz, D., Intille, S. S., Korhone, I., Nilsen, W. J., & Pavel, M. (2013). Moving the science of behavior change into the 21st century: novel solutions to prevent disease and promote health. *IEEE pulse, 4*, 22-24.
- Saunders, B., & Inzlicht, M. (2018). An appraisal framework to understand why negative affect is both good and bad for self-control. *PsyArXiv Preprints*.
- Savage, L. J. (1954). *The Foundations of Statistics*. New York, Wiley.

- Scherbaum, S., Dshemuchadse, M., Fischer, R., & Goschke, T. (2010). How decisions evolve: The temporal dynamics of action selection. *Cognition*, *115*, 407-416.
- Scherbaum, S., & Kieslich, P. J. (2018). Stuck at the starting line: How the starting procedure influences mouse-tracking data. *Behavior Research Methods*, *50* 1-14.
- Schimmack, U. (2003). Affect measurement in experience sampling research. *Journal of Happiness Studies*, *4*, 79-106.
- Schmittmann, V. D., Cramer, A. O., Waldorp, L. J., Epskamp, S., Kievit, R. A., & Borsboom, D. (2013). Deconstructing the construct: A network perspective on psychological phenomena. *New Ideas in Psychology*, *31*, 43-53.
- Schneider, C., Weinmann, M., & vom Brocke, J. (2018). Digital nudging: guiding online user choices through interface design. *Communications of the ACM*, *61*, 67-73.
- Shadlen, M. N., & Shohamy, D. (2016). Decision making and sequential sampling from memory. *Neuron*, *90*, 927-939.
- Shahrestani, A., Van Gorp, P., Le Blanc, P., Greidanus, F., de Groot, K., & Leermakers, J. (2017). Unified Health Gamification can significantly improve well-being in corporate environments. In *2017 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (pp. 4507-4511). IEEE.
- Sheridan, T. B., & Verplank, W. L. (1978). *Human and computer control of undersea teleoperators*. Massachusetts Inst of Tech Cambridge Man-Machine Systems Lab.
- Shrout, P. E., & Lane, S. P. (2011). Psychometrics. In M. R. Mehl, & T. A. Conner (Eds.), *Handbook of research methods for studying daily life* (pp. 302-320). New York, NY: Guilford Press.
- Silver, N. (2012). *The signal and the noise: why so many predictions fail--but some don't*. Penguin.
- Smaldino, P. E. (2017). Models are stupid, and we need more of them. In Vallacher, R. R., Read, S. J. & Nowak, A. (Eds.) *Computational Social Psychology*, 311-331. Routledge, New York.
- Smolders, K. C. H. J., De Kort, Y. A. W., & Van den Berg, S. M. (2013). Daytime light exposure and feelings of vitality: Results of a field study during regular weekdays. *Journal of Environmental Psychology*, *36*, 270-279.
- Song, J. H., & Nakayama, K. (2009). Hidden cognitive states revealed in choice reaching tasks. *Trends in Cognitive Sciences*, *13*, 360-366.
- Spivey, M. J., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences*, *102*, 10393-10398.

