TIME IN MIND: UNDERSTANDING THE ROLE OF EPISODIC FUTURE THINKING IN INTERTEMPORAL CHOICE

TIME IN MIND: UNDERSTANDING THE ROLE OF EPISODIC FUTURE THINKING IN INTERTEMPORAL CHOICE

By ISAAC KINLEY, BSc

A Thesis Submitted to the School of Graduate Studies in Partial Fulfilment of the Requirements for the Degree of Doctor of Philosophy

McMaster University DOCTOR OF PHILOSOPHY (2023) Hamilton, Ontario (Psychology, Neuroscience & Behaviour

TITLE: Time in mind: understanding the role of episodic future thinking in intertemporal choice

AUTHOR: Isaac Kinley, BSc (McMaster University)

SUPERVISOR: Professor Suzanna Backer

NUMBER OF PAGES: xv, 123

Lay Abstract

Pursuing our goals for the future usually means sacrificing immediate gratification, yet we often make decisions that are not in our best interest over the long term. This is because we assign lower subjective value to future rewards the further they are from the present. Individuals differ in how much they devalue future rewards, and these differences are related to many real-world outcomes. Our tendency to devalue future rewards is reduced when we vividly imagine the future in a process called "episodic future thinking," and this thesis seeks to understand how this effect occurs. The most obvious explanation would seem to be that episodic future thinking "simulates" the experience of future rewards and allows us to recognize their value in the present. However, using results from several experimental studies, I argue that this may not be the best explanation after all, and I develop an alternative.

Abstract

Humans and other animals systematically discount the value of future rewards as a function of their delay, and individual differences in the steepness of this “delay discounting” are predictive of a range of important real-world outcomes. Episodic future thinking, the mental simulation of episodes in the personal future, is one means by which to curb delay discounting. This thesis seeks to contribute to our understanding of how this effect occurs. The account that predominates in the literature is that episodic future thinking simulates the experience of future rewards, enabling their undiscounted value to be appreciated in the present. This thesis takes this account as a starting point, formalizing it in a mathematical model and carrying out several experimental studies to test its predictions. We find that key predictions are not borne out and develop an alternative account in which simulated experience plays a less central role.

Acknowledgements

I would first like to thank my supervisor, Dr Sue Becker, for her inestimably valuable mentorship. Since my initial involvement with the lab 7 years ago, I have appreciated the atmosphere of intellectual rigour and interpersonal warmth that she cultivates. I always felt supported in pursuing whatever projects I found interesting, even up to the point of completely changing the topic of this thesis a year in. I am fortunate to have had the chance to draw on Sue's experience and expertise as I developed these projects, and she sets a brilliant example of the type of scientist and mentor I hope to someday be. Thank you for taking a chance on that nervous undergraduate so many years ago—it has made all the difference for me.

I am grateful also to the members past and present of my supervisory committee, Drs John Connolly, Michael Mazurek, Michael Amlung, and Margaret McKinnon, for their many years of insight and support. As I navigated this program of research, their well-honed sense of the lines of inquiry worth pursuing helped me avoid many dead ends and illuminated my path forward.

I received a great deal of help from many people who, as undergraduate students, collected data for the research here and for projects that did not make it into this thesis but nonetheless informed its development. In roughly chronological order, thanks to Connor Pires, Lauren Iafrate, Morgan Porteous, Joseph Oluwasola, Soroush Bagheri, Jana Lau, Katriel Read, and Tess Vosper.

Thanks also to my friends and colleagues in PNB for years of camaraderie. Some of my best graduate school memories are from those pre-lockdown Friday nights in PC 205. Thanks particularly to my labmates past and present, and thanks to the friends who have been around since this was all just a twinkle in a high schooler's eye.

Many thanks to Michal, who was the first to read much of this thesis. So much of my approach to life, both within academia and beyond, has been transformed by our time together. I am so lucky to have you by my side.

Finally, I owe so much beyond this thesis to the love of my parents and sister. I am so grateful to have had them to turn to during this long and sometimes arduous chapter of my life. As it comes to a close at last, I hope I have made you proud.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Humans are arguably unique in our ability to cast our minds into the future (Suddendorf and Corballis, 2007). This ability, called episodic future thinking (Atance and O'Neill, 2001), enables us to predict and imagine our futures, to set goals, and to make plans (Szpunar et al., 2014a). Episodic future thinking also enables flexible decision making between actions that lead to smaller immediate rewards versus larger delayed rewards (i.e., flexibility during *intertemporal choice*; Bulley et al., 2016), and is of particular interest for its apparent ability to encourage deferral of gratification (i.e., *intertemporal patience*; Rösch et al., 2022). Modern life increasingly demands intertemporal patience (Crawford, 2015), and intertemporal *im*patience is an important predictor of a range of unhealthy behaviours (Bickel et al., 2012, 2019) and psychiatric conditions (Amlung et al., 2019), particularly gambling and substance use disorders (Amlung et al., 2017).

Given the important relationship between episodic future thinking and intertemporal choice, recent research has begun to explore future thinking interventions for disorders and unhealthy behaviours in which intertemporal impatience is heavily implicated. For example, several studies have explored the effect of episodic future thinking on alcohol demand (Snider et al., 2016; Patel and Amlung, 2020) and consumption (Athamneh et al., 2022) in people with alcohol use disorders. Similar work has been done on cigarette smoking (Chiou and Wu, 2017; Stein et al., 2016) and food consumption in obesity (Daniel et al., 2013). Intertemporal impatience is usually taken to be the target of these interventions, so to maximize their effectiveness, it is important to understand exactly how and when episodic future thinking encourages deferred gratification. As we will see, despite a large body of empirical work, there has been relatively little progress toward a theoretical account of these effects, and much of this experimental work does little to constrain possible theoretical accounts. This thesis represents an attempt to move toward a more detailed theoretical account of the effect of episodic future thinking on intertemporal choice. To provide the appropriate background, I begin by separately discussing intertemporal choice and episodic future thinking before exploring their intersection.

1.1 Delay discounting

Humans and other species exhibit a near-universal tendency to undervalue future rewards as a function of their delay, a phenomenon called *delay discounting* (Odum, 2011a). There are a number of good reasons to discount future rewards: most obviously, as finite beings, it would not be rational to have an infinite temporal horizon of self-interest (Yaari, 1965). More practically, delay is confounded with risk—it is always possible that a promised future reward will become unavailable even if we are still around to receive it (Sozou, 1998). Moreover, in formal mathematical models of decision making, some discounting is required to make decision problems computationally tractable (Naik et al., 2019). Finally, the philosopher Derek Parfit argued that it may be rational to care less about future outcomes when one feels less connected to one's future self than one's present self (Parfit, 1984), as nearly all of us do (Bartels and Rips, 2010).

While some discounting is justified in theory, as discussed above, excessively steep discounting (in which the subjective value of future rewards declines rapidly with increasing delay) can be maladaptive and is considered a facet of impulsivity (Dalley and Robbins, 2017). Moreover, empirically, temporal discounting departs in interesting ways from normative theories of so-called "rational discounting", in which the subjective value of a future reward is expected to "decay" exponentially with increasing delay (Rachlin, 2006). Contrary to these expectations, humans (and other animals) are widely observed to exhibit *hyperbolic discounting* as described by the following equation (Mazur, 1987):

$$\frac{V_{\text{subj.}}}{V_{\text{obj.}}} = \frac{1}{1 + k \times \text{delay}}$$

where $V_{\text{subj.}}$ (the subjective value of a future reward), as a proportion of $V_{\text{obj.}}$ (the objective value of said reward), declines according to its delay and an individual-differences parameter k. This function does not always fit better than alternatives for individual decision makers (Franck et al., 2015), but is nonetheless widely used to quantify discounting.

An important consequence of hyperbolic discounting (indeed, of any non-exponential discounting function; Kurth-Nelson and Redish, 2012) is the phenomenon of intertemporal preference reversals, in which relative preferences between two possible future rewards do not remain constant over time (Strotz, 1955; Bulley and Schacter, 2020). For example, a student looking ahead to the next semester may clearly see the value in beginning their assignments early, but may end up choosing to spend time with friends when the semester actually begins. That is, the student's preference for an early start over a vibrant social life might reverse over time. Preference reversals, in turn, give rise to the phenomenon of "precommitment" (Kurth-Nelson and Redish, 2012), in which an agent anticipates a preference reversal and chooses to guarantee ahead of time that the smaller-sooner reward will be unavailable when the preference reversal occurs. The story often used to illustrate precommitment is that of Ulysses, who knows that his current preference not to be shipwrecked will reverse when he hears the Sirens' song. Thus he tells his crew to bind him to the mast of his ship and plug their own ears, ensuring he will have no choice but to stay the course at the critical moment (Elster, 1979). A less mythopoeic expression of precommitment might be ensuring that unhlealthy food will be unavailable at the moment of temptation by not purchasing it in the first place.

Although discounting rates are sometimes considered a trait variable (Odum, 2011b), there is a range of manipulations that can produce shallower discounting rates during an experiment (Rung and Madden, 2018; Scholten et al., 2019). For example, discounting is decreased by specifying the dates on which future rewards would occur rather than simply their distance from the present (e.g., "Jan 1" vs "in 2 months"; DeHart and Odum, 2015) and by making explicit that choosing an immediate reward means receiving no reward in the future (the "explicit zero" effect; Radu et al., 2011). Similarly, participants are less likely to choose immediate rewards when future rewards are concrete (e.g., a travel voucher) vs abstract (e.g., an equivalent amount of money; Kim et al., 2013). As we will see, an episodic future thinking cueing approach numbers among these patience-encouraging manipulations.

1.1.1 Neuroscience of delay discounting

It has long been argued that internally consistent decision making requires behaving as though one is employing a single scale to represent the value of different rewards (Samuelson et al., 1983). Activation of the ventromedial prefrontal cortex correlates with subjective valuations of a wide range of reward types (e.g., visual, olfactory, gustatory, and monetary; Peters and Büchel, 2010b), prompting the suggestion that ventromedial prefrontal activation is the "common currency" into which all rewards are converted for comparison (Levy and Glimcher, 2012). The prefrontal cortex is anatomically well-positioned for such a role, as a hub of integration for signals from across the rest of the brain; indeed, the "appraisal-by-content" hypothesis (Dixon et al., 2017) proposes that different medial prefrontal subregions assign value to different types of information depending on the type of input they preferentially receive (e.g., exteroceptive information, episodic memories, etc.). This appears to describe the situation in intertemporal choice well, where immediate and delayed reward values are represented by distinct, anticorrelated subpopulations of neurons in the medial prefrontal cortex, with delayed reward representations lying more anteriorly along an established posterior-to-anterior abstraction gradient (Wang et al., 2014; Dixon et al., 2017).

An account of intertemporal choice in which both the immediate and delayed reward values are represented in the medial prefrontal cortex contrasts with dual-systems models of delay discounting. For example, McClure et al. (2004) argue for the "beta-delta" model in which only immediate rewards are represented by a limbic–medial prefrontal "beta" system, whereas delayed reward values are represented in the dorsolateral prefrontal cortex. This is based in part on the observation that dorsolateral prefrontal activation often accompanies the choice of larger delayed rewards (McClure et al., 2004). Peters and Büchel (2011) instead argue that lateral prefrontal cortical activation does not index reward valuation *per se*, but rather reflects the engagement of cognitive control or "willpower", which acts by modulating medial prefrontal representations of reward value (Hare et al., 2009).

In either case, the medial prefrontal cortex is clearly implicated in intertemporal choices. As we will see, this structure is a site of neural overlap between reward valuation and episodic future thinking, suggesting an avenue by which our choices might be influenced by our imagery of the future.

1.2 Episodic future thinking

Episodic future thinking, the process of imagining in detail events that may happen in the future, is a potentially powerful modulator of delay discounting. Though they may seem to be opposites, there is a deep connection between episodic future thinking and episodic memory. The *constructive episodic simulation hypothesis* (Schacter and Addis, 2007a) argues that episodic future thinking uses the same constructive processes and content as episodic memory to flexibly recombine perceptual details into novel imagined future scenarios. This is proposed to explain the well-known fallibility of memory: the same flexibility that enables episodic future thinking also produces memory errors such as the erroneous incorporation of details into event memories (Schacter and Addis, 2007b). Thus, the distortions that inevitably accompany constructive memory can be seen as the price of prospection (Schacter, 2012). Addis (2020) goes even further: rather than seeing episodic future thinking as a secondary benefit of a system whose fundamental function is memory, we can see both episodic future thinking and episodic memory as products of a single "simulation system", the evolutionary "purpose" of which would fundamentally be its orientation toward the future rather than the past (Klein, 2013).

Numerous lines of behavioural evidence point to an overlap in component processes between episodic memory and future thinking. First, both abilities emerge around the same time in development, between three and five years of age (Schacter et al., 2012). Moreover, various conditions that are accompanied by impairments of episodic memory are also accompanied by episodic future thinking deficits, including depression (Addis et al., 2016) and post-traumatic stress disorder (Brown et al., 2014). This occurs most dramatically in the case of medial temporal lobe amnesia, in which patients are more or less altogether unable to either remember the past or imagine the future (Hassabis et al., 2007; Klein et al., 2002; Race et al., 2011).

Healthy aging is also associated with concomitant declines in episodic memory and future thinking abilities (Schacter et al., 2013). These are often measured using the *autobiographical interview* (Levine et al., 2002), a method for scoring the episodic detail of event narratives. However, while older adults provide less episodic detail than younger adults for both memories and imagined future events, the same pattern of results holds for a picture description control task (Gaesser et al., 2011), suggesting that ostensible age-related "declines" in episodic memory and future thinking could be driven in part by changes in communicative style rather than simply by cognitive deficits. As we will see, communicative style is one of several factors outside of constructive processes that could produce shared variance between episodic memory and future thinking performance. Nonetheless, a great deal of existing research has investigated the shared reliance of these two abilities on constructive processes, as described by the constructive episodic simulation hypothesis.

Recently, an experimental procedure called the *episodic specificity induction* has been shown to enhance episodic richness for descriptions of remembered episodes and imagined future episodes without any similar effect for control description tasks not thought to depend on constructive processes (Madore et al., 2014; Madore and Schacter, 2016). Thus, this induction is argued to selectively enhance constructive processes (Schacter et al., 2017) and can provide clues as to the precise nature of these processes and the cognitive abilities beyond memory and future thinking in which they are implicated. For example, the specificity induction enhances means–end problem solving (Madore and Schacter, 2014; Jing et al., 2016) and creativity (Madore et al., 2015). Moreover, the specificity induction increases the number of details integrated into an imagined scene without necessarily increasing the spatial *coherence* of that scene (Madore et al., 2019). In contrast, the "scene construction" account of episodic memory (Hassabis and Maguire, 2007; Maguire and Mullally, 2013) might suggest that the specificity induction would target a process of constructing a spatial scene.

Of course, scene construction could still play a role in both episodic memory and future thinking without being targeted by the specificity induction. Other basic abilities shared between episodic memory and future thinking, but not thought to be targeted by the specificity induction, include visual mental imagery (D'Argembeau and Van der Linden, 2006) and working memory (Zavagnin et al., 2016). Note, however, that working memory may be more heavily taxed by episodic future thinking than memory (Hill and Emery, 2013). Indeed, episodic future thinking is generally reported by participants to be more difficult (McDonough and Gallo, 2010) and less vivid (D'Argembeau and Van der Linden, 2004) than episodic memory.

Semantic processes have also received increasing attention as an area of overlap between episodic memory and future thinking. Autobiographical memory has long been argued to represent a dynamic interplay between semantic autobiographical knowledge and episodic detail (Conway and Pleydell-Pearce, 2000). A given memory can be thought of as a pattern of activation over a hierarchically organized structure representing lifetime periods, event types, and episodic details, and the same process is argued to underlie episodic future thinking (Conway et al., 2019). Indeed, construction of imagined future episodes seems to involve first accessing personal semantic knowledge and only later incorporating perceptual details (D'Argembeau and Mathy, 2011). Similarly, Irish and Piguet (2013) argue that semantic memory provides a "scaffold" upon which episodic detail can be elaborated, based in part on the observation that semantic dementia impairs detail generation in episodic future thinking. Addis (2020) proposes that this semantic scaffolding process consists of the activation of relevant *schemas* as templates for different categories of events (Addis, 2018), along with *relational reasoning* (Alexander, 2016) to guide the integration of plausible details into memories and imagined future episodes.

1.2.1 Neuroscience of episodic future thinking

The broad overlap in component processes between episodic memory and future thinking is mirrored in a similarly striking overlap in their neural substrates. Both processes evoke activity in the brain's "default mode" network, initially observed as a set of structures that become less active during a range of attention-demanding, externally oriented tasks (Raichle et al., 2001; Raichle, 2015). Core nodes of this network include the medial prefrontal cortex, medial temporal lobe, and medial parietal cortex (Smallwood et al., 2021).

The default mode network is functionally heterogeneous (i.e., its component parts are involved in a range of task domains; Anderson et al., 2013), and a great deal of work has attempted to identify the specific contributions of its substructures to episodic future thinking. Andrews-Hanna et al. (2010) suggest that the network can be divided into a "midline core", active during self-relevant affective decisions, and a medial temporal system involved in scene imagery. Other authors have emphasized the valuational role of the medial prefrontal cortex (Benoit et al., 2014), especially for self-relevant valuational targets (D'Argembeau, 2013) and personal goals (Stawarczyk and D'Argembeau, 2015). In contrast, some authors point to its role in social cognition, specifically the representation of familiar others in imagined scenarios (Szpunar et al., 2014b). Of course, familiar others are usually socially close, and medial prefrontal activity is sensitive to the overlapping constructs of social closeness (Krienen et al., 2010), personal relevance (Abraham, 2013), and value (Peters and Büchel, 2010b), suggesting that existing research has largely converged on an essentially valuational role for the medial prefrontal cortex in episodic future thinking.

Interpretations of the role of the medial temporal subsystem have tended to be less mutually compatible. Some authors argue for a purely spatial interpretation, in which the hippocampus' essential role is the construction of mental scenes (*scene construction*, as described earlier; Hassabis and Maguire, 2007; Maguire and Mullally, 2013) where imagined events take place (Palombo et al., 2016). This accords with a range of findings implicating medial temporal structures in representing imagined self-location (Bellmund et al., 2016; Bergouignan and Paz-Alonso, 2022) and imagined navigation (Horner et al., 2016). Similarly, the hippocampus can be seen as representing spatial relationships that "scaffold" memory and imagination (Robin et al., 2016; Robin, 2018), and perhaps even allow non-spatial (e.g., social or temporal) relationships to be represented according to spatial analogies (as reviewed in Becker, 2023). In contrast, others argue that the essential role of the hippocampus is *relational processing*, i.e., binding together both the spatial and non-spatial (e.g., emotions, dialogue) elements of an imagined episode, particularly when the conjunction of these elements is novel (Roberts et al., 2018). This is thought to explain why hippocampal engagement increases with the level of imagined detail (as there is more to bind; Addis et al., 2011) and decreases as an episode is repeatedly imagined (as the conjunction of its elements becomes less novel; van Mulukom et al., 2013).

Of course, there is no need to restrict the entire hippocampus to a single function: indeed, experimentally, distinct hippocampal sub-regions appear to be recruited for relational processing and scene construction (Dalton et al., 2018). It seems likely, then, that episodic future thinking engages the hippocampus in the service of both of these processes simultaneously. An intriguing possibility raised by Eichenbaum (2017) is that the hippocampus is well-suited to represent relationships in general, whether these are relationships of novel co-occurrence, transitive order relationships (Dusek and Eichenbaum, 1997), adjacencies in an abstract graph (Schapiro et al., 2016), or distances in metric spaces defined by social (Schafer and Schiller, 2018), object-featural (Constantinescu et al., 2016), or physical spatial (O'keefe and Nadel, 1979) dimensions. Perhaps, then, the precise role of the hippocampus in episodic future thinking depends on the specific content of the imagined future event and the nature of the relationships being represented.

Regardless of the precise characterization of default mode activity, episodic simulation clearly involves the co-activation of regions involved in both valuation-like and scene-imagery-like processes. Importantly, episodic future thinking involves a special degree of influence of the latter regions on the former (Campbell et al., 2018), suggesting a route by which imagery of the future might influence intertemporal choice.

1.3 Episodic future thinking and decision making

It is by now well established that providing cues to elicit episodic future thinking during intertemporal choices reduces delay discounting (Rösch et al., 2022). fMRI work has revealed that the extent of this effect depends on the degree of evoked coactivation of medial temporal and medial prefrontal structures (Peters and Büchel, 2010a; Benoit et al., 2011; Sasse et al., 2015), suggesting that medial prefrontal representations of reward can be driven by hippocampus-dependent episodic simulations in much the way they can be driven by dorsolateral-prefrontal-dependent willpower (Peters and Büchel, 2011). However, beyond the core notion that episodic future thinking influences intertemporal choice by influencing reward valuation, these results do little to constrain theories about the precise nature of this process.

Is there a need for such theories? The effect of episodic future thinking on delay discounting may seem obvious to the point of tautology—no surprise that thinking ahead should make us more future-oriented. However, different studies ascribe widely ranging magnitudes to this effect (Hollis-Hansen et al., 2019), meaning there is a question of *when*, and thus *how*, the effect occurs. There have been some attempts to explicitly theorize about the role of episodic future thinking in intertemporal choice: Boyer (2008) argues that episodic future thinking bypasses discounting by providing an imagined preview of future reward, a view echoed by others (Benoit et al., 2011; Bulley et al., 2016). Hoerl and McCormack (2016) point out that, according to this account, the *simulation* of future events is critical to the effect of episodic future thinking on decision making. From this, we can draw the prediction that episodic future thinking will influence decision making to the extent that it represents a *successful* simulation of the future.

The degree of "simulation success" in episodic future thinking could be measured through a number of variables. Much as in the case of episodic memory, episodic future thinking is considered distinct from its semantic counterpart in that it involves autonoetic consciousness (Atance and O'Neill, 2001), which is the sense that one is "re-experiencing" a remembered episode (Tulving, 1985) or "pre-experiencing" an imagined future episode (D'Argembeau and Van der Linden, 2004). Self-reported degree of autonoetic consciousness is strongly correlated with factors such as overall vividness, level of detail, and clarity (D'Argembeau and Van der Linden, 2012), any of which could be taken as an index of the degree to which an attempted mental simulation of a future episode has succeeded.

If "simulation success" is indeed necessary for the effect of episodic future thinking, we can further ask whether whether it influences decision making through some mediating factor. For example, (successfully) mentally simulating the future may impact one of the many known factors that influence delay discounting (Rung and Madden, 2018; Scholten et al., 2019), reducing episodic future thinking to a special case of construal level or perhaps framing effects. Alternatively, the role of episodic future thinking may be more unique: mentally simulating the future may make it seem closer in time (Hoerl and McCormack, 2016) or less uncertain (Bulley et al., 2016; Kinley et al., 2022), or may provoke a heightened sense of connection to one's future self (Hoerl and McCormack, 2016). The remainder of the thesis follows from this line of reasoning in attempting to clarify the role of episodic future thinking in intertemporal choice.

1.4 Summary of remaining chapters

In the second chapter of this thesis, we take a "first-principles" approach, using formal reinforcement learning models of decision making to capture the essential features of intertemporal choice and to explore how the role of episodic future thinking can be understood in such a setting. The chapter is framed in the context of a computational model of addiction, a key element of which is steep delay discounting. It relies on the distinction between a fast, habitual, "model-free" decision making system and a competing slow, deliberative, "model-based" system, where episodic future thinking is identified with the latter. We propose that, under appropriate assumptions, episodic future thinking might decrease the apparent uncertainty of future rewards and thus their probability of being chosen over immediate rewards. The broader topic of the chapter is a critique of the "habit theory" of addiction, which sees addiction as resulting from a dominance of habitual over goal-directed decision making. We show that, in order for a model based on this theory to "work" (i.e., to give rise to addiction-like behaviours), we are required to assume that suboptimal addictive rewards are overvalued by the habit system and not by the planning system. This assumption is called into question by recent experimental evidence.

In chapter three, we attempted to identify empirical correlates of the effect of episodic future thinking on delay discounting, with the goal of determining whether "simulation success" indeed amplifies the effect of episodic future thinking on delay discounting. Following the now-standard paradigm introduced by Peters and Büchel (2010a), participants completed a delay discounting monetary choice task, with some trials including cues to previously-described future events. A variety of analytic approaches were used to determine whether any self-reported aspect of these imagined events (e.g., vividness, emotional valence/intensity) correlated with the effect of episodic future thinking cues on choices, but no correlation appeared. These results are described in detail in the chapter, along with several novel findings concerning visual perspective and trait dissociativity. The experiment used a novel pair of questionnaire items to query visual perspective (first- vs third-person), finding that participants often report switching between these perspectives and, moreover, the degree to which they report this switching is correlated with their trait dissociativity, as measured by the Dissociative Experiences Scale (Bernstein and Putnam, 1986). Thus, the chapter is framed in terms of these visual perspective results. The general discussion of this thesis will describe the methodological and theoretical insights that can be derived from the pattern of delay discounting results in this chapter.

The fourth chapter represents a second attempt to identify correlates of the episodic future thinking cue effect on delay discounting. It seemed possible that the lack of identified correlations in the third chapter was a result of insufficient variance in the predictor variables. That is, perhaps there was no correlation between simulation success and cue effect because all participants more or less successfully simulated future events. Thus, in the fourth chapter we took an experimental approach to try to increase the variance in simulation success: participants were assigned at random to an episodic future thinking group, in which they imagined personal future events, or a semantic future thinking group, in which they imagined non-self-relevant future events. This novel experimental manipulation was successful in that participants in the episodic group gave higher self-reported ratings of vividness and autonoetic consciousness. Moreover, only participants in the episodic group exhibited a cue effect. However, the difference in the cue effect between groups was not mediated by any difference in self-report ratings. This pattern of results held across two experiments with a collective sample size of 451. Thus, we again did not find evidence that simulation success is essential to the effect of episodic future thinking on discounting, despite a large sample size and strong experimental manipulation. We speculate that the self-relevance of imagined future events, the factor operationally differentiating the episodic and semantic future thinking groups, may have also been the factor giving rise to the larger effect on decision making in the episodic future thinking group. In the general discussion of this thesis, I will pick up on this thread and attempt to specify how episodic future thinking might influence delay discounting beyond the "simulation success" view outlined above.

Bibliography

Abraham, A. (2013). The world according to me: personal relevance and the medial prefrontal cortex. *Frontiers in human neuroscience*, 7:341.

Addis, D. R. (2018). Are episodic memories special? on the sameness of remembered and imagined event simulation. *Journal of the Royal Society of New Zealand*, 48(2-3):64–88.

Addis, D. R. (2020). Mental time travel? a neurocognitive model of event simulation. *Review of Philosophy and Psychology*, 11:233–259.

Addis, D. R., Cheng, T., P. Roberts, R., and Schacter, D. L. (2011). Hippocampal contributions to the episodic simulation of specific and general future events. *Hippocampus*, 21(10):1045–1052.

Addis, D. R., Hach, S., and Tippett, L. J. (2016). Do strategic processes contribute to the specificity of future simulation in depression? *British Journal of Clinical Psychology*, 55(2):167–186.

Alexander, P. A. (2016). Relational thinking and relational reasoning: harnessing the power of patterning. *NPJ science of learning*, 1(1):1–7.

Amlung, M., Marsden, E., Holshausen, K., Morris, V., Patel, H., Vedelago, L., Naish, K. R., Reed, D. D., and McCabe, R. E. (2019). Delay discounting as a transdiagnostic process in psychiatric disorders: A meta-analysis. *JAMA psychiatry*, 76(11):1176–1186.

Amlung, M., Vedelago, L., Acker, J., Balodis, I., and MacKillop, J. (2017). Steep delay discounting and addictive behavior: A meta-analysis of continuous associations. *Addiction*, 112(1):51–62.

Anderson, M. L., Kinnison, J., and Pessoa, L. (2013). Describing functional diversity of brain regions and brain networks. *Neuroimage*, 73:50–58.

Andrews-Hanna, J. R., Reidler, J. S., Sepulcre, J., Poulin, R., and Buckner, R. L. (2010). Functional-anatomic fractionation of the brain's default network. *Neuron*, 65(4):550–562.

Atance, C. M. and O'Neill, D. K. (2001). Episodic future thinking. *Trends in cognitive sciences*, 5(12):533–539.

Athamneh, L. N., Brown, J., Stein, J. S., Gatchalian, K. M., LaConte, S. M., and Bickel, W. K. (2022). Future thinking to decrease real-world drinking in alcohol use disorder: Repairing reinforcer pathology in a randomized proof-of-concept trial. *Experimental and Clinical Psychopharmacology*, 30(3):326.

Bartels, D. M. and Rips, L. J. (2010). Psychological connectedness and intertemporal choice. *Journal of Experimental Psychology: General*, 139(1):49.

Becker, S. (2023). Memory for space. In Kahana, M. J. and Wagner, A. D., editors, *The Oxford Handbook of Human Memory*, chapter 3.3. Oxford University Press, Oxford.

Bellmund, J. L., Deuker, L., Navarro Schröder, T., and Doeller, C. F. (2016). Grid-cell representations in mental simulation. *Elife*, 5:e17089.

Benoit, R. G., Gilbert, S. J., and Burgess, P. W. (2011). A neural mechanism mediating the impact of episodic prospection on farsighted decisions. *Journal of Neuroscience*, 31(18):6771–6779.

Benoit, R. G., Szpunar, K. K., and Schacter, D. L. (2014). Ventromedial prefrontal cortex supports affective future simulation by integrating distributed knowledge. *Proceedings of the National Academy of Sciences*, 111(46):16550–16555.

Bergouignan, L. and Paz-Alonso, P. M. (2022). Simulating the situated-self drives hippocampo-cortical engagement during inner narration of events. *Cerebral Cortex*, 32(24):5716–5731.

Bernstein, E. M. and Putnam, F. W. (1986). Development, reliability, and validity of a dissociation scale. *Journal of Nervous and Mental Disease*.

Bickel, W. K., Athamneh, L. N., Basso, J. C., Mellis, A. M., DeHart, W. B., Craft, W. H., and Pope, D. (2019). Excessive discounting of delayed reinforcers as a trans-disease process: Update on the state of the science. *Current Opinion in Psychology*, 30:59–64.

Bickel, W. K., Jarmolowicz, D. P., Mueller, E. T., Koffarnus, M. N., and Gatchalian, K. M. (2012). Excessive discounting of delayed reinforcers as a trans-disease process contributing to addiction and other disease-related vulnerabilities: emerging evidence. *Pharmacology & therapeutics*, 134(3):287–297.

Boyer, P. (2008). Evolutionary economics of mental time travel? *Trends in cognitive sciences*, 12(6):219–224.

Brown, A. D., Addis, D. R., Romano, T. A., Marmar, C. R., Bryant, R. A., Hirst, W., and Schacter, D. L. (2014). Episodic and semantic components of autobiographical memories and imagined future events in post-traumatic stress disorder. *Memory*, 22(6):595–604.

Bulley, A., Henry, J., and Suddendorf, T. (2016). Prospection and the present moment: The role of episodic foresight in intertemporal choices between immediate and delayed rewards. *Review of General Psychology*, 20(1):29–47.

Bulley, A. and Schacter, D. L. (2020). Deliberating trade-offs with the future. *Nature human behaviour*, 4(3):238–247.

Campbell, K. L., Madore, K. P., Benoit, R. G., Thakral, P. P., and Schacter, D. L. (2018). Increased hippocampus to ventromedial prefrontal connectivity during the construction of episodic future events. *Hippocampus*, 28(2):76–80.

Chiou, W.-B. and Wu, W.-H. (2017). Episodic future thinking involving the non-smoking self can induce lower discounting and cigarette consumption. *Journal of Studies on Alcohol and Drugs*, 78(1):106–112.

Constantinescu, A. O., O'Reilly, J. X., and Behrens, T. E. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, 352(6292):1464–1468.

Conway, M. A., Justice, L. V., and D'Argembeau, A. (2019). The self-memory system revisited. *The organization and structure of autobiographical memory*, pages 28–51.

Conway, M. A. and Pleydell-Pearce, C. W. (2000). The construction of autobiographical memories in the self-memory system. *Psychological review*, 107(2):261.

Crawford, M. B. (2015). *The world beyond your head: On becoming an individual in an age of distraction.* Farrar, Straus and Giroux.

Dalley, J. W. and Robbins, T. W. (2017). Fractionating impulsivity: neuropsychiatric implications. *Nature Reviews Neuroscience*, 18(3):158–171.

Dalton, M. A., Zeidman, P., McCormick, C., and Maguire, E. A. (2018). Differentiable processing of objects, associations, and scenes within the hippocampus. *Journal of Neuroscience*, 38(38):8146–8159.

Daniel, T. O., Stanton, C. M., and Epstein, L. H. (2013). The future is now: reducing impulsivity and energy intake using episodic future thinking. *Psychological science*, 24(11):2339–2342.

D'Argembeau, A. and Mathy, A. (2011). Tracking the construction of episodic future thoughts. *Journal of experimental psychology: General*, 140(2):258.

DeHart, W. B. and Odum, A. L. (2015). The effects of the framing of time on delay discounting. *Journal of the experimental analysis of behavior*, 103(1):10–21.

Dixon, M. L., Thiruchselvam, R., Todd, R., and Christoff, K. (2017). Emotion and the prefrontal cortex: An integrative review. *Psychological bulletin*, 143(10):1033.

Dusek, J. A. and Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences*, 94(13):7109–7114.

D'Argembeau, A. (2013). On the role of the ventromedial prefrontal cortex in self-processing: the valuation hypothesis. *Frontiers in human neuroscience*, 7:372.

D'Argembeau, A. and Van der Linden, M. (2004). Phenomenal characteristics associated with projecting oneself back into the past and forward into the future: Influence of valence and temporal distance. *Consciousness and cognition*, 13(4):844–858.

D'Argembeau, A. and Van der Linden, M. (2006). Individual differences in the phenomenology of mental time travel: The effect of vivid visual imagery and emotion regulation strategies. *Consciousness and cognition*, 15(2):342–350.

D'Argembeau, A. and Van der Linden, M. (2012). Predicting the phenomenology of episodic future thoughts. *Consciousness and cognition*, 21(3):1198–1206.

Eichenbaum, H. (2017). The role of the hippocampus in navigation is memory. *Journal of neurophysiology*, 117(4):1785–1796.

Elster, J. (1979). *Ulysses and the sirens: Studies in rationality and irrationality*. Cambridge University Press.

Franck, C. T., Koffarnus, M. N., House, L. L., and Bickel, W. K. (2015). Accurate characterization of delay discounting: A multiple model approach using approximate bayesian model selection and a unified discounting measure. *Journal of the experimental analysis of behavior*, 103(1):218–233.

Gaesser, B., Sacchetti, D. C., Addis, D. R., and Schacter, D. L. (2011). Characterizing age-related changes in remembering the past and imagining the future. *Psychology and aging*, 26(1):80.