- Spruijt-Metz, D., Hekler, E., Saranummi, N., Intille, S., Korhonen, I., Nilsen, W., ... & Sanna, A. (2015). Building new computational models to support health behavior change and maintenance: new opportunities in behavioral research. *Translational Behavioral Medicine*, 5, 335-346.
- Stawarz, K. M. (2017). *Towards better medication adherence apps: Preventing forgetfulness by facilitating the formation of routine-based remembering strategies*. PhD dissertation, University College London.
- Stillman, P. E., Medvedev, D., & Ferguson, M. J. (2017). Resisting Temptation: Tracking How Self-Control Conflicts Are Successfully Resolved in Real Time. *Psychological Science*, 28, 1240-1258.
- Stillman, P. E., Shen, X., & Ferguson, M. J. (2018). How mouse-tracking can advance social cognitive theory. *Trends in Cognitive Sciences*, 22, 531-543.
- Stone, A. A., Schwartz, J. E., Schkade, D., Schwarz, N., Krueger, A., & Kahneman, D. (2006). A population approach to the study of emotion: Diurnal rhythms of a working day examined with the Day Reconstruction Method. *Emotion*, 6, 134-139.
- Story, G., Vlaev, I., Seymour, B., Darzi, A., & Dolan, R. (2014). Does temporal discounting explain unhealthy behavior? A systematic review and reinforcement learning perspective. *Frontiers in Behavioral Neuroscience*, 8:76.
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, 8, 220-247.
- Strauss, M. E., & Smith, G. T. (2009). Construct validity: Advances in theory and methodology. *Annual Review of Clinical Psychology*, 5, 1-25.
- Sullivan, N. J., & Huettel, S. A. (2018). Dietary self-control depends on the latency and rate of information accumulation. *bioRxiv*, 465393.
- 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. *Psychological Science*, 26, 122-134.
- Sullivan, N. J., Hutcherson, C., Harris, A. M., & Rangel, A. (2019). A response to Zhang et al.(2018), "Can Mouse-tracking Reveal Attribute Processing Speeds in Dietary Self-control? Commentary on Sullivan et al. (2015) and Lim et al. (2018)". *bioRxiv*, 572974.
- Summerfield, C., & Tsetsos, K. (2012). Building bridges between perceptual and economic decision-making: neural and computational mechanisms. *Frontiers in Neuroscience*, 6, 70.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction* (Vol. 1, No. 1). Cambridge: MIT press.

- Szaszi, B., Palfi, B., Szollosi, A., Kieslich, P. J., & Aczel, B. (2018). Thinking dynamics and individual differences: Mouse-tracking analysis of the denominator neglect task. *Judgment and Decision Making, 13*, 23-32.
- Tangney, J. P., Baumeister, R. F., & Boone, A. L. (2004). High self-control predicts good adjustment, less pathology, better grades, and interpersonal success. *Journal of Personality, 72*, 271-324.
- Thai, S. & Page-Gould, E. (2018). ExperienceSampler: An open-source scaffold for building smartphone apps for experience sampling. *Psychological Methods, 23*, 729-739.
- Thakkar, J., Kurup, R., Laba, T. L., Santo, K., Thiagalingam, A., Rodgers, A., ... & Chow, C. K. (2016). Mobile telephone text messaging for medication adherence in chronic disease: a meta-analysis. *JAMA Internal Medicine, 176*, 340-349.
- Thaler, R., & Sunstein, C. (2008). *Nudge: The gentle power of choice architecture*. New Haven, Conn.: Yale.
- Thomas, J. G., & Bond, D. S. (2015). Behavioral response to a just-in-time adaptive intervention (JITAI) to reduce sedentary behavior in obese adults: Implications for JITAI optimization. *Health Psychology, 34*(S), 1261-1267.
- Thomas, M., Sing, H., Belenky, G., Holcomb, H., Mayberg, H., Dannals, R., ... & Welsh, A. (2000). Neural basis of alertness and cognitive performance impairments during sleepiness. I. Effects of 24 h of sleep deprivation on waking human regional brain activity. *Journal of Sleep Research, 9*, 335-352.
- Thorndike, E. L. (1932). *The fundamentals of learning*. New York, NY, US: Teachers College Bureau of Publications.
- Tobias, R. (2009). Changing behavior by memory aids: a social psychological model of prospective memory and habit development tested with dynamic field data. *Psychological Review, 116*, 408-438.
- Triandis, H. C. (1977). *Interpersonal behavior*. Brooks/Cole Pub. Co..
- Tricomi, E., Balleine, B. W., & O'Doherty, J. P. (2009). A specific role for posterior dorsolateral striatum in human habit learning. *European Journal of Neuroscience, 29*, 2225-2232.
- Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological Review, 117*, 440-463.
- Trueblood, J. S., Brown, S. D., & Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological Review, 121*, 179-205.

- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, *79*, 281-299.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*, 453-458.
- Tybur, J. M., Bryan, A. D., & Hooper, A. E. C. (2012). An evolutionary perspective on health psychology: New approaches and applications. *Evolutionary Psychology*, *10*, 855-867.
- Tybur, J. M., & Griskevicius, V. (2013). Evolutionary psychology: a fresh perspective for understanding and changing problematic behavior. *Public Administration Review*, *73*, 12-22.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, *108*, 550-592.
- Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, *111*, 757-769.
- Usher, M., Olami, Z., & McClelland, J. L. (2002). Hick's law in a stochastic race model with speed-accuracy tradeoff. *Journal of Mathematical Psychology*, *46*, 704-715.
- van Bree, R. J., van Stralen, M. M., Mudde, A. N., Bolman, C., de Vries, H., & Lechner, L. (2015). Habit as mediator of the relationship between prior and later physical activity: A longitudinal study in older adults. *Psychology of Sport and Exercise*, *19*, 95-102.
- van Breukelen, G. J. (2006). ANCOVA versus change from baseline had more power in randomized studies and more bias in nonrandomized studies. *Journal of Clinical Epidemiology*, *59*, 920-925.
- van der Linden, D., Frese, M., & Meijman, T. F. (2003). Mental fatigue and the control of cognitive processes: effects on perseveration and planning. *Acta Psychologica*, *113*, 45-65.
- van Dijk, E. T., Westerink, J. H., Beute, F., & IJsselsteijn, W. A. (2015). In sync: the effect of physiology feedback on the match between heart rate and self-reported stress. *BioMed Research International*, *2015*, 1-9.
- Verplanken, B., Aarts, H., van Knippenberg, A., & van Knippenberg, C. (1994). Attitude versus general habit: Antecedents of travel mode choice. *Journal of Applied Social Psychology*, *24*, 285-300.
- Verplanken, B., Aarts, H., & Van Knippenberg, A. (1997). Habit, information acquisition, and the process of making travel mode choices. *European Journal of Social Psychology*, *27*, 539-560.
- Verplanken, B., & Melkevik, O. (2008). Predicting habit: The case of physical exercise. *Psychology of Sport and Exercise*, *9*, 15-26.

- Verplanken, B., & Orbell, S. (2003). Reflections on past behavior: a self-report index of habit strength 1. *Journal of Applied Social Psychology, 33*, 1313-1330.
- Vlaev, I., & Dolan, P. (2015). Action change theory: A reinforcement learning perspective on behavior change. *Review of General Psychology, 19*, 69-95.
- von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behavior*, 2nd rev. Princeton, NJ, US: Princeton University Press.
- Walker, I., Thomas, G. O., & Verplanken, B. (2015). Old habits die hard: Travel habit formation and decay during an office relocation. *Environment and Behavior, 47*, 1089-1106.
- Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology, 54*, 1063-1070.
- Webb, N. M., Shavelson, R. J., & Haertel, E. H. (2006). 4 reliability coefficients and generalizability theory. *Handbook of Statistics, 26*, 81-124.
- Webb, T. L., Joseph, J., Yardley, L., & Michie, S. (2010). Using the internet to promote health behavior change: a systematic review and meta-analysis of the impact of theoretical basis, use of behavior change techniques, and mode of delivery on efficacy. *Journal of Medical Internet Research, 12*(1): e4.
- West, R. (2006). *Theory of addiction*. Oxford: Blackwell.
- Westin, A (1967). *Privacy and freedom*. New York: Atheneum.
- Wiedemann, A. U., Gardner, B., Knoll, N., & Burkert, S. (2014). Intrinsic rewards, fruit and vegetable consumption, and habit strength: A three-wave study testing the associative-cybernetic model. *Applied Psychology: Health and Well-Being, 6*, 119-134.
- Wilt, J., Funkhouser, K., & Revelle, W. (2011). The dynamic relationships of affective synchrony to perceptions of situations. *Journal of Research in Personality, 45*, 309-321.
- Wirth, R., Pfister, R., & Kunde, W. (2016). Asymmetric transfer effects between cognitive and affective task disturbances. *Cognition and Emotion, 30*, 399-416.
- Wood, C., & Magnello, M. E. (1992). Diurnal changes in perceptions of energy and mood. *Journal of the Royal Society of Medicine, 85*, 191-194.
- Wood, W., & Neal, D. T. (2007). A new look at habits and the habit-goal interface. *Psychological Review, 114*, 843-863.
- Wood, W., & Rünger, D. (2016). Psychology of habit. *Annual Review of Psychology, 67*, 289-314.

- Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nature Reviews Neuroscience*, 7, 464-476.
- Yin, H. H., Knowlton, B. J., & Balleine, B. W. (2004). Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *European Journal of Neuroscience*, 19, 181-189.
- Yin, H. H., Knowlton, B. J., & Balleine, B. W. (2005). Blockade of NMDA receptors in the dorsomedial striatum prevents action–outcome learning in instrumental conditioning. *European Journal of Neuroscience*, 22, 505-512.
- Zhang, C., Smolders, K. C., Lakens, D., & IJsselsteijn, W. A. (2018). Two experience sampling studies examining the variation of self-control capacity and its relationship with core affect in daily life. *Journal of Research in Personality*, 74, 102-113.
- Zhang, C., Starczewski, A. P., Lakens, D., & IJsselsteijn, W. A. (2018). A decision-making perspective on coaching behavior change: a field experiment on promoting exercise at work. In: Ham J., Karapanos E., Morita P., Burns C. (Eds.) *Persuasive Technology. PERSUASIVE 2018. Lecture Notes in Computer Science*, vol 10809. Springer, Cham.
- Zhang, C., van Wissen, A., Lakens, D., Vanschoren, J., De Ruyter, B., & IJsselsteijn, W. A. (2016, September). Anticipating habit formation: a psychological computing approach to behavior change support. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct* (pp. 1247-1254). ACM.
- Zhang, C., Willemsen, M., & Lakens, D. (2018). Can Mouse-tracking Reveal Attribute Processing Speeds in Dietary Self-control? Commentary on Sullivan et al. (2015) and Lim et al. (2018) with A Simulation Study. *PsyArXiv*. doi: 10.31234/osf.io/725vp
- Zhao, J., Freeman, B., & Li, M. (2016). Can mobile phone apps influence people’s health behavior change? An evidence review. *Journal of Medical Internet Research*, 18(11): e287.
- Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Profile Books.
- Zuckerman, O., & Gal-Oz, A. (2014). Deconstructing gamification: evaluating the effectiveness of continuous measurement, virtual rewards, and social comparison for promoting physical activity. *Personal and Ubiquitous Computing*, 18, 1705-1719.

Summary

Towards a Psychological Computing Approach to Digital Lifestyle Interventions

The rapid development of digital technologies has led to the belief that digital lifestyle intervention systems can alleviate the global problems of unhealthy lifestyles and chronic diseases. Yet, the lack of long-term effectiveness and user engagement suggests that the research and development of such systems are still in the early stage. Many researchers advocated the importance of applying psychological theories in digital lifestyle interventions, but the realization of this proposal is hindered by a lack of theory integration and the limitations of traditional behavior change theories. Based on the recent development in computational modeling of human behavior, a psychological computing approach was proposed, in which theory-based computational models are implemented in digital systems to predict and intervene in lifestyle behaviors. Research reported in this thesis takes the first few steps to bridge the gap between psychological theories and digital lifestyle interventions, and to move towards a psychological computing approach.

In Chapter 2, the psychological computing approach is introduced by developing an integrated theoretical framework of lifestyle behavior change, called the *adaptive decision-making framework*. This was done by reviewing relevant individual theories of learning and decision-making in the psychology literature, and then integrating the theoretical ideas into a two-level representation of lifestyle behaviors, including both daily decisions and periodical reflections. Common digital intervention techniques were mapped to the framework. The framework offers a much needed theory integration that is dynamic and matches with the temporal granularity of digital data and interventions. It can be used by intervention designers to select theory-based behavior change techniques, and by scientists as a scaffold to develop computational models of behavior change. The chapter also identified habit formation and self-control as the two main topics for the rest of the thesis.

Following a sequential sampling approach, Chapter 3 transformed the action-level decision-making and learning processes in the adaptive decision framework into a computational model. The model focused on explaining how habits and goals interact with each other to influence value-based decisions. Through three simulation studies, it was shown that the model could reproduce important laboratory findings from the literature of habit-goal interaction, including effects in the devaluation and reversal learning paradigms, and predict gradual changes in decision time. Our model challenges the current approach of arbitration models and provides a more parsimonious explanation for habit-goal conflicts. By incorporating the process of option generation, our approach can be extended to model habit formation in daily environments.