Hare, T. A., Camerer, C. F., and Rangel, A. (2009). Self-control in decision-making involves modulation of the vmpfc valuation system. *Science*, 324(5927):646–648.

Hassabis, D., Kumaran, D., Vann, S. D., and Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *Proceedings of the National Academy of Sciences*, 104(5):1726–1731.

Hassabis, D. and Maguire, E. A. (2007). Deconstructing episodic memory with construction. *Trends in cognitive sciences*, 11(7):299–306.

Hill, P. F. and Emery, L. J. (2013). Episodic future thought: Contributions from working memory. *Consciousness and cognition*, 22(3):677–683.

Hoerl, C. and McCormack, T. (2016). Making decisions about the future. *Seeing the future: Theoretical perspectives on future-oriented mental time travel*, page 241.

Hollis-Hansen, K., O'Donnell, S. E., Seidman, J. S., Brande, S. J., and Epstein, L. H. (2019). Improvements in episodic future thinking methodology: Establishing a standardized episodic thinking control. *PLoS One*, 14(3):e0214397.

Horner, A. J., Bisby, J. A., Zotow, E., Bush, D., and Burgess, N. (2016). Grid-like processing of imagined navigation. *Current Biology*, 26(6):842–847.

Irish, M. and Piguet, O. (2013). The pivotal role of semantic memory in remembering the past and imagining the future. *Frontiers in behavioral neuroscience*, 7:27.

Jing, H. G., Madore, K. P., and Schacter, D. L. (2016). Worrying about the future: An episodic specificity induction impacts problem solving, reappraisal, and well-being. *Journal of Experimental Psychology: General*, 145(4):402.

Kim, H., Schnall, S., and White, M. P. (2013). Similar psychological distance reduces temporal discounting. *Personality and Social Psychology Bulletin*, 39(8):1005–1016.

Kinley, I., Amlung, M., and Becker, S. (2022). Pathologies of precision: A bayesian account of goals, habits, and episodic foresight in addiction. *Brain and Cognition*, 158:105843.

Klein, S. B. (2013). The temporal orientation of memory: It's time for a change of direction. *Journal of Applied Research in Memory and Cognition*, 2(4):222–234.

Klein, S. B., Loftus, J., and Kihlstrom, J. F. (2002). Memory and temporal experience: The effects of episodic memory loss on an amnesic patient's ability to remember the past and imagine the future. *Social Cognition*, 20(5):353–379.

Krienen, F. M., Tu, P.-C., and Buckner, R. L. (2010). Clan mentality: evidence that the medial prefrontal cortex responds to close others. *Journal of Neuroscience*, 30(41):13906–13915.

Kurth-Nelson, Z. and Redish, A. D. (2012). Don't let me do that!–models of precommitment. *Frontiers in Neuroscience*, 6:138.

Levine, B., Svoboda, E., Hay, J. F., Winocur, G., and Moscovitch, M. (2002). Aging and autobiographical memory: dissociating episodic from semantic retrieval. *Psychology and aging*, 17(4):677.

Levy, D. J. and Glimcher, P. W. (2012). The root of all value: a neural common currency for choice. *Current opinion in neurobiology*, 22(6):1027–1038.

Madore, K. P., Addis, D. R., and Schacter, D. L. (2015). Creativity and memory: Effects of an episodic-specificity induction on divergent thinking. *Psychological science*, 26(9):1461–1468.

Madore, K. P., Gaesser, B., and Schacter, D. L. (2014). Constructive episodic simulation: dissociable effects of a specificity induction on remembering, imagining, and describing in young and older adults. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(3):609.

Madore, K. P., Jing, H. G., and Schacter, D. L. (2019). Episodic specificity induction and scene construction: Evidence for an event construction account. *Consciousness and Cognition*, 68:1–11.

Madore, K. P. and Schacter, D. L. (2014). An episodic specificity induction enhances means-end problem solving in young and older adults. *Psychology and aging*, 29(4):913.

Madore, K. P. and Schacter, D. L. (2016). Remembering the past and imagining the future: Selective effects of an episodic specificity induction on detail generation. *Quarterly Journal of Experimental Psychology*, 69(2):285–298.

Maguire, E. A. and Mullally, S. L. (2013). The hippocampus: a manifesto for change. *Journal of Experimental Psychology: General*, 142(4):1180.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. *Quantitative analyses of behavior*, 5:55–73.

McClure, S. M., Laibson, D. I., Loewenstein, G., and Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science*, 306(5695):503–507.

McDonough, I. M. and Gallo, D. A. (2010). Separating past and future autobiographical events in memory: Evidence for a reality monitoring asymmetry. *Memory & Cognition*, 38:3–12.

Naik, A., Shariff, R., Yasui, N., Yao, H., and Sutton, R. S. (2019). Discounted reinforcement learning is not an optimization problem. *arXiv preprint arXiv:1910.02140*.

Odum, A. L. (2011a). Delay discounting: I'm ak, you're ak. *Journal of the experimental analysis of behavior*, 96(3):427–439.

Odum, A. L. (2011b). Delay discounting: trait variable? *Behavioural processes*, 87(1):1–9.

O'keefe, J. and Nadel, L. (1979). Précis of o'keefe & nadel's the hippocampus as a cognitive map. *Behavioral and Brain Sciences*, 2(4):487–494.

Palombo, D., Hayes, S., Peterson, K., Keane, M., and Verfaellie, M. (2016). Medial temporal lobe contributions to episodic future thinking: Scene construction or future projection? *Cerebral Cortex*, 1:12.

Parfit, D. (1984). *Reasons and persons*. OUP Oxford.

Patel, H. and Amlung, M. (2020). Acute and extended exposure to episodic future thinking in a treatment seeking addiction sample: A pilot study. *Journal of Substance Abuse Treatment*, 116:108046.

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

Peters, J. and Büchel, C. (2010b). Neural representations of subjective reward value. *Behavioural brain research*, 213(2):135–141.

Peters, J. and Büchel, C. (2011). The neural mechanisms of inter-temporal decision-making: understanding variability. *Trends in cognitive sciences*, 15(5):227–239.

Race, E., Keane, M. M., and Verfaellie, M. (2011). Medial temporal lobe damage causes deficits in episodic memory and episodic future thinking not attributable to deficits in narrative construction. *Journal of Neuroscience*, 31(28):10262–10269.

Rachlin, H. (2006). Notes on discounting. *Journal of the experimental analysis of behavior*, 85(3):425–435.

Radu, P. T., Yi, R., Bickel, W. K., Gross, J. J., and McClure, S. M. (2011). A mechanism for reducing delay discounting by altering temporal attention. *Journal of the experimental analysis of behavior*, 96(3):363–385.

Raichle, M. E. (2015). The brain's default mode network. *Annual review of neuroscience*, 38:433–447.

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., and Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the national academy of sciences*, 98(2):676–682.

Roberts, R. P., Schacter, D. L., and Addis, D. R. (2018). Scene construction and relational processing: separable constructs? *Cerebral Cortex*, 28(5):1729–1732.

Robin, J. (2018). Spatial scaffold effects in event memory and imagination. *Wiley Interdisciplinary Reviews: Cognitive Science*, 9(4):e1462.

Robin, J., Wynn, J., and Moscovitch, M. (2016). The spatial scaffold: The effects of spatial context on memory for events. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(2):308.

Rösch, S. A., Stramaccia, D. F., and Benoit, R. G. (2022). Promoting farsighted decisions via episodic future thinking: A meta-analysis. *Journal of Experimental Psychology: General*, 151(7):1606.

Rung, J. M. and Madden, G. J. (2018). Experimental reductions of delay discounting and impulsive choice: A systematic review and meta-analysis. *Journal of experimental psychology: general*, 147(9):1349.

Samuelson, P. A., Samuelson, P. A., Samuelson, P. A., and Samuelson, P. A. (1983). *Foundations of economic analysis*. Harvard university press Cambridge, MA.

Sasse, L. K., Peters, J., Büchel, C., and Brassen, S. (2015). Effects of prospective thinking on intertemporal choice: The role of familiarity. *Human brain mapping*, 36(10):4210–4221.

Schacter, D. L. (2012). Constructive memory: past and future. *Dialogues in Clinical Neuroscience*, 14(1):7.

Schacter, D. L. and Addis, D. R. (2007a). The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481):773–786.

Schacter, D. L. and Addis, D. R. (2007b). The ghosts of past and future. *Nature*, 445(7123):27–27.

Schacter, D. L., Addis, D. R., Hassabis, D., Martin, V. C., Spreng, R. N., and Szpunar, K. K. (2012). The future of memory: remembering, imagining, and the brain. *Neuron*, 76(4):677–694.

Schacter, D. L., Benoit, R. G., and Szpunar, K. K. (2017). Episodic future thinking: Mechanisms and functions. *Current opinion in behavioral sciences*, 17:41–50.

Schacter, D. L., Gaesser, B., and Addis, D. R. (2013). Remembering the past and imagining the future in the elderly. *Gerontology*, 59(2):143–151.

Schafer, M. and Schiller, D. (2018). Navigating social space. *Neuron*, 100(2):476–489.

Schapiro, A. C., Turk-Browne, N. B., Norman, K. A., and Botvinick, M. M. (2016). Statistical learning of temporal community structure in the hippocampus. *Hippocampus*, 26(1):3–8.

Scholten, H., Scheres, A., De Water, E., Graf, U., Granic, I., and Luijten, M. (2019). Behavioral trainings and manipulations to reduce delay discounting: A systematic review. *Psychonomic Bulletin & Review*, 26:1803–1849.

Smallwood, J., Bernhardt, B. C., Leech, R., Bzdok, D., Jefferies, E., and Margulies, D. S. (2021). The default mode network in cognition: a topographical perspective. *Nature reviews neuroscience*, 22(8):503–513.

Snider, S. E., LaConte, S. M., and Bickel, W. K. (2016). Episodic future thinking: Expansion of the temporal window in individuals with alcohol dependence. *Alcoholism: clinical and experimental research*, 40(7):1558–1566.

Sozou, P. D. (1998). On hyperbolic discounting and uncertain hazard rates. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 265(1409):2015–2020.

Stawarczyk, D. and D'Argembeau, A. (2015). Neural correlates of personal goal processing during episodic future thinking and mind-wandering: An ale meta-analysis. *Human brain mapping*, 36(8):2928–2947.

Stein, J. S., Wilson, A. G., Koffarnus, M. N., Daniel, T. O., Epstein, L. H., and Bickel, W. K. (2016). Unstuck in time: episodic future thinking reduces delay discounting and cigarette smoking. *Psychopharmacology*, 233:3771–3778.

Strotz, R. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23(3):165–180.

Suddendorf, T. and Corballis, M. C. (2007). The evolution of foresight: What is mental time travel, and is it unique to humans? *Behavioral and brain sciences*, 30(3):299–313.

Szpunar, K. K., Spreng, R. N., and Schacter, D. L. (2014a). A taxonomy of prospection: Introducing an organizational framework for future-oriented cognition. *Proceedings of the National Academy of Sciences*, 111(52):18414–18421.

Szpunar, K. K., St. Jacques, P. L., Robbins, C. A., Wig, G. S., and Schacter, D. L. (2014b). Repetition-related reductions in neural activity reveal component processes of mental simulation. *Social Cognitive and Affective Neuroscience*, 9(5):712–722.

Tulving, E. (1985). Memory and consciousness. *Canadian Psychology/Psychologie canadienne*, 26(1):1.

van Mulukom, V., Schacter, D. L., Corballis, M. C., and Addis, D. R. (2013). Re-imagining the future: repetition decreases hippocampal involvement in future simulation. *PloS one*, 8(7):e69596.

Wang, Q., Luo, S., Monterosso, J., Zhang, J., Fang, X., Dong, Q., and Xue, G. (2014). Distributed value representation in the medial prefrontal cortex during intertemporal choices. *Journal of Neuroscience*, 34(22):7522–7530.

Yaari, M. E. (1965). Uncertain lifetime, life insurance, and the theory of the consumer. *The Review of Economic Studies*, 32(2):137–150.

Zavagnin, M., De Beni, R., Borella, E., and Carretti, B. (2016). Episodic future thinking: the role of working memory and inhibition on age-related differences. *Aging clinical and experimental research*, 28:109–119.

Chapter 2

Compulsivity and impulsivity in Bayesian reinforcement learning models of addiction: a computational critique of the habit theory

Isaac Kinley & Suzanna Becker

Introductory note

Mathematical models are invaluable to every field of science, and psychology is no exception. Formalizing a theory using a mathematical model makes that theory unambiguous and gives us a rigorous way to draw out its implications and hidden assumptions. This chapter draws on mathematical models of decision making from the reinforcement learning literature to formalize the habit theory of addiction. Addiction is characterized by impulsivity (a strong preference for immediate rewards) and compulsivity (an insensitivity to the negative consequences of certain actions), and the habit theory sees these as arising from an extreme dominance of habit over goal-directed action. In the chapter's model, goal-directed action is analogized to model-based reinforcement learning, in which a simulated agent learns an internal model of its environment and uses this to plan ahead. In contrast, habit is analogized to model-free reinforcement learning, in which a simulated agent lacks such a model and simply repeats previously rewarded actions. We can then test whether the model-free system is more susceptible than the model-based system to impulsivity and compulsivity, as the habit theory of addiction would suggest.

Episodic future thinking interventions encourage detailed imagination of future contingencies, and so can be seen as targeting the model-based system. The chapter uses a Bayesian formulation of reinforcement learning in which the model-based and model-free systems track the uncertainty of their reward expectancies, and behaviour is controlled by the system that exhibits lower uncertainty. Thus, with a Bayesian model, we can explore how an episodic future thinking intervention could not only encourage the model-based system to favour future rewards, but also increase the model-based system's confidence in these rewards and thus its control over behaviour. This can provide a framework for further empirical exploration of the effect of episodic future thinking on delay discounting.

The main author (I.K.) designed and ran the simulations, and primarily wrote the paper. It was commissioned for inclusion in a volume provisionally titled *Habits: Their Definition, Neurobiology, and Role in Addiction* edited by Youna Vandaele and to be published by Springer Nature: Kinley, I. & Becker, S. (2023). Compulsivity and impulsivity in Bayesian reinforcement learning models of addiction: a computational critique of the habit theory. In Y. Vandaele (Ed.), *Habits: Their definition, neurobiology, and role in addiction.* Springer Nature.

Abstract

Addiction is sometimes argued to represent an extreme dominance of habitual behaviour, driven by stimulus–response associations, over goal-directed behaviour, involving planning based on action–outcome contingencies. In this chapter, we formalize a recent elaboration on this "habit theory" of addiction using Bayesian reinforcement algorithms as models of habit and planning. In these models, compulsivity and intertemporal impatience, both considered important elements of addiction, can arise through a dominance of habit over planning, but only on the assumption that the planning system does not overvalue addictive rewards. That is, the habit theory of addiction implicitly assumes that the planning system, in contrast to the habit system, ascribes appropriately high values to non-addictive rewards and appropriately low values to addictive rewards. However, recent evidence suggests that goal-directed overvaluation of addictive rewards is a key driver of addiction, which presents a significant challenge for the habit theory. We discuss whether this challenge will prove to be insurmountable.

2.1 Introduction

The constructs of compulsivity and intertemporal impatience (a strong preference for immediate over delayed rewards, which is an important facet of impulsivity) are related but distinct: both appear to be characterized by a relative lack of consideration of future consequences, but while intertemporal impatience reflects a goal of receiving immediate rewards, compulsive behaviour is that which is not directed toward any apparent goal (though compulsions may temporarily relieve anxiety; Cuzen and Stein, 2014; Gillan et al., 2016b). Both constructs appear to be important elements of addiction: elevated compulsivity and intertemporal impatience have been found in both behavioural and substance addictions (MacKillop et al., 2011; Lee et al., 2019), and intertemporal impatience is correlated with the severity of substance use disorders (Amlung et al., 2017). Moreover, the "habit theory" of addiction (Everitt and Robbins, 2005, 2016) arises from the observation that addictive substances can elicit compulsive seeking behaviour in animal models that persists even in the face of aversive consequences such as electrical foot shocks, providing an analogue to the apparent consequence-insensitivity seen in human addiction.

Bayesian models of behaviour have been proposed to explain various aspects of addiction. These models posit decision-making agents that maintain beliefs about various parameters in their environment and, crucially, varying degrees of uncertainty about said beliefs. Addiction is proposed to arise in these models through some dysfunction of precision, the inverse of uncertainty. On the one hand, Friston (2012) argues that compulsivity can arise when repeated exposure to an addictive reward creates hyper-precise beliefs about the value of reward-seeking behaviours. These behaviours then become inflexible, such that they continue even when they are no longer optimal (for example, because the reward is no longer available). On the other hand, Schwartenbeck et al. (2015) argues that intertemporal impatience can manifest due to hypo-precise beliefs about the value of a long-term behavioural policy of deferring immediate reward. Such a policy of deferral is then unlikely to be sustained, resulting in intertemporal impatience.

As pointed out by Mollick and Kober (2020), these models appear to be in tension. It is unclear how compulsivity and intertemporal impatience can manifest simultaneously in addiction if one arises through hyper-precise beliefs and the other arises through hypo-precise beliefs. In a previous paper (Kinley et al., 2022), we argued that the two Bayesian models might be reconciled through an appeal to the aforementioned habit theory of addiction. This theory relies on the distinction between habitual behaviour, which is driven by stimulus–response associations, and goal-directed behaviour, which involves planning on the basis of known action–outcome contingencies (Dolan and Dayan, 2013). According to the habit theory of addiction, addictive behaviours arise through an extreme dominance of habit over goal-directed action (Everitt and Robbins, 2005, 2016). This is proposed to explain why these behaviours can persist despite negative consequences and explicit goals of cessation. The balance of control between habit and goal-directed behaviour is thought to be arbitrated by the precision of each system's reward predictions, such that the more precise system is more likely to drive decisions (Daw et al., 2005; Lee et al., 2014). Thus, we argued (Kinley et al., 2022) that addiction could be characterized by both increased precision in the habitual system and reduced precision in the goal-directed system. This was argued to shift the balance of control in favour of habit, producing compulsive reward-seeking behaviours as well as intertemporal impatience, as the planning necessary to defer gratification no longer drives decision making.

The distinction between habitual and goal-directed behaviour corresponds closely to the distinction in reinforcement learning between model-free and model-based algorithms (Sutton and Barto, 2018). Model-free algorithms are those in which no internal representation of the environment is learned. Instead, the value of each action in each state is represented by a separate parameter, and these parameters are updated according to experience. In contrast, model-based algorithms learn an internal model of the environment and the dependence of outcomes on previous actions. Thus, whereas the decisions of model-free algorithms are driven by static parameters, model-based algorithms are able to dynamically search through a tree of action–outcome contingencies to select the optimal next action. As models of true human decision making, these algorithms are gross simplifications, but are nonetheless valuable as polar extremes defining a continuum along which human decision making does exhibit meaningful variance. In this chapter, we will formalize the account put forth in Kinley et al. (2022), using reinforcement learning models of decision making to explore how the balance between goal-directed and habitual behaviour could give rise to compulsivity and intertemporal impatience. In keeping with our prior emphasis on precision, we will use Bayesian variants of standard model-free and model-based algorithms, which maintain beliefs about action values along with varying degrees of uncertainty about these beliefs. Formal models such as these are valuable in that they can reveal assumptions implicit in "verbal theories" (i.e., theories expressed only in words; van Rooij and Blokpoel, 2020). Made explicit, these assumptions and the theory that rests on them can then be critically evaluated. As we will see, the account put forth in Kinley et al. (2022) did indeed contain various implicit assumptions, the validity of which we will discuss in this chapter.

2.2 Bayesian reinforcement learning

In this section, we give an intuitive description of the reinforcement learning models used throughout this chapter. The mathematical details can be found in the appendix and the source code for all simulations and visualizations can be found at `github.com/kinleyid/bayesian-rl-chapter`.

2.2.1 The environment

The simulated environments in which reinforcement learning algorithms operate are called Markov decision processes (MDPs). An MDP consists of a set of possible states and a set of actions available at each state. Each action from a given state is associated with a probability distribution over subsequent states. That is, a particular action could lead always to the same subsequent state or could lead probabilistically to one of several possible subsequent states. The MDPs we will consider will have an "episodic" structure, meaning that, at a subset of "terminal" states, no action is possible any longer and the agent instead begins back in the initial state. Moreover, in the MDPs we will consider, arrival in each terminal state may be accompanied by a state-specific reward, but no reward will be available at nonterminal states. The goal of a reinforcement learning algorithm is to learn the long-run reward associated with each action in each state. Actions are more likely to be taken when they have higher estimated long-run rewards (Eq. 2.2). As we will see next, model-free and model-based learners take different approaches to estimating these long-run rewards.

2.2.2 Model-free learner

The model-free algorithm we will use is based on an approach called Q-learning (Watkins and Dayan, 1992). A classical Q-learning agent has a set of cached reward estimates for each action in each state that are initialized to 0. Upon taking some action, the agent arrives in a subsequent state and receives the associated reward, if any. Rewarded actions are reinforced (i.e., assigned higher cached reward values) and are more likely to be repeated when the same state/stimulus is encountered in the future (Fig. 2.1, Eq. 2.3). The goal of the agent is to estimate not merely the

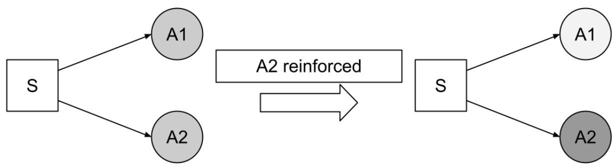

Figure 2.1: An illustration of model-free learning. Initially the agent estimates the values of actions A1 and A2 to be 0 and so is equally likely to take each of them (left, represented by equal shading). If a reward is received after taking action A1, this action becomes more valued and thus more likely upon encountering state/stimulus S (right, represented by darker shading for A2 and lighter shading for A1). Note that the agent has no explicit expectation about the subsequent states likely to be reached after any action.

immediate rewards associated with actions, but their long-term values. However, these cannot be directly observed ahead of time. To estimate the long-term value of an action that was just taken, the agent combines the immediate reward of that action with its current estimate of the long-term reward of the ostensible best subsequent action. The difference between this combined reward value and the agent's current estimate of the value of the just-taken action can be thought of as a prediction error learning signal. If the learning signal is positive, the estimated reward value of the just-taken action is revised upward, and vice versa.

This form of learning is recursive, with updates to estimated action values dependent on the same estimates for subsequent actions. Thus, learning signals can be thought of as slowly propagating backward through a chain of prior actions. This learning can be accelerated using "eligibility traces", which track the recency of actions and thus the eligibility of their associated value estimates for updates based on new observations. When a given action is taken, its eligibility trace is set to 100% and then decays over time (Eq. 2.6). The estimate of the action's value is then updated based on subsequent actions, using learning signals that are scaled down according to the action's eligibility trace (Eq. 2.7). Thus, the estimated

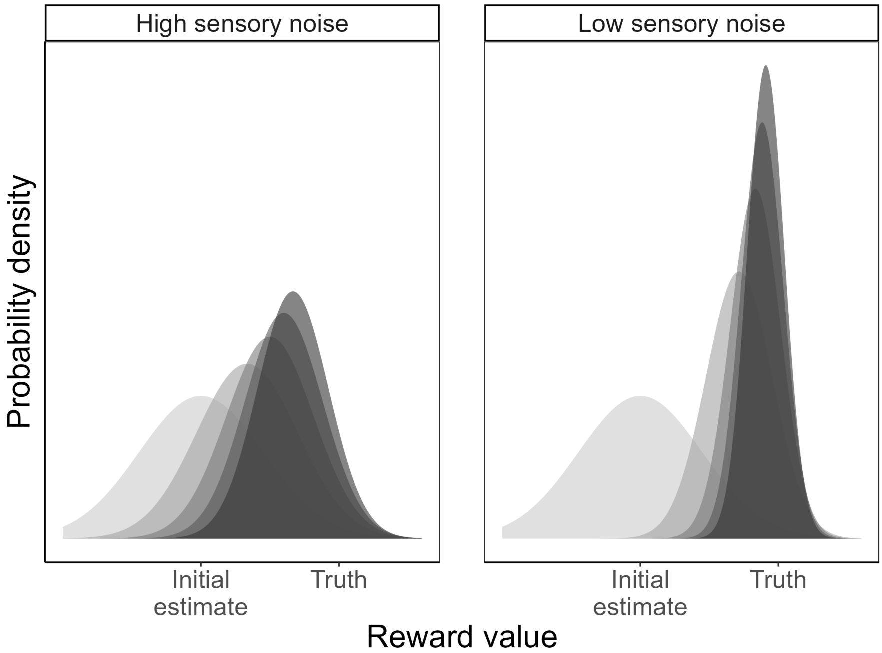

Figure 2.2: The agent's belief about the long-term reward associated with an action is encoded in a probability distribution. The change in this belief over a series of observations is visualized here by overlaying a series of increasingly dark gray distributions. Initial uncertainty, seen in the wide spread of the lightest gray distribution, decreases as observations accumulate and the agent's belief converges on the true parameter value. When sensory noise is lower (right panel compared to left panel), a greater weight is accorded to new observations, and thus the initial uncertainty decreases more quickly.

reward of an action is "eligible" for larger updates based on immediately following actions compared to more distant downstream actions. The rate of decay of the eligibility traces is denoted by the parameter λ. After each step in the task, the eligibility traces for each action are multiplied by λ. Thus, if $\lambda = 0.6$, the eligibility traces decay by 40% after each step. For higher values of λ, the agent has a longer "memory" for preceding state–action pairs. When $\lambda = 0$, the situation is equivalent to using no eligibility traces at all.

In the current setting, we are interested in tracking the uncertainty of the agent's reward estimates. Thus, we use a Bayesian variant of Q-learning that builds on the one introduced by Dearden et al. (1998). Rather than using a static estimated value for the reward associated with an action, these agents assume that the reward is a random variable that follows a normal probability distribution (Eq. 2.8). The goal is then to estimate the mean of this distribution while tracking its uncertainty about the mean. With each new observation, this uncertainty decreases (Fig. 2.2, left panel).

Importantly, this model includes a "sensory noise" factor that controls the precision (inverse variance) of observations. This factor determines the weight accorded to new observations relative to prior expectations when updating reward estimates; when new observations are highly precise, they are weighted more heavily (Fig. 2.2, right panel). As in the non-Bayesian formulation, learning follows a recursive calculation, with updates to reward estimates dependent on provisional estimates for subsequent actions. Thus, uncertainty about the value of subsequent actions imposes a limit on the possible precision of estimates for earlier actions (Eq. 2.18).

Finally, to account for possible changes in the environment (e.g., old rewards may become unavailable and new rewards may appear), the agent is designed to gradually forget previous observations as their relevance declines. Specifically, the parameters controlling the agent's estimate of an action's reward exponentially decays toward its default value, and the uncertainty of this estimate gradually increases in the absence of new data (Eqs. 2.24, 2.25). As with the eligibility traces, the rate of decay of the agent's "memory" is controlled by a factor w (Eq. 2.26).

2.2.3 Model-based learner

In contrast to a model-free agent, a model-based agent attempts to learn an explicit model of its environment, consisting of estimated action–outcome contingencies (i.e., state–action–state transition probabilities) and estimated reward values for each state (Fig. 2.3). At decision time, the expected reward of each available action is computed by traversing a tree of estimated action–outcome contingencies (Eq. 2.36). That is, the model-based agent estimates the values of actions through "simulated experience" of the long-term outcomes of those actions. As we will see shortly, there are important differences between the behaviour of model-based and model-free algorithms that arise because the model-based algorithms do not *directly* estimate the values of actions, but instead estimate the values of specific outcomes and the relationships between outcomes and prior actions.

In a Bayesian formulation, the beliefs of the model-based agent, like those of the model-free agent, include some uncertainty. The model-based agent assumes that the reward associated with each terminal state follows a normal distribution with an unknown mean. Updates to its estimate of the mean of this distribution follow the same Bayesian-inferential formulation as the model-free agent's direct estimates of action values (Fig. 2.2, Eqs. 2.28, 2.29). However, eligibility traces are no longer used given that the reward associated with subsequent states is not informative about the value of a prior state considered in isolation.

To learn action–outcome contingencies, the agent keeps count of the subsequent states reached after each action in each state (i.e., of the state–action–state transitions). These counts encode a probability distribution over possible transition functions from each state to each subsequent state (Eq. 2.31). For example, in Fig. 2.4, prior to any observations, a model-based agent is uncertain about the probability that a given action will lead to one of two possible subsequent states. With increasing numbers of observations, the agent's belief converges on the true transition probability.

Like the model-free agent, the model-based agent is more sensitive to recent observations than distant ones. The estimates of the reward available at each state

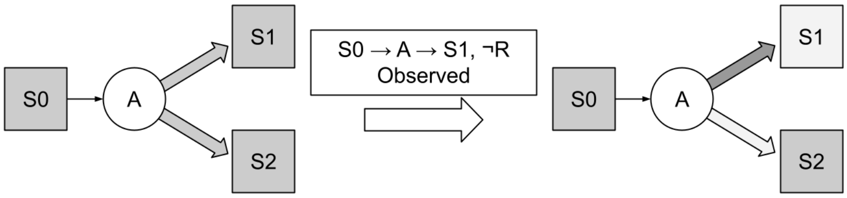

Figure 2.3: An illustration of model-based learning in a 3-state environment. The agent initially assumes that taking action A in state S0 is equally likely to lead to S1 or S2, the two possible subsequent states (left, represented by equal shading of the arrows from A to S1 and S2). After action A leads to state S1, the agent increases its estimate of the probability of this state–action–state transition (right, represented by darker arrow from A to S1 and a lighter arrow from A to S2). Also, when the arrival in state S1 is not accompanied by a reward (denoted ¬R), the agent lowers its estimate of the reward associated with state S1 (right, represented by lighter shading for S1). Note that this does not affect reward estimates for S0 or S2. Note also that estimates of transition probabilities and estimates of rewards are separate, but can be combined to estimate optimal actions.

decay toward their initial values over time, as do the state transition counts. This allows the model-based agent to be flexible in the face of changing reward values of different states as well as changing state–action–state transition probabilities.

2.3 Measuring model-based and model-free decision making

The behavioural effects of model-based or model-free control become evident in a carefully designed experimental task. One such task, introduced by Daw et al. (2011) and illustrated in Fig. 2.5, involves two stages: at the first stage, the two available actions are each principally associated with a different state in the second stage. Each action will usually (70% of the time) lead to an associated second-stage state, but occasionally (30% of the time) leads to the other second-stage state. Each second-stage state presents two actions, each of which has some probability of returning a reward. These probabilities drift independently over time such that the task requires continual learning (and continual forgetting).

To illustrate how this task differentiates between model-based and model-free strategies, suppose an agent has learned that action 2aa currently offers a high probability of reward. Suppose that, in state 1, action 1a is taken, which usually leads to state 2a. However, this time, the action leads to state 2b. The agent takes action 2bb and receives no reward, and the next trial then begins with the agent back in state 1. What is the optimal next action, according to a model-based or model-free strategy? A model-based agent understands the transition structure of the task, and thus understands that the transition from to state 2b after action

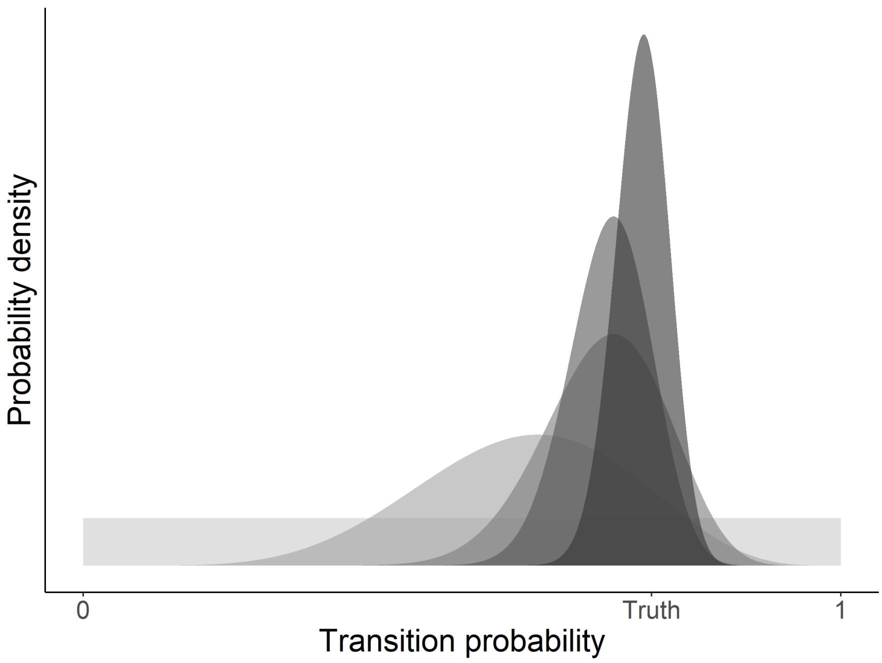

Figure 2.4: The model-based agent's belief about the probability of an action leading to one of two subsequent states is encoded in a probability distribution. The change in this belief over a series of observations is visualized here by overlaying a series of increasingly dark gray distributions. Initially, the distribution is uniform, meaning the agent believes the action may, with equal probability, have anywhere from a 0% chance to a 100% chance of leading to leading to the given subsequent state. With increasing observations, the agent's estimate of this probability converges on its true value of 75%.

1a was something of a fluke. Still hoping to end up in state 2a, from which the rewarding action 2aa is available, this agent would likely take action 1a again, *even though action 1a did not lead to reward on the previous trial.* In contrast, a model-free agent does not understand the transition structure of the task, instead knowing only that action 1a did not end up resulting in reward. Thus the agent would be more likely to take action 1b, *even though action 1b is likely to lead to the same undesired outcome as last time.*

This difference in strategies manifests in different decision profiles for model-free and model-based learners (Fig. 2.6). The probability of the model-free learner repeating the same action in stage 1 is driven by a main effect of the reward received on the last trial: rewarded actions are on average more likely to be repeated than unrewarded actions. For the model-based learner, this probability is driven by an interaction between the previous reward received and whether the previous stage 1-to-stage 2 transition was common or rare: as described above, knowledge of the transition structure of the task is used in order to maximize reward.