In Chapter 4, the habit formation part of the computational model developed in Chapter 3 was evaluated in two real-world digital intervention trials of changing dental behavior. Data collected in two studies were used to understand the reciprocal relationship between habit strength, attitude, and behavior, and also to test whether habit strength computed by the computational model could improve behavior prediction. Statistical analyses showed that habit strength influenced actual toothbrushing behavior besides attitude, but its moderation on the attitude-behavior relationship was only found in Study 2. Predictive modeling results suggested that models based on computed habit strength predicted behavior better than models based on self-reported behavioral determinants, and in Study 2 also better than the models based on past behavior rate. In line with the psychological computing approach, the theory-based computational model of habit formation has the potential to be implemented in digital systems to predict and intervene in lifestyle behaviors.

Chapters 5 and 6 explored two methods of understanding self-control in daily lifestyle decisions. Through two laboratory food-choice experiments, Chapter 5 evaluated whether a mouse-tracking technique could be used in digital systems to measure users' decision conflicts experienced in self-control dilemmas. Results showed that while the mouse-tracking paradigm could be transferred from a desktop setting to touch-screen devices, the correlations between mouse-tracking parameters and the decision-conflict strength were too small to be practically useful. Combined with simulation studies, the chapter also helped to scrutinize an existing method of using mouse-tracking data to reveal the cognitive mechanisms underlying self-control, and to clarify the plausibility of various theoretical assumptions associated with the mouse-tracking paradigm.

In Chapter 6, the variations of self-control capacity in people's daily lives are studied using an experience sampling method. Data from two field studies were used to understand how self-control capacity varied inter- and intra-individually and how these variations were related to changes in people's affective states. Results showed that variations in self-control capacity were attributed largely to individual differences, slightly to diurnal patterns, but very little to day-to-day changes. State self-control capacity correlated more with affective valence and tense arousal inter-individually, but more with alertness and energetic arousal intra-individually. We also discussed the challenges of defining and measuring self-control capacity as a unified construct and the implications of the results for designing lifestyle interventions.

Finally, in the concluding chapter, the relevance of the psychological computing approach to the problem of behavior change was positioned in a larger context, and discussed from both an evolutionary and a historical perspective. Contributions of the thesis, future research directions, and ethical issues were discussed with reference to the psychological computing approach.

Acknowledgements

I have always been wondering what's it like to complete a PhD and now I know it! Looking back from where I am now, it almost feels like a miracle that a person could work on a single project with all the commitment and dedication for more than 4 years. The journey was a tremendous one, and was never short of excitements, frustrations, and pushes for personal limits. This thesis would not be possible without all the people who accompanied me in one way or another.

Wijnand, thank you for everything. I am so grateful that you trusted me and gave me the job at the first place. Thank you also for always supporting my research decisions and helping me to become an independent scientist. Certainly sometimes I would wish to have more discussions with you, which are always motivating and inspiring, but I guess meeting my promoter once or even twice a month was already exceptional. Thank you for giving me your valuable time. You are one of the most resourceful and warm-hearted person I know.

Daniel, it was wonderful to have you on my side. Of course, I have known you long before the PhD and I still really miss the days when we could discuss many crazy experiment ideas back in 2011. You have certainly become busier and busier over the years, but I am grateful that you are always there when I need you. You supported me not only on specific research problems, but also on a personal level how to live with the ups and downs in the journey. You are always a role-model for me as a researcher.

Joaquin, thank you for your intellectual inputs and kind supports. Even though the project did not go to the direction that would best match your interests, the overall theme of the thesis and a key chapter were in debt to my conversations with you. Your immense knowledge on machine learning was always something I could rely on and talking to someone from a very different background was always refreshing experience.

Arlette, thank you so much for accepting to be my supervisor and taking care of me at Philips in the last three years or so. The amount of time you spent on me could easily qualify you as my daily supervisor! In terms of research, you helped so much to sharpen my thinking on habit formation and my thesis would be nothing like what it is now if I did not read your thesis. I hope to continue our collaborations for the years to come.

I would also like to thank some other colleagues I worked with at Philips Research. Ron, it was my privilege to have you in my team for a year. It is a pity that the close connection was so short, because you are one of the smartest person I ever met. You really helped me to delineate my research questions. Dietwig, thanks for taking care of me when Arlette was away and warmly inviting me to the Human Modeling research group. Boris, thank you for helping me to get started at Philips. Although Arlette took your role eventually, I am grateful to have your inputs from the early meetings and I am always impressed by your witty style. And

Hanne, thank you for initiating the nice collaboration. It was a difficult time for me physically, but I enjoyed working with you in those late hours unpacking design toothbrushes.