Humans exhibit decision profiles between the two extremes illustrated in Fig. 2.6 (Daw et al., 2011). The degree to which an individual's decisions are driven by the aforementioned main effect or interaction can quantify the contribution of

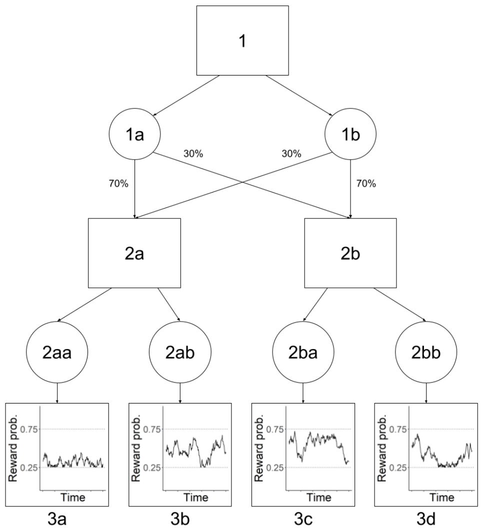

Figure 2.5: Abstract representation of the 2-stage decision making task (Daw et al., 2011). Rectangles represent states and circles represent actions. Each action in the first stage is most likely (but not certain) to lead to a different second-stage state. Each of the 4 possible actions in the second stage leads to a particular terminal state with a reward probability that independently drifts according to reflected Brownian motion over the course of the task such that continuous learning (and continuous forgetting) is required for optimal performance.

model-free and model-based strategies, respectively. This has enabled a large body of research revealing correlates of individual differences in reinforcement learning strategies. For example, a bias toward model-free decision making appears to be generally associated with compulsivity (Voon et al., 2015; Gillan et al., 2016a) and higher working memory capacity is associated with more model-based decision making in younger adults (Eppinger et al., 2013). Fig. 2.6 demonstrates that the Bayesian models used here, while novel in some respects, nonetheless conform to the operational definitions of model-based and model-free learning used in behavioural research, thus allowing us to connect the computational results described in this chapter to research with humans.

2.4 Model-free decision making, habit, and compulsivity

The behavioural hallmark of habitual control is devaluation insensitivity, i.e., the continuation of a conditioned behaviour after the reinforcer has been devalued (Dickinson, 1985). In a typical devaluation experiment using animal models, a behaviour is trained using some reinforcer over an extended period of time, after which the reinforcer is devalued, such as by being paired with a nausea-inducing injection of lithium chloride (Adams, 1982). Such an intervention distinguishes between behaviour driven by stimulus–response associations (i.e., habit)

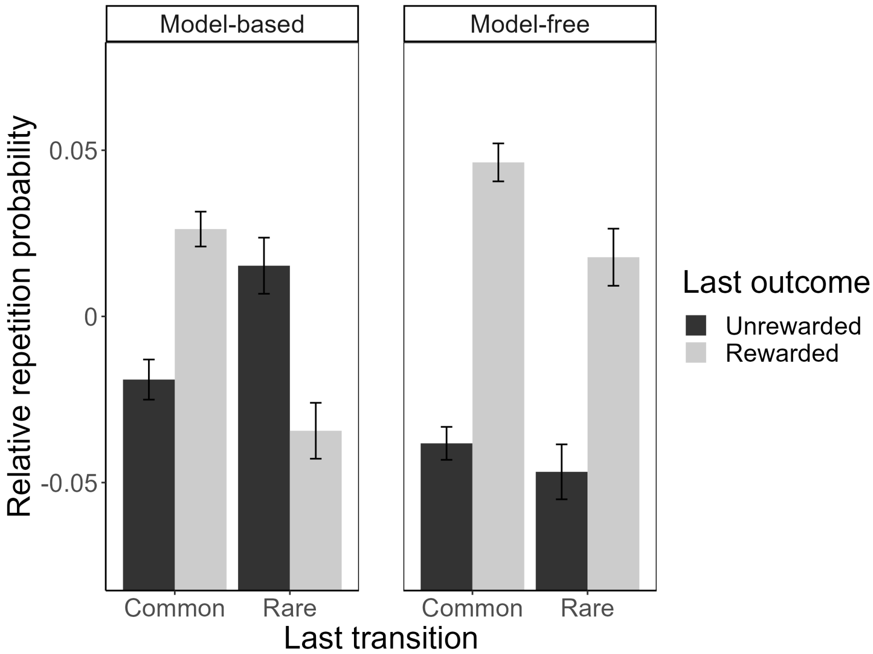

Figure 2.6: Probability of repeating the first-stage action that was taken on the previous trial (relative to each agent type's baseline probability of such repetition) as a function of reward received on the previous trial and whether the previous transition from first- to second-stage was common (1a to 2a or 1b to 2b; 70% probability) or rare (1a to 2b or 1b to 2a; 30% probability). See Fig. 2.5. For the model-free agent, repetition probability is driven by a simple main effect of previous reward, whereas the model-based agent exhibits an interaction between previous reward and previous transition probability.

and behaviour driven by learned action–outcome association (i.e., goals; Rozeboom, 1958): if an animal is acting on the basis of action–outcome contingencies, and the expected outcome is no longer valued, then the associated action is unlikely to be taken. However, devaluing a reward does not weaken the association between the conditioned behaviour and the conditioned stimulus. Thus, behaviour driven by stimulus–response associations will continue after devaluation. Instrumental behaviours are sensitive to devaluation after brief but not extended training, implying that initially goal-directed behaviours become habitual after sufficient repetition (Dickinson, 1985).

2.4.1 Asymptotic uncertainty and the transition from planning to habit

Following Daw et al. (2005), we can model a devaluation experiment using the Markov decision process shown in Fig. 2.7. At each non-terminal state, the agent can either press a lever or attempt to access the reward. The agent begins in the initial state s_0 and, if the lever is pressed, enters a "reward delivered" state. From here the agent can access the reward to enter the "consumption state". If the agent deviates from this sequence of actions, it reaches a terminal state in which no reward is received. After the receipt or non-receipt of a reward, the episode ends and the task begins again.

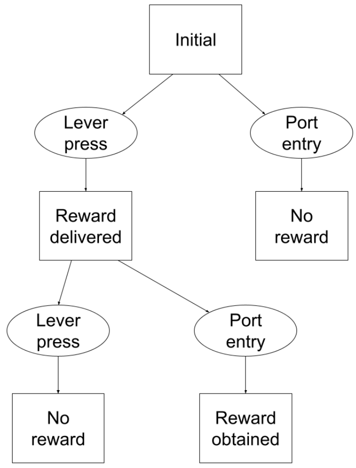

Figure 2.7: Abstract representation of a typical devaluation task (Daw et al., 2005). Rectangles represent states and ovals represent actions. The agent must press a lever in order for a reward to be delivered and then must enter the “consumption state” to obtain the reward. Any deviation from this sequence of actions leads to a terminal “no reward” state.

In computational terms, the devaluation manipulation impacts the reward value associated with the consumption state. For a model-based system, this means updating the estimated reward associated with this state to reflect the new subjective value. Thus, when planning the next action, the model-based agent simulates the now-devalued experience of the consumption state and acts accordingly. However, the model-free system does not engage in planning of the same kind, and because it has not accessed the devalued reward via the actions illustrated in Fig. 2.7, it has not had the chance to devalue those actions. Thus, the model-free learner continues the instrumental behaviour (the level press) at the same rate, while the model-based learner does not (Fig. 2.8).

Daw et al. (2005) propose an explanation for the transition from planning to habit. Early in learning, the goal-directed/model-based system makes more efficient use of information: experience of a reward is immediately used during the next decision when planning ahead to the same point. In contrast, during learning in a habitual/model-free system, reinforcement slowly propagates backward through the chain of preceding actions. However, provided that reward contingencies are static, the model-free system is asymptotically more precise, being free of the uncertainty inherent in planning forward through action–outcome contingencies. Thus, in uncertainty-weighted competition over behavioural control, the model-free system eventually wins out (Fig. 2.9).

It is worth noting that this pattern of results is sensitive to the choice of parameters for the current model. For example, using eligibility traces that do not decay

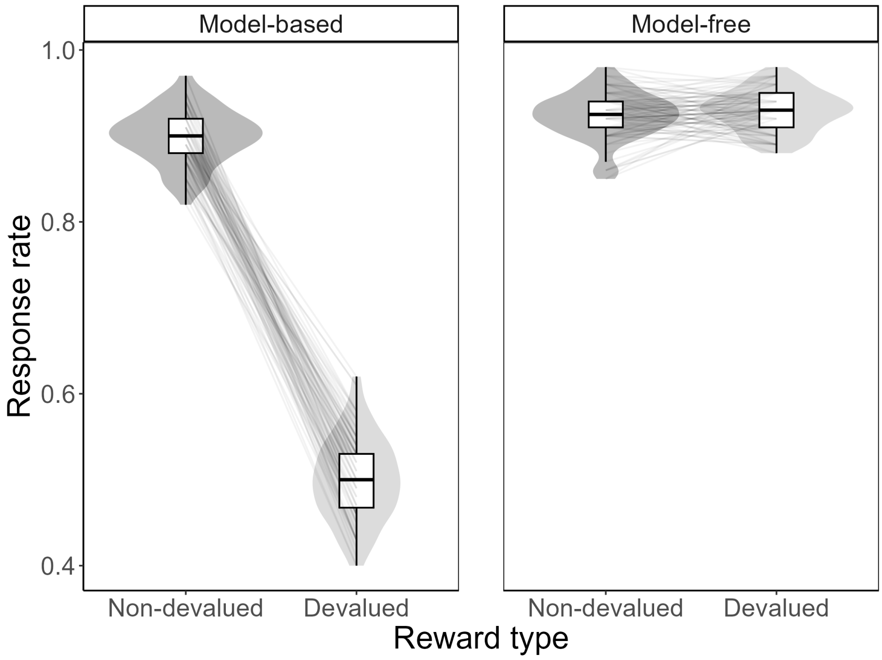

Figure 2.8: Rate of lever pressing before and after devaluation for the model-free and model-based learners. The model-based agent is sensitive to changes in reward value, reducing its rate of lever pressing after devaluation, whereas the model-free agent is not.

over time (i.e., setting $\lambda = 1$) accelerates the decrease in the model-free agent's uncertainty during the first few training iterations. Similarly, the asymptotic uncertainty of the model-based agent can be made arbitrarily close to its minimal value[1] by reducing its initial uncertainty (i.e., choosing a sufficiently small θ_0, the initial "pseudocount" of state transitions prior to any actual observations; Eq. 2.32.). Nonetheless, the current model is meant to build on the asymptotic uncertainty theory of the transition of behaviour from goal-directed to habitual. Thus, in order to take this theory as an assumption and explore its implications, the parameters of the current model were selected so as to reproduce the pattern of results reported in Daw et al. (2005).

2.4.2 Habit is not sufficient for compulsivity

The simple model described above of habit built up through extended experience is insufficient to explain the phenomenon of compulsivity for the simple reason that, after new experience with the devalued reward, the strength of the model-free agent's stimulus–response associations will weaken accordingly. A model-free algorithm is not *categorically incapable* of learning to devalue a reward; it is simply slower to do so than a model-based algorithm. That is, it is *relatively* devaluation-insensitive over a short time horizon. In contrast, compulsive be-

[1]This minimal value is $\sigma_r^2 + \epsilon$, where ϵ depends on the forgetting rate w such that $w = 1$ implies $\epsilon = 0$.

Figure 2.9: Uncertainty of each agent's reward estimate at the initial state over the course of learning. The model-based agent initially makes more efficient use of experience and learns more quickly, but the model-free agent is asymptotically less uncertain.

haviours are marked by their prolonged persistence, even after extended exposure to negative consequences. Thus, if compulsivity is to be explained by a dominance of habit, then we must additionally assume that there is something impeding the learning that would usually occur in the habit system and lead to the cessation of the compulsive behaviour.

Many addictive substances increase mesolimbic dopamine transmission (Pierce and Kumaresan, 2006), and some evidence also implicates dopamine transmission in behavioural addictions (Boileau et al., 2014; Poletti et al., 2013) (though there does not appear to be a straightforward overlap between behavioural and chemical addictions; Sinclair et al., 2016). Classically, phasic dopamine transmission is thought to encode reward prediction error (δ in the current learning equations; Schultz et al., 1993; Schultz, 1998): dopamine transmission increases when an unexpected reward is received, but gradually declines as the reward comes to be expected. Similarly, when an expected reward fails to materialize, dopamine transmission drops below its baseline level.

Redish (2004) points out that, if dopamine encodes a reward prediction error, and if addictive substances continue to produce such a signal even after extended experience, then subjective valuations of addictive rewards can in theory increase without bound. Thus, addictive rewards become overvalued and reward-seeking behaviours become inelastic to costs imposed on them. Equivalent results could be obtained for the current learning equations. However, an alternative view sees dopamine as instead encoding the *precision* (i.e., inverse variance) of prediction errors (rather than their magnitude) and thereby determining their influence when updating beliefs about the environment (Friston et al., 2012).[2] According to this precision-based view of dopamine, Friston (2012) proposes a model of addiction in which prolonged exposure to reward in a hyper-dopaminergic state creates hyper-precise beliefs about the value of the behaviour that leads to it. This behaviour then persists even after it no longer leads to reward, because the agent's belief in its value is strong enough to overshadow experience.

However, both this model and the model of simple habit described above fail to account for a crucial aspect of compulsivity in addiction, namely, that an addictive reward continues to be available *in conjunction with* its consequences. Thus any learning about these consequences takes place alongside the continued pathological learning effects of heightened dopamine transmission. To capture this dynamic, we can use the same task as described in Fig 2.7. After some number of learning trials, we can introduce a "punishment" (a negative reward) that the agent observes *separately* from the (positive) reward. Practically, this means that the learning equations (Eqs. 2.21–2.23, 2.28, 2.29) are applied twice upon reaching the terminal state, once for the reward and once for the punishment. When the reward is addictive, the learning equations use a lower sensory noise parameter σ_r, reflecting heightened dopamine transmission. Thus, a model-free agent over-weights the observation of the addictive reward and does not sufficiently devalue the sequence of actions leading to it when a punishment is added (Fig. 2.10, left panel).

2.4.3 Is habit necessary for compulsivity?

In order to capture the long-term compulsivity seen in addiction, we are required to posit a process of pathological reward learning rather than merely an imbalance between model-free and model-based control of behaviour. Indeed, Everitt and Robbins (2016) acknowledge that dominance of habit alone is not a sufficient condition for compulsivity to arise in addiction. However, the model of compulsivity presented here may prompt us to further ask whether a dominance of habit is even a *necessary* condition for compulsivity: as shown in Fig. 2.10 (right panel), the model-based learner also exhibits insensitivity to punishment in the presence of an addictive reward. If a model-based agent can exhibit compulsivity, it follows that a dominant model-free system is not necessary to explain compulsive behaviour.

This poses a significant problem for a habit-based account of addiction. The only apparent recourse is to posit that pathological reward learning in addiction is limited to the habit system, and some models do exactly this. For example, Keramati and Gutkin (2013) show how addiction can arise through overvaluation of addictive rewards specifically within the lower levels of a hierarchical reinforcement learning architecture, where lower levels are identified with neural structures underlying model-free learning and higher levels with structures underlying model-based learning (Haruno and Kawato, 2006). The Bayesian model presented here

[2] On this view, dopamine transmission can be taken to encode $1/\sigma_r^2$ rather than δ in the current learning equations. Thus, with increased dopamine transmission, σ_r^2 decreases and α, the "learning rate", increases.

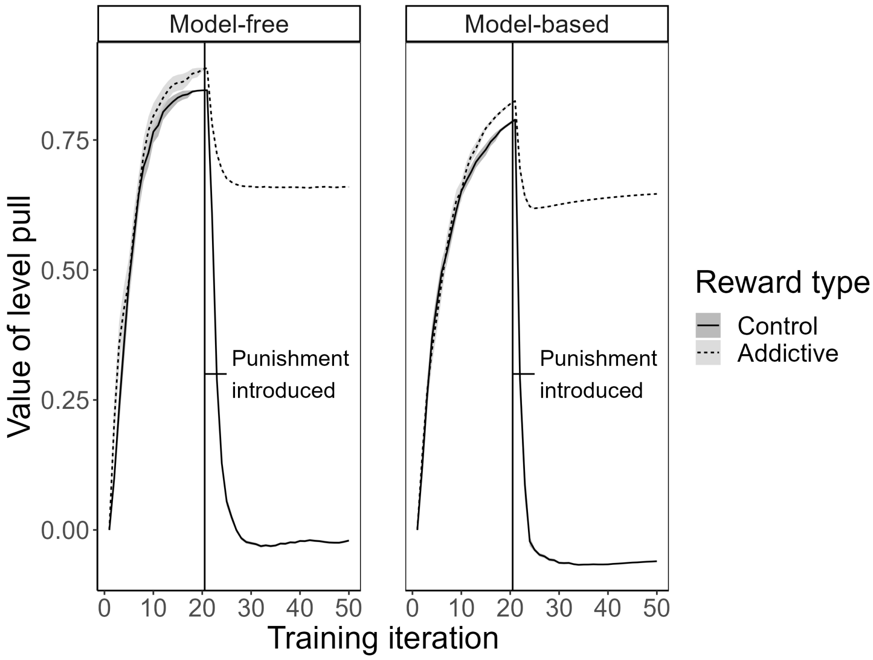

Figure 2.10: Subjective valuations for the lever press action in the initial state over the course of an experiment. After the 20th training iteration, a punishment is introduced to accompany the reward. When the lever press leads to the non-addictive reward, it is quickly devalued. However, the lever press is still considered a valuable action when it might lead to the addictive reward. This pattern of results holds for both model-free and model-based agents.

could similarly be altered so that the model-based agent is less susceptible to over-valuing addictive rewards. For example, if the model-based system maintained separate valuations of the reward and the punishment, then the estimated value of the addictive reward could be hyper-precise without overshadowing the negative value of the punishment. We will postpone for now the discussion of whether such an alteration to the present computational model to avoid overvaluation in the planning system would be justified by existing evidence, and instead simply note that the habit theory of addiction requires us to assume overvaluation of the addictive reward in the habit system and an absence of overvaluation in the planning system.

In the next section, we explore the intertemporal facet of addiction. In the task used here, the consequences of addictive behaviour are immediate, whereas in more realistic contexts, these consequences (and the benefits of abstinence) are typically delayed (Ognibene et al., 2019). We will next explore whether the strong preference for immediate over delayed rewards often seen in addiction can be explained by a dominance of habit over planning.

2.5 Model-based decision making and intertemporal impatience

Humans and other animals systematically discount the value of future rewards as a function of their delay. The most commonly used function to describe this discounting is a hyperbolic curve (Mazur, 1987):

$$V_{\text{subj}} = \frac{V_{\text{obj}}}{1 + kD} \tag{2.1}$$

where V_{subj} is the subjective value of a reward, V_{obj} is its objective value, D is the reward's delay (arbitrary units), and k is an individual differences parameter that quantifies intertemporal impatience: for higher values of k, a reward's value decays more steeply as a function of its delay. With a delay of 0, the subjective value V_{subj} matches the objective value V_{obj}.

Numerous studies have associated steeper discounting with addiction severity (for a review, see Amlung et al., 2017), and some evidence even suggests that steep discounting could be a risk factor for developing addiction (Audrain-McGovern et al., 2009). Bickel et al. (2014) argue that excessively steep delay discounting is so closely connected to addiction as to qualify as a behavioural marker, predicting the risk of developing addiction and indexing its severity and response to treatment. Beyond addiction, steeper discounting is associated with a range of unhealthy behaviours (Daugherty and Brase, 2010) and is thought to be a trans-disease process contributing to other disorders such as ADHD and depression (Bickel et al., 2014, 2019). Given these findings, it seems tempting to suggest that shallower discounting is always more desirable, or even that discounting is irrational. However, some discounting is necessary given the inherent uncertainty of the future—there is always some risk that a delayed reward will not materialize (Sozou, 1998). Even if future rewards are guaranteed, it cannot be rational to have an infinite temporal horizon of self-interest as a finite being (Yaari, 1965). Moreover, some discounting is necessary to make reinforcement learning with non-episodic MDPs computationally tractable, so that the expected future reward is finite (Naik et al., 2019).

2.5.1 A reinforcement learning model of intertemporal choice

In the current setting, we are interested less in temporal discounting that is guaranteed because it is built into the basic learning equations as a numerical constant (Eq. 2.3), and more in temporal discounting as an emergent phenomenon arising due to the environment and the form of learning and decision making (model-free or model-based) that takes place. Thus, we do not build discounting into the learning equations (i.e., we set $\gamma = 1$; Eq. 2.3).

We can design an impulsivity assay as shown in Fig. 2.11. The agent begins in an initial "0th" waiting state, from which it can decide to take a small immediate reward r_{imm} or can decide to wait for a larger reward later $r_{\text{del}} > r_{\text{imm}}$. The latter action will bring the agent either to the next (1st) waiting state or to a terminal no-reward state. The same options are available from the next waiting state, such that the agent can progress through a series of waiting states or can, at any point, opt for r_{imm}. At the final waiting state r_{del} becomes available. Learning in this task can be seen as a race between two processes. On the one hand, the agent must learn the value of the delayed reward through experience that can be gained only by forgoing the immediate reward. On the other hand, as the agent learns the value of the immediate reward, it is less likely to forgo it. The agent will be a steeper or shallower discounter depending on which of these processes is dominant.

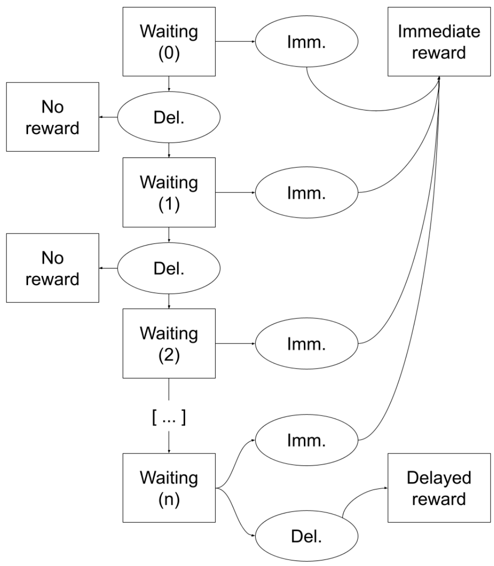

Figure 2.11: Abstract representation of an intertemporal choice task. At each of a sequence of waiting stages, the agent can either continue to wait or opt for an immediate reward. Continuing to wait is associated with a small risk of transitioning to a terminal "no reward" state. At the final waiting stage, continued waiting leads to the larger delayed reward.

2.5.2 Planning is necessary for patience

With the default parameters used for the other simulations described in this chapter, the model-free agent exhibits steeper discounting than the model-based agent (Fig. 2.12, top panel). It is worth noting that this result is sensitive to the choice of parameters: maximizing the model-free agent's eligibility traces (i.e., setting $\lambda = 1$; Eq. 2.6) and increasing the model-based agent's uncertainty about state transitions (i.e., setting θ_0 to a high value; Eq. 2.32) reverses the pattern of results (Fig. 2.12, middle panel). This is because higher eligibility traces better enable the model-free agent to learn the value of a long chain of "wait" choices and whereas increasing the model-based agent's uncertainty about state transitions increases it's asymptotic uncertainty about the value of waiting for the delayed reward.

However, increasing the value of θ_0 places a somewhat arbitrary impediment on the model-based agent's learning. By default, θ_0 had a value of 1 in the simulations presented in this chapter, which encodes a uniform distribution over possible

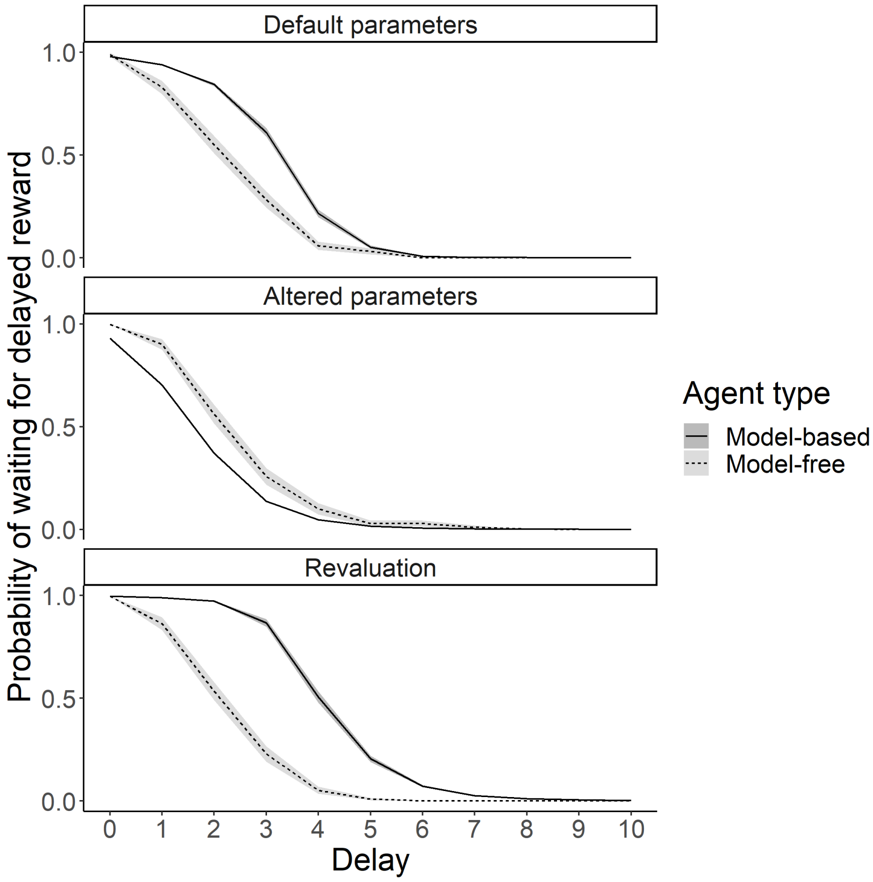

Figure 2.12: Probability of waiting the requisite number of times to receive the larger delayed reward as a function of agent type and number of waiting steps. The top and middle panels demonstrate that the choice of parameters determines whether the model-based or model-free agent exhibits greater intertemporal impatience (though the differences are not dramatic). In the bottom panel, the model-based agent subjectively assigns the delayed reward its full value and the immediate reward a value of 0, while the model-free agent again uses the default parameters from the top panel.

transition functions from each action to each subsequent state (Fig. 2.4). This implies maximal uncertainty on the part of the agent prior to any observations which, according to the principal of maximum entropy (Jaynes, 1957a,b), is the best representation of the agent's initial (lack of) knowledge. Thus, the choice of parameters producing shallower discounting in the model-based learner is not an arbitrary one.

Moreover, we can additionally assume that the model-based system's reward valuations do not come from direct experience alone. Instead, as the term "goal-directed" suggests, model-based reward valuations can be assumed to arise from the agent's goals, which may be liable to change. Story et al. (2014) propose a reinforcement learning model of intertemporal impatience in which, after extended experience with an immediate reward, the agent "re-evaluates [its] goals" such that the delayed reward is now valued highly and the immediate reward devalued by the model-based system. We can model this by adjusting the model-based system's estimate of r_{del} to the correct value and its estimate of r_{imm} to 0. With this revaluation, the model-based learner exhibits considerably shallower discounting than the model-free learner, even when the latter's parameters chosen are for minimal discounting (Fig. 2.12, bottom panel). That is, the model-free agent always exhibits relatively steep discounting and model-based planning is necessary for intertemporal patience. The discounting that the model-based learner does still show is due to uncertainty about the value of waiting.

In our previous work (Kinley et al., 2022), we described factors that might increase the uncertainty in reward estimates produced by a goal-directed system in addiction. For example, some addictive substances are associated with grey matter reductions in the prefrontal cortex (Wang et al., 2012; Matochik et al., 2003), an area broadly associated with model-based control (Huang et al., 2020). Similarly, stress is a known precipitant of relapse (Mantsch et al., 2016) and has been shown to reduce model-based control (Radenbach et al., 2015). The process of generating low-uncertainty simulations of future rewards is argued to be cognitively taxing (Gershman and Bhui, 2020), and thus could be adversely affected by stress and prefrontal grey matter reductions. There is, then, a certain overdetermination to the intertemporal impatience that could arise through uncertainty in the model-based system: such uncertainty would not only shift the balance of control to the (more impatient) model-free system, but would also make the model-based system itself exhibit steeper discounting.

2.5.3 Is planning sufficient for patience?

Given that model-free habits lead to steep discounting, an agent that exhibits intertemporal patience must do so through model-based planning. But is such planning *sufficient* for patience? That is, does model-based control necessarily lead to farsighted decisions? Arguably not: if we are willing to grant that model-based valuations might change to reflect a goal of obtaining the delayed reward (Fig. 2.12, bottom panel; Story et al., 2014), there is no reason in principle why a model-based agent could not also preferentially value immediate rewards and exhibit steep discounting despite intact planning abilities. Thus, steep discounting might reflect the dominance of a short-sighted habit system over a future-oriented planning system, but in principle might instead reflect the operation of a short-sighted planning system. If the impulsivity seen in addiction is to be explained as a dominance of habit over planning, as we argued in our elaboration on the habit theory of addiction (Kinley et al., 2022), we are required to assume that the planning system does not overvalue immediate rewards. We will discuss shortly whether such an assumption is warranted by existing evidence. In the meantime, we will note that a psychological intervention widely used to decrease discounting both in healthy participants and in individuals with addiction appears to target the model-based system. This intervention is called "episodic future thinking" and, as we argued previously (Kinley et al., 2022), its effectiveness might be understood in terms of decreased certainty about future rewards within the model-based system.

We will next elaborate on this proposal with reference to the computational model presented here.

2.5.4 Episodic future thinking in light of the model

Episodic future thinking (the mental simulation of the personal future; Atance and O'Neill, 2001) is arguably one expression of model-based planning in humans (Kinley et al., 2022). Imagining a future experience enables a process of valuation (D'Argembeau, 2013) similar to that which occurs in model-based decision making when a trajectory of future states is simulated. Indeed, the hippocampus, which in rodents encodes spatial representations of possible paths forward from a decision point (Johnson and Redish, 2007), is critical in humans for both episodic future thinking (Hassabis et al., 2007) and model-based control (Vikbladh et al., 2019).

Structured episodic future thinking tasks have been established as reliable methods of reducing intertemporal impatience (Rösch et al., 2022) and have been explored on this basis as interventions for addiction (Patel and Amlung, 2020; García-Pérez et al., 2022; Snider et al., 2016; Forster et al., 2021). These interventions typically encourage vivid, detailed, and emotionally positive imagery of the future, which can be interpreted in terms of the model-based learning equations presented here. For example, imagining a more positive event could be seen as simulating a future state with a higher value of $\hat{\mu}_r$ (Eq. 2.27). Similarly, vivid and detailed imagery could reflect precise expectations of what a future state will entail (i.e., a low value of $\sigma^2_{\hat{\mu}_r}$; Eq. 2.27) and perhaps a greater degree of certainty that waiting for the immediate reward will not lead to the "no reward" end state (encoded in the relative values of $\boldsymbol{\theta}$ introduced in Eq. 2.30).

As illustrated in Fig. 2.13a, precise expectations of highly valuable future rewards can produce arbitrarily shallow discounting in the model-based agent. Fig. 2.13b demonstrates that the episodic future thinking intervention also lowers the agent's uncertainty. Thus, an episodic future thinking intervention not only decreases discounting within the model-based system, but also makes this system more likely to win out in uncertainty-weighted competition with its model-free counterpart. The expected overall result of this dynamic is shallower discounting, as seen in empirical episodic future thinking interventions (Rösch et al., 2022).

We can speculate as to how the effects of this episodic future thinking intervention might occur in other implementations of model-based learning. The learning equations used here assume that planning exhaustively searches possible future trajectories, but this is somewhat unrealistic—outside of simple episodic tasks of the type used here, exhaustive planning is prohibitively computationally intensive. An alternative family of approaches instead compute expected values for only a sample of possible future trajectories (Moore and Atkeson, 1993; Dearden et al., 1999). A structured episodic future thinking intervention could be seen as biasing this sample, such that certain (vivid, positive) future events are over-represented and thus appear to be highly likely. Indeed, experimental evidence suggests that repeatedly imagining a future event increases subjective estimates of its probability of actually occurring (Szpunar and Schacter, 2013). As with the results shown in Fig. 2.13, a future state that is both rewarding (because imagined positively)

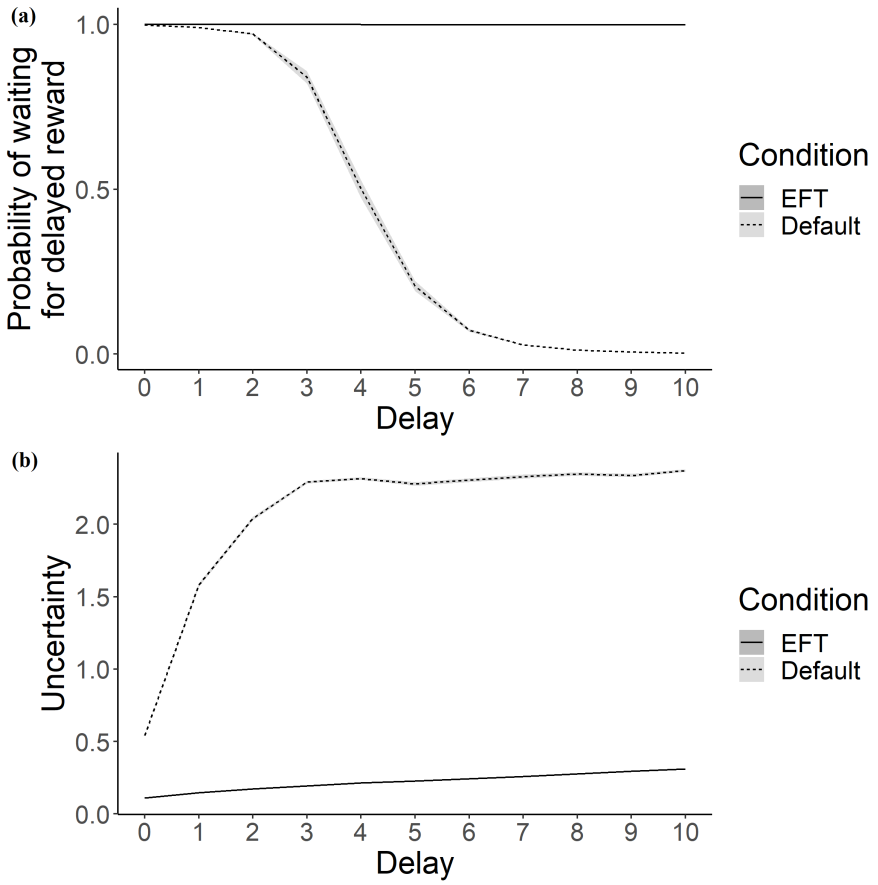

Figure 2.13: Effects of the episodic future thinking (EFT) intervention in the model-based agent. (a) The EFT intervention can produce arbitrarily shallow discounting with an appropriate change of parameters, particularly to the agent's estimate of the uncertainty of the delayed reward via the hazard rate. (b) Without the EFT intervention, the model-based agent exhibits high uncertainty, especially as the delayed reward becomes increasingly distant. In contrast, with the EFT intervention, the agent maintains precise expectations of the value of the delayed reward across a range of delays, making it more likely to win out in precision-weighted competition with the model-free agent.

and probable (because imagined frequently) would be likely to be sought after to the exclusion of smaller immediate rewards.