And Martijn, my 5th supervisor, thank you for your help and your loyalty to be in all those monthly meetings. I really enjoyed working with you on the topic of mouse-tracking and decision-making models. Although our intense discussions haven't led to any output yet, I am sure they will. Thank you also for involving me in supervising Master theses.

I am also thankful for all the members in my defense committee – Denise de Ridder, Panos Markopoulos, Wilhelm Hofmann, and Rik Crutzen, who spent their precious time to read my thesis in full and shared their thoughtful comments. Thanks to professor Cees Midden for accepting to be the chair of my defense.

Of course, the PhD journey was not only about doing research. I want to thank all my HTI colleagues and fellow PhD students for making the journey enjoyable. I have witnessed people coming and going and the move of the group from IPO to Atlas, but our group always feels like home, because of all of you. Thanks to our dear secretaries and technical staff for all your patience and helps. Thanks also to all the HTI and P&T students who chose to work with me, and the time I spent supervising your projects was the most rewarding part of my job. Thanks to those who accompanied me in the conference trips abroad, Wijnand, Els, Peter, Elcin, and especially Hanne and Sofia, who were with me in the wonderful place called Kitchener to enjoy storms and waterfalls. Special thanks to Zhenzhen (famously known as Kong), my long-term roommate, for your emotional support and all the late evenings working together in the office. Also thanks to Huihui, for your kindness and humor that always shine.

I am equally grateful to my Data Science Flagship colleagues and friends – Linda, Marta, Mustafa, and Saskia from CPH, and also Alberto, Alok, Bart, Heleen, Humberto, Mark, and Niek. Despite all the problems that the flagship problem had, our group has always been strong and we supported each other in and outside research. The journey would be lonely and not as fun without you. It's the time that all of us finishes our PhDs and takes new adventures, but I wish the bond of the flagship remains.

To all my friends in Eindhoven and beyond, thank you for your support and the time you spent with me. There is nothing I should take for granted – playing football, playing tennis, going to concerts together, enjoying nice food, travelling, going to movies, playing games, or simply supporting each other at the most critical moments. There was certainly a difficult period when I just came back to Eindhoven, and without your company, I would not be at where I am now. Thanks also to some of my best friends in China. We cannot be with each other as often as we want, but whenever I returned home you welcomed me with love and delicious food. Special thanks to Lichao Pan and Kaipeng Hu for accepting to be my seconds, and to Qi Sun for proof-reading my thesis.

To my mom and dad, you loved me unconditionally and always supported me for whatever decisions I make. You are the ones who taught me to be an honest and good person. To my grandma, you are one of the most important person in my life. It's a pity that you cannot come for the defense, but it was great to have you for a summer in Europe. Also thanks to my uncle, aunt, and cousin, who supported me with affection and came to see me in Europe.

And to Yichen, I am so lucky to have you in this journey and beyond. Thank you for all the wonderful time we had together and the care and love you gave me. It is never easy to live together with a PhD, but I assure you that you did an amazing job.

Chao Zhang

October 28, 2019, in Eindhoven

Curriculum Vitae

Chao Zhang was born on November 4th 1987 in Wuhan, Hubei province, China. He studied industrial design at Huazhong University of Science and Technology and psychology at Central China Normal University, both in his hometown. After receiving his Bachelor degrees in 2010, he came to Eindhoven, the Netherlands to study in the Human-Technology Interaction Master program at Eindhoven University of Technology, with the sponsor of a talent scholarship from the university. In 2012, he obtained his Master of Science degree with Cum Laude. The Master thesis on presence in telecommunication, supervised by Daniël Lakens and Wijnand IJsselsteijn, was published as a journal article.

After briefly working in Tübingen, Germany as a research scientist, and in Shanghai, China as a user researcher, he returned in March 2015 to Eindhoven University of Technology to pursue his PhD. The project was embedded in the Data Science Flagship and part of the project was done at Philips Research. His research focused on modeling the cognitive processes in human learning and decision-making and applying the computational models to digital health interventions. His research outputs have been published in international conferences and journals. Since April 2019, he has worked as a post-doc researcher in the Human-Technology Interaction group at Eindhoven University of Technology.