2.6 Discussion

In this chapter, we have attempted to formalize the elaboration on the habit theory of addiction (Everitt and Robbins, 2005, 2016) put forth by Kinley et al. (2022). On this account, addiction involves hyper-precise model-free reward predictions and hypo-precise model-based reward predictions, resulting in the dominance of habit over planning and consequently both the compulsivity and intertemporal impatience seen in addiction. However, we have seen how formalizing this theory illuminates its implicit assumptions. In principle, habit is not necessary for compulsivity because a planning system can exhibit compulsive reward seeking. Similarly, planning is not in principle sufficient for patience because a planning system could preferentially value immediate rewards. Thus, to explain compulsive reward seeking and intertemporal impatience as a dominance of model-free habit over model-based planning, we are required to assume as a contingent fact that the brain's model-based system does not overvalue addictive rewards or immediate rewards even though, in principle, it could. Is this assumption justified by empirical evidence?

Hogarth (2020) marshals various lines of evidence to argue that drug choice is driven by goal-directed overvaluation of drug rewards in addiction. For example, in rodents that choose drug rewards over food, the value of drugs appears to be represented by the orbitofrontal cortex, an area implicated in goal-directed action (Guillem et al., 2018; Guillem and Ahmed, 2018). Similarly, increased drug seeking following withdrawal (Hutcheson et al., 2001) can be interpreted as goal-directed behaviour similar to the goal-directed seeking of salinated water after sodium depletion (Krieckhaus and Wolf, 1968) or goal-directed avoidance of a devalued outcome, as discussed earlier. Moreover, in human studies, economic demand (i.e., willingness to pay) for drug rewards is associated with severity of addiction (Bruner and Johnson, 2014). Crucially, economic demand is argued to reflect expected values, i.e., model-based valuations (Hogarth, 2020). Thus, contrary to the habit theory of addiction, it appears that goal-directed overvaluation of drug rewards is a key driver of addiction.

Moreover, contrary to the expectation that steep delay discounting arises through the dominance of habit over planning, evidence for a relationship between measures of model-based control and delay discounting has been mixed. One study found that steeper discounting is predictive of reduced model-based control, but only weakly ($R^2 = 0.03$; Shenhav et al., 2017), while others have found no significant relationship at all (Hunter et al., 2018; Solway et al., 2017). Similarly, while tyrosine supplementation (which can increase dopamine and epinephrine transmission) increases model-based control and reduces discounting in healthy participants (Mathar et al., 2022), exposing regular slot machine users to a gambling environment increases discounting but also increases model-based control (Wagner et al., 2022). Thus, the two constructs do not always covary. Interestingly, Hunter et al. (2018) found that participants higher in model-based control spent more time deliberating about their intertemporal choices but were no shallower than other participants in their discounting, suggesting that their choices reflected a goal-directed preference for immediate rewards. The gambling environment described by Wagner et al. (2022) might have evoked a similar goal-directed preference for immediate rewards, conducive to continued gambling.

Thus, it appears that we cannot safely assume that the goal-directed system never overvalues addictive or immediate rewards. Does this mean the habit theory of addiction is necessarily untenable? Epstein (2020) argues that, ultimately, it is unlikely that all available evidence in the addiction literature will be explained by either a variant of the habit theory or a competitor. Such a "winner-take-all" approach to theory ignores the possibility of heterogeneous psychological phenomena—for example, addiction might include a habit-driven subtype (Epstein, 2020). Similarly, certain types of choices (e.g., sequential rather than simultaneous; Vandaele and Ahmed, 2021) might lend themselves to the expression of overpowering habits. Redish et al. (2008) enumerate a range of "vulnerabilities" that exist within a decision making system comprised of habit and planning subsystems, including overvaluation by the habit subsystem, overvaluation by the planning subsystem, imbalance between planning and habit, and "incorrect" (e.g., truncated) planning, among other processes we have not considered here. There is no need to assume these vulnerabilities are mutually exclusive or occur in a specific conjunction for all individuals at all times. The computational results reported here clarify that the habit theory of addiction assumes one particular conjunction, namely, an imbalance between a habit subsystem that overvalues addictive rewards and a planning system that does not overvalue addictive or immediate rewards. However, given the foregoing discussion, we can simply add to these results the important caveat that this conjunction is likely not universal and thus that the explanatory power of the habit theory of addiction is not total.

2.7 Appendix: learning equations

In this appendix, we give the mathematical details of the reinforcement learning models used throughout this chapter. The source code for the implementation of these models is available at `github.com/kinleyid/bayesian-rl-chapter`.

2.7.1 The environment

Before describing specific reinforcement learning algorithms, it is necessary to explain the simulated environments in which they operate. These are called Markov decision processes (MDPs). An MDP consists of, first, a set S of possible states and a set $A(s)$ of actions available at each state s. Each state–action pair (s, a) is associated with a probability distribution over successor states, $p(s'|s, a)$. The MDPs we will consider will have an "episodic" structure, meaning that, at a subset of "terminal" states, no action is possible and the agent begins back in the initial state. Moreover, in the MDPs we will consider, arrival in each terminal state may be accompanied by a state-specific reward—no reward will be available at nonterminal states.

The goal of a reinforcement learning algorithm is to learn the long-run reward associated with each state–action pair (s, a), denoted $Q(s, a)$. These Q values are used to select which action to take in state s. In the algorithms we will consider here, actions are selected according to a Boltzmann/"softmax" distribution over Q values for possible actions:

$$p(a) \propto \exp \frac{Q(s, a)}{\tau} \tag{2.2}$$

where τ is a "temperature" parameter that controls the explore–exploit trade-off: the highest-value option is always most likely to be chosen, but it is most likely by less for higher values of τ. That is, for higher values of τ, the agent is more likely to explore actions it considers sub-optimal. As we will see, model-free and model-based algorithms take different approaches to learning Q values.

2.7.2 Model-free learner

A classical Q-learning agent has a set of cached Q values for each state–action pair that are initialized to 0. Upon taking action a in state s, the agent arrives at s' and observes reward $r(s')$. The state–action value is the updated as follows:

$$[Q(s,a)]' = Q(s,a) + \alpha \left[r(s') + \gamma \max_{a'} Q(s',a') - Q(s,a) \right] \tag{2.3}$$

where a' is the next action and $\max_{a'} Q(s',a')$ is the value of the optimal next action (0 if no further action is possible). The notation $[\cdot]'$ is used to denote updated values. Thus the value of $Q(s,a)$ is gradually adjusted toward the value of the observed reward plus the expected values of all subsequent rewards, discounted at each time step by a factor γ. We will set $\gamma = 1$ so that there is no discounting built into the learning equations, but we will see later that discounting can emerge even when $\gamma = 1$. The amount by which $Q(s,a)$ is adjusted can be written in terms of a prediction error δ:

$$\delta(s,a) \coloneqq r(s') + \max_{a'} Q(s',a') - Q(s,a) \tag{2.4}$$

$$[Q(s,a)]' = Q(s,a) + \alpha\delta(s,a) \tag{2.5}$$

This learning equation is recursive, with updates to the value of the current state–action pair dependent on learned values of subsequent pairs. Thus, learning signals slowly propagate backward through the chain of preceding state–action pairs. This learning can be accelerated using "eligibility traces", which track the recency of encountered state–action pairs and thus the eligibility of their associated Q values for updates based on new observations. For each state–action pair (s,a) we define an eligibility quantity $z(s,a)$. When action a is taken in state s, the value of $z(s,a)$ is set to 1. After each subsequent action, $z(s,a)$ decays by some factor λ:

$$[z(s,a)]' = \lambda z(s,a) \tag{2.6}$$

Then, for each subsequent state–action pair (s',a'), the value of $Q(s,a)$ is adjusted according to the eligibility trace $z(s,a)$ and the just-observed prediction error $\delta(s',a')$:

$$[Q(s,a)]' = Q(s,a) + z(s,a)\alpha\delta(s',a') \tag{2.7}$$

For higher values of λ, the agent has a longer "memory" for preceding state–action pairs. When $\lambda = 0$, the learning equations reduce to the ones described earlier that use no eligibility traces.

In the current setting, we are interested in tracking the uncertainty of the agent's reward estimates. Thus, we use a Bayesian variant of Q learning that builds on the one introduced by Dearden et al. (1998). Let $R(s,a)$ be a random variable denoting the total value of taking action a in state s and following an optimal policy thereafter. For simplicity, following Daw et al. (2005), we will build into the model the implicit knowledge that reward is only available at terminal states: the agent assumes that $R(s,a)$ is normally distributed with variance σ_r^2 (a "sensory noise" factor) and an unknown mean. Let $Q(s,a)$ be the agent's best estimate of $\mathbb{E}\left[R(s,a)\right]$. The agent's belief in this estimate includes some uncertainty, and is encoded in a normal probability density function over possible estimates q, centred on the best estimate $Q(s,a)$:

$$p\left(q\right) = \mathcal{N}\left(q \middle| Q(s,a), \sigma^2_{Q(s,a)}\right) \tag{2.8}$$

The variance term $\sigma^2_{Q(s,a)}$ quantifies the agent's uncertainty about $Q(s,a)$. This prior distribution is updated with experience: suppose that the agent takes action a in state s and arrives in state s'. If s' is terminal, the agent observes reward r. Then the standard Bayesian updates to the parameters of the prior distribution are as follows:

$$\left[Q(s,a)\right]'_r = \left(\frac{1}{\sigma^2_{Q(s,a)}} + \frac{n}{\sigma^2_r}\right)^{-1} \left(\frac{Q(s,a)}{\sigma^2_{Q(s,a)}} + \frac{nr}{\sigma^2_r}\right) \tag{2.9}$$

$$\left[\sigma^2_{Q(s,a)}\right]'_r = \left(\frac{1}{\sigma^2_{Q(s,a)}} + \frac{n}{\sigma^2_r}\right)^{-1} \tag{2.10}$$

where n, the number of observations, is always 1. The notation $\left[\cdot\right]'_x$ denotes values updated according to some specific observation x. As in the non-Bayesian learning equations, the latter update can be rewritten in terms of a learning rate $\alpha(s,a)$ and a prediction error $\delta(s,a)$:

$$\alpha(s,a) \coloneqq \frac{n\sigma^2_{Q(s,a)}}{\sigma^2_r + n\sigma^2_{Q(s,a)}} \tag{2.11}$$

$$\delta(s,a) \coloneqq r - Q(s,a) \tag{2.12}$$

$$\left[Q(s,a)\right]'_r = Q(s,a) + \alpha(s,a)\delta(s,a) \tag{2.13}$$

However, if s' is nonterminal, the agent observes no immediate reward. It instead "observes" only its current estimate of the expected total value of following an optimal policy from state s' onward, denoted $Q(s', a'^{*})$, where a'^{*} is the ostensible best next action $\mathrm{argmax}_{a'} Q(s', a')$. To account for uncertainty about $Q(s', a'^{*})$, when s' is not terminal, Dearden et al. (1998) suggest computing a mixture of all possible updated densities, with each one weighted by the estimated relative probability of the observation that would justify it:

$$\left[p\left(q\right)\right]' = \int_{-\infty}^{\infty} \mathcal{N}\left(q \middle| \left[Q(s,a)\right]'_x, \left[\sigma^2_{Q(s,a)}\right]'_x\right) \mathcal{N}\left(x \middle| Q(s', a'^{*}), \sigma^2_{Q(s',a'^{*})} + \sigma^2_r\right) dx \tag{2.14}$$

where $\mathcal{N}\left(x\middle|Q(s', a'^{*}), \sigma^2_{Q(s',a'^{*})} + \sigma^2_r\right)$ is the estimated probability density over values of $R(s', a'^{*})$, evaluated at x. To solve this integral, we note that

$$\mathcal{N}\left(q\middle|[Q(s,a)]'_x, \left[\sigma^2_{Q(s,a)}\right]'_x\right) = \frac{1}{\alpha(s,a)}\mathcal{N}\left(x\middle|\frac{q}{\alpha(s,a)} + Q(s,a)\left(1 - \frac{1}{\alpha(s,a)}\right), \frac{\left[\sigma^2_{Q(s,a)}\right]'_x}{\alpha(s,a)^2}\right)$$

We can then use the following identity (Coelho, 2013)

$$\int_{-\infty}^{\infty} \mathcal{N}\left(x\middle|\mu_1, \sigma_1^2\right)\mathcal{N}\left(x\middle|\mu_2, \sigma_2^2\right)dx = \mathcal{N}\left(\mu_1\middle|\mu_2, \sigma_1^2 + \sigma_2^2\right) \tag{2.15}$$

to obtain the result

$$[p(q)]' = \mathcal{N}\Big(q\Big|Q(s,a) + \alpha(s,a)\left(Q(s', a'^{*}) - Q(s,a)\right), \left[\sigma^2_{Q(s,a)}\right]'_x + \alpha(s,a)^2\left(\sigma^2_{Q(s',a'^{*})} + \sigma^2_s\right)\Big) \tag{2.16}$$

Thus, the mixture distribution $[p\,(q)]'$ is a normal probability density with the following mean and variance parameters:

$$[Q(s,a)]' = Q(s,a) + \alpha(s,a)\delta(s,a) \tag{2.17}$$

$$\left[\sigma^2_{Q(s,a)}\right]' = \left(\frac{1}{\sigma^2_{Q(s,a)}} + \frac{n}{\sigma^2_s}\right)^{-1} + \alpha(s,a)^2\left(\sigma^2_{Q(s',a'^{*})} + \sigma^2_s\right) \tag{2.18}$$

where $\delta(s,a) \coloneqq Q(s', a'^{*}) - Q(s,a)$. Thus, accounting for uncertainty about $Q(s', a'^{*})$ does not change the mean of the updated distribution, but only adds to its variance.

Defining the parameter updates in terms of prediction errors $\delta(s,a)$ allows us to extend the notion of eligibility traces to a Bayesian learner. The same eligibility factor $z(s,a)$ as defined earlier can be incorporated into updates of $\sigma^2_{Q(s,a)}$ and $Q(s,a)$ after taking action a' in state s' and subsequently arriving in state s''. First, the "learning rate" $\alpha(s,a)$ and the prediction error $\delta(s', a')$ are set to

$$\alpha(s,a) \coloneqq \frac{z(s,a)\sigma^2_{Q(s,a)}}{\sigma^2_r + z(s,a)\sigma^2_{Q(s,a)}} \tag{2.19}$$

$$\delta(s', a') \coloneqq \begin{cases} r - Q(s', a'), & \text{if } s'' \text{ is terminal} \\ Q(s'', a''^{*}) - Q(s', a'), & \text{if } s'' \text{ is nonterminal} \end{cases} \tag{2.20}$$

where r is the reward observed upon arriving in state s'' and a''^{*} is the ostensible best next action from state s'', i.e., $\mathrm{argmax}_{a''}Q(s'', a'')$. $Q(s,a)$ can then be updated as follows, regardless of whether s'' is terminal or nonterminal:

$$[Q(s,a)]' = Q(s,a) + \alpha(s,a)\delta(s', a') \tag{2.21}$$

When s'' is terminal, the update to $\sigma^2_{Q(s,a)}$ is

$$\left[\sigma^2_{Q(s,a)}\right]' = \left(\frac{1}{\sigma^2_{Q(s,a)}} + \frac{z(s,a)}{\sigma^2_s}\right)^{-1} \tag{2.22}$$

When s'' is nonterminal, this update is

$$\left[\sigma^2_{Q(s,a)}\right]' = \left(\frac{1}{\sigma^2_{Q(s,a)}} + \frac{z(s,a)}{\sigma^2_s}\right)^{-1} + \alpha(s,a)^2 \left(\sigma^2_{Q(s'',a''*)} + \sigma^2_s\right) \tag{2.23}$$

Thus, in much the way the learning rate is reduced using the eligibility factor in the classical learning equations, the weight n accorded to the observation $\delta(s', a')$ is multiplied by $z(s,a)$ in this Bayesian formulation (n does not explicitly appear in the above equations because it always has a value of 1).

Finally, some flexibility is built into the system using an exponential forgetting procedure (Daw et al., 2005). Before updating $Q(s,a)$ and $\sigma^2_{Q(s,a)}$ according to a new observation, both factors decay toward their initial values μ_0 and σ^2_0 by a forgetting factor $w(s,a)$:

$$Q(s,a) \leftarrow w(s,a)Q(s,a) + (1 - w(s,a))\,\mu_0 \tag{2.24}$$

$$\sigma^2_{Q(s,a)} \leftarrow w(s,a)\left(\sigma^2_{Q(s,a)}\right) + (1 - w(s,a))\,\sigma^2_0 \tag{2.25}$$

where $w(s,a)$ is some proportion of a maximal forgetting factor w_0 determined based on the eligibility factor such that slower learning entails slower forgetting:

$$w(s,a) = 1 - z(s,a)(1 - w_0) \tag{2.26}$$

In state s, the overall uncertainty of the model-free agent's reward estimate measured as $\sigma^2_{Q(s,a^*)} + \sigma^2_r$, where a^* is the ostensible best next action $\text{argmax}_a Q(s,a)$. This is the variance of its estimated probability density function over possible values of $R(s,a^*)$.

2.7.3 Model-based learner

In contrast to a model-free agent, a model-based agent attempts to learn an explicit model of its environment. At decision time, the expected reward of each possible action is computed recursively by traversing a tree of estimated action–outcome contingencies. Like the model-free agent, the model-based agent's beliefs about action–outcome contingencies and about the reward associated with each terminal state include some uncertainty. Identically to the model-free agent, the model-based agent assumes the reward $r(s)$ associated with a terminal state s is normally distributed with unknown mean and fixed variance σ^2_r. The agent's current best estimate of $\mathbb{E}[r(s)]$ is denoted $\hat{\mu}_{r(s)}$. The model-based agent's belief in this estimate is encoded in a normal probability density function over possible estimates μ centred on the current best estimate $\hat{\mu}_{r(s)}$:

$$p(\mu) = \mathcal{N}\left(\mu \middle| \hat{\mu}_{r(s)}, \sigma^2_{\hat{\mu}_{r(s)}}\right) \tag{2.27}$$

where $\sigma^2_{\hat{\mu}_{r(s)}}$ quantifies the agent's uncertainty about $\hat{\mu}_{r(s)}$. After observing some reward r, this density is updated according to

$$\left[\hat{\mu}_{r(s)}\right]' = \hat{\mu}_{r(s)} + \frac{\sigma^2_{\hat{\mu}_{r(s)}}}{\sigma^2_r + \sigma^2_{\hat{\mu}_{r(s)}}} \left(r - \hat{\mu}_{r(s)}\right) \tag{2.28}$$

$$\left[\sigma^2_{\hat{\mu}_{r(s)}}\right]' = \left(\frac{1}{\sigma^2_{\hat{\mu}_{r(s)}}} + \frac{1}{\sigma^2_r}\right)^2 \tag{2.29}$$

These updates are identical to the model-free agent's updates to $Q(s,a)$ and $\sigma^2_{Q(s,a)}$ when a leads to a terminal state.

To learn action–outcome contingencies, the agent tracks the number of times each state–action pair (s,a) leads to a given successor state s'. This count is denoted $\theta_{s'}$ and the vector of counts for each possible successor following (s,a) can be denoted $\boldsymbol{\theta}$. Given a set of state-action-successor observations, the best estimate of the probability of a successor state when action a is taken in state s is

$$\hat{t}\left(s'|s,a\right) = \frac{\theta_{s'}}{\sum \theta} \tag{2.30}$$

where $\sum \theta$ is the total number of times (s,a) has occurred. In other words, the estimated probability of a successor s'_k following state–action pair (s,a) is simply its so-far-observed relative frequency. As with the estimate of $\mathbb{E}\left[r(s)\right]$, the agent's belief in these estimated transition probabilities contains some uncertainty, and is encoded in a Dirichlet probability density function over possible successor distributions

$$p\left(\mathbf{t}\right) = \mathrm{Dir}\left(\mathbf{t}|\boldsymbol{\theta}\right) \tag{2.31}$$

where $\mathbf{t}$ is a vector representing a possible probability distribution over successor states following (s,a). The parameters of this Dirichlet distribution are simply the counts $\theta_{s'}$ of each successor state s' following (s,a). Before making any observations, each count $\theta_{s'}$ is set to the same initial value θ_0, reflecting a prior belief that each possible successor is equally likely:

$$\boldsymbol{\theta} \leftarrow \theta_0 \mathbf{1} \tag{2.32}$$

Like the model-free agent, the model-based agent is more sensitive to recent observations than distant ones through a process of exponential forgetting. Before updating the parameters of the distribution over possible values of $\mathbb{E}[r(s)]$, these decay toward their initial values σ_0^2 and μ_0 by a forgetting rate w (equal to w_0 in the model-free learning equations):

$$\sigma^2_{\hat{\mu}_{r(s)}} \leftarrow w\sigma^2_{\hat{\mu}_{r(s)}} + (1-w)\sigma_0^2 \tag{2.33}$$

$$\hat{\mu}_{r(s)} \leftarrow w\hat{\mu}_{r(s)} + (1-w)\mu_0 \tag{2.34}$$

Similarly, following a state–action pair, the corresponding counts $\boldsymbol{\theta}$ for each possible successor decayed toward their initial values θ_0 of before the count corresponding to the observed successor was incremented:

$$\boldsymbol{\theta} \leftarrow w\boldsymbol{\theta} + (1-w)\theta_0 \tag{2.35}$$

To estimate the expected value of taking some action a in state s, the agent computes the following:

$$Q(s,a) = \sum_{s'} \max_{a'} Q(s',a')\hat{t}\,(s'|s,a) \tag{2.36}$$

where $Q(s',a')$, the estimated expected value of taking action a' in state s', is computed recursively according to the same equation. If s' is terminal, then $\max_{a'} Q(s',a') = \hat{\mu}_{r(s')}$.

The agent's estimate of the distribution over total reward value received from state s onward can be denoted $f_s(r)$, and is given as follows:

$$f_s(r) = \int_{\mathbf{t}} \left(\sum_{s'} f_{s'}(r) t(s'|s,a^*) \right) \text{Dir}(\mathbf{t}|\boldsymbol{\theta})\ d\mathbf{t} \tag{2.37}$$

where a^* is the ostensible best next action. $f_s(r)$ is a mixture distribution of sub-mixture distributions $\sum_{s'} f_{s'}(r)t(s'|s,a^*)$, where each sub-mixture is weighted by its relative probability $\text{Dir}(\mathbf{t}|\boldsymbol{\theta})$. To compute the variance of the overall mixture, we need to compute its moments. Let R_s be a random variable denoting the total reward received from state s onward. The nth moment of R_s is

$$\begin{aligned} \mathbb{E}\left[R_s^n\right] &= \int_r r^n \int_{\mathbf{t}} \left(\sum_{s'} f_{s'}(r) t(s'|s,a^*) \right) \text{Dir}(\mathbf{t}|\boldsymbol{\theta})\ d\mathbf{t}\ dr \\ &= \int_{\mathbf{t}} \int_r r^n \left(\sum_{s'} f_{s'}(r) t(s'|s,a^*) \right)\ dr\ \text{Dir}(\mathbf{t}|\boldsymbol{\theta})\ d\mathbf{t} \end{aligned} \tag{2.38}$$

i.e., a weighted average of the nth moment of each possible sub-mixture with weight vector $\mathbf{t}$. This can be written in terms of the nth moments of possible successors:

$$\begin{aligned} \mathbb{E}\left[R_s^n\right] &= \int_{\mathbf{t}} \left(\sum_{s'} t(s'|s,a^*) \int_r r^n f_{s'}(r) dr \right) \text{Dir}(\mathbf{t}|\boldsymbol{\theta})\ d\mathbf{t} \\ &= \int_{\mathbf{t}} \left(\sum_{s'} t(s'|s,a^*) \mathbb{E}[R_{s'}^n] \right) \text{Dir}(\mathbf{t}|\boldsymbol{\theta})\ d\mathbf{t} \end{aligned} \tag{2.39}$$

This is the expected value of the dot product of a Dirichlet-distributed random vector $\mathbf{t}$ with a constant vector comprised of the nth moment of $R_{s'}$ for each possible successor state s'. This is simply the dot product of the expected value of $\mathbf{t}$ with this latter vector of nth moments:

$$\begin{aligned} \mathbb{E}\left[R_s^n\right] &= \sum_{s'} \mathbb{E}\left[R_{s'}^n\right] \int_{\mathbf{t}} t(s'|s,a^*) \text{Dir}(\mathbf{t}|\boldsymbol{\theta})\ d\mathbf{t} \\ &= \sum_{s'} \mathbb{E}[R_{s'}^n] \mathbb{E}\left[t(s'|s,a^*)\right] \\ &= \sum_{s'} \mathbb{E}[R_{s'}^n] \hat{t}(s'|s,a^*) \end{aligned} \tag{2.40}$$

where $\hat{t}(s'|s, a^*)$ is defined as above in terms of observed counts and $\mathbb{E}[R^n_{s'}]$ are computed recursively. In the "base case" where s' is terminal,

$$\mathbb{E}[R^2_{s'}] = \mu^2_{r(s')} + \sigma^2_{\mu_{r(s')}} + \sigma^2_r \tag{2.41}$$

The variance of R_s, and therefore the uncertainty of the model-based agent's reward predictions, can be computed using the first two moments as $\mathrm{Var}[R_s] = \mathbb{E}[R^2_s] - \mathbb{E}[R_s]^2$.[3] These two moments, in turn, are computed by value iteration (Sutton and Barto, 2018).

Bibliography

Adams, C. D. (1982). Variations in the sensitivity of instrumental responding to reinforcer devaluation. *The Quarterly Journal of Experimental Psychology Section B*, 34(2b):77–98.

Amlung, M., Vedelago, L., Acker, J., Balodis, I., and MacKillop, J. (2017). Steep delay discounting and addictive behavior: A meta-analysis of continuous associations. *Addiction*, 112(1):51–62.

Atance, C. M. and O'Neill, D. K. (2001). Episodic future thinking. *Trends in cognitive sciences*, 5(12):533–539.

Audrain-McGovern, J., Rodriguez, D., Epstein, L. H., Cuevas, J., Rodgers, K., and Wileyto, E. P. (2009). Does delay discounting play an etiological role in smoking or is it a consequence of smoking? *Drug and alcohol dependence*, 103(3):99–106.

Bickel, W. K., Athamneh, L. N., Basso, J. C., Mellis, A. M., DeHart, W. B., Craft, W. H., and Pope, D. (2019). Excessive discounting of delayed reinforcers as a trans-disease process: Update on the state of the science. *Current Opinion in Psychology*, 30:59–64.

Bickel, W. K., Koffarnus, M. N., Moody, L., and Wilson, A. G. (2014). The behavioral-and neuro-economic process of temporal discounting: A candidate behavioral marker of addiction. *Neuropharmacology*, 76:518–527.

Boileau, I., Payer, D., Chugani, B., Lobo, D., Houle, S., Wilson, A., Warsh, J., Kish, S., and Zack, M. (2014). In vivo evidence for greater amphetamine-induced dopamine release in pathological gambling: a positron emission tomography study with [11c]-(+)-phno. *Molecular psychiatry*, 19(12):1305–1313.

Bruner, N. R. and Johnson, M. W. (2014). Demand curves for hypothetical cocaine in cocaine-dependent individuals. *Psychopharmacology*, 231:889–897.

[3]It is perhaps surprising that $\mathrm{Var}[R_s]$ depends only on the relative values of $\theta_{s'}$, i.e., that multiplying $\boldsymbol{\theta}$ by some scalar $k > 1$ does not decrease $\mathrm{Var}[R_s]$. This is because the variance is a linear function of the mixture weights; the variances of the beta-binomial and Dirichlet-multinomial distribution are similarly dependent only on the relative values of their concentration parameters when the number of trials is 1.

Coelho, L. P. (2013). Integral of the product of two gaussians - elm-pages.

Cuzen, N. L. and Stein, D. J. (2014). Behavioral addiction: The nexus of impulsivity and compulsivity. In *Behavioral addictions*, pages 19–34. Elsevier.

Daugherty, J. R. and Brase, G. L. (2010). Taking time to be healthy: Predicting health behaviors with delay discounting and time perspective. *Personality and Individual differences*, 48(2):202–207.

Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., and Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69(6):1204–1215.

Daw, N. D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*, 8(12):1704–1711.

Dearden, R., Friedman, N., and Andre, D. (1999). Model based bayesian exploration. In *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, pages 150–159.

Dearden, R., Friedman, N., and Russell, S. (1998). Bayesian q-learning. *Aaai/iaai*, 1998:761–768.

Dickinson, A. (1985). Actions and habits: the development of behavioural autonomy. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 308(1135):67–78.

Dolan, R. J. and Dayan, P. (2013). Goals and habits in the brain. *Neuron*, 80(2):312–325.

D'Argembeau, A. (2013). On the role of the ventromedial prefrontal cortex in self-processing: the valuation hypothesis. *Frontiers in human neuroscience*, 7:372.

Eppinger, B., Walter, M., Heekeren, H. R., and Li, S.-C. (2013). Of goals and habits: age-related and individual differences in goal-directed decision-making. *Frontiers in neuroscience*, 7:253.

Epstein, D. H. (2020). Let's agree to agree: a comment on hogarth (2020), with a plea for not-so-competing theories of addiction. *Neuropsychopharmacology*, 45(5):715–716.

Everitt, B. J. and Robbins, T. W. (2005). Neural systems of reinforcement for drug addiction: from actions to habits to compulsion. *Nature neuroscience*, 8(11):1481–1489.

Everitt, B. J. and Robbins, T. W. (2016). Drug addiction: updating actions to habits to compulsions ten years on. *Annual review of psychology*, 67:23–50.

Forster, S. E., Steinhauer, S. R., Ortiz, A., and Forman, S. D. (2021). Evaluating effects of episodic future thinking on valuation of delayed reward in cocaine use disorder: a pilot study. *The American journal of drug and alcohol abuse*, 47(2):199–208.

Friston, K. (2012). Policies and priors. *Computational neuroscience of drug addiction*, pages 237–283.

Friston, K. J., Shiner, T., FitzGerald, T., Galea, J. M., Adams, R., Brown, H., Dolan, R. J., Moran, R., Stephan, K. E., and Bestmann, S. (2012). Dopamine, affordance and active inference. *PLoS computational biology*, 8(1):e1002327.

García-Pérez, Á., Aonso-Diego, G., Weidberg, S., and Secades-Villa, R. (2022). Effects of episodic future thinking on reinforcement pathology during smoking cessation treatment among individuals with substance use disorders. *Psychopharmacology*, 239(2):631–642.

Gershman, S. J. and Bhui, R. (2020). Rationally inattentive intertemporal choice. *Nature communications*, 11(1):3365.

Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A., and Daw, N. D. (2016a). Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. *elife*, 5:e11305.

Gillan, C. M., Robbins, T. W., Sahakian, B. J., van den Heuvel, O. A., and van Wingen, G. (2016b). The role of habit in compulsivity. *European Neuropsychopharmacology*, 26(5):828–840.

Guillem, K. and Ahmed, S. H. (2018). Preference for cocaine is represented in the orbitofrontal cortex by an increased proportion of cocaine use-coding neurons. *Cerebral cortex*, 28(3):819–832.

Guillem, K., Brenot, V., Durand, A., and Ahmed, S. H. (2018). Neuronal representation of individual heroin choices in the orbitofrontal cortex. *Addiction biology*, 23(3):880–888.

Haruno, M. and Kawato, M. (2006). Heterarchical reinforcement-learning model for integration of multiple cortico-striatal loops: fmri examination in stimulus-action-reward association learning. *Neural networks*, 19(8):1242–1254.

Hassabis, D., Kumaran, D., Vann, S. D., and Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *Proceedings of the National Academy of Sciences*, 104(5):1726–1731.

Hogarth, L. (2020). Addiction is driven by excessive goal-directed drug choice under negative affect: translational critique of habit and compulsion theory. *Neuropsychopharmacology*, 45(5):720–735.

Huang, Y., Yaple, Z. A., and Yu, R. (2020). Goal-oriented and habitual decisions: Neural signatures of model-based and model-free learning. *NeuroImage*, 215:116834.

Hunter, L. E., Bornstein, A. M., and Hartley, C. A. (2018). A common deliberative process underlies model-based planning and patient intertemporal choice. *bioRxiv*, page 499707.

Hutcheson, D., Everitt, B., Robbins, T., and Dickinson, A. (2001). The role of withdrawal in heroin addiction: enhances reward or promotes avoidance? *Nature neuroscience*, 4(9):943–947.

Jaynes, E. T. (1957a). Information theory and statistical mechanics. *Physical review*, 106(4):620.

Jaynes, E. T. (1957b). Information theory and statistical mechanics. ii. *Physical review*, 108(2):171.

Johnson, A. and Redish, A. D. (2007). Neural ensembles in ca3 transiently encode paths forward of the animal at a decision point. *Journal of Neuroscience*, 27(45):12176–12189.

Keramati, M. and Gutkin, B. (2013). Imbalanced decision hierarchy in addicts emerging from drug-hijacked dopamine spiraling circuit. *PloS one*, 8(4):e61489.

Kinley, I., Amlung, M., and Becker, S. (2022). Pathologies of precision: A bayesian account of goals, habits, and episodic foresight in addiction. *Brain and Cognition*, 158:105843.

Krieckhaus, E. and Wolf, G. (1968). Acquisition of sodium by rats: interaction of innate mechanisms and latent learning. *Journal of comparative and Physiological Psychology*, 65(2):197.

Lee, R. S., Hoppenbrouwers, S., and Franken, I. (2019). A systematic meta-review of impulsivity and compulsivity in addictive behaviors. *Neuropsychology review*, 29:14–26.

Lee, S. W., Shimojo, S., and O'Doherty, J. P. (2014). Neural computations underlying arbitration between model-based and model-free learning. *Neuron*, 81(3):687–699.

MacKillop, J., Amlung, M. T., Few, L. R., Ray, L. A., Sweet, L. H., and Munafò, M. R. (2011). Delayed reward discounting and addictive behavior: a meta-analysis. *Psychopharmacology*, 216:305–321.

Mantsch, J. R., Baker, D. A., Funk, D., Lê, A. D., and Shaham, Y. (2016). Stress-induced reinstatement of drug seeking: 20 years of progress. *Neuropsychopharmacology*, 41(1):335–356.

Mathar, D., Erfanian Abdoust, M., Marrenbach, T., Tuzsus, D., and Peters, J. (2022). The catecholamine precursor tyrosine reduces autonomic arousal and decreases decision thresholds in reinforcement learning and temporal discounting. *PLOS Computational Biology*, 18(12):e1010785.

Matochik, J. A., London, E. D., Eldreth, D. A., Cadet, J.-L., and Bolla, K. I. (2003). Frontal cortical tissue composition in abstinent cocaine abusers: a magnetic resonance imaging study. *Neuroimage*, 19(3):1095–1102.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. *Quantitative analyses of behavior*, 5:55–73.

Mollick, J. A. and Kober, H. (2020). Computational models of drug use and addiction: A review. *Journal of abnormal psychology*, 129(6):544.

Moore, A. W. and Atkeson, C. G. (1993). Prioritized sweeping: Reinforcement learning with less data and less time. *Machine learning*, 13:103–130.

Naik, A., Shariff, R., Yasui, N., Yao, H., and Sutton, R. S. (2019). Discounted reinforcement learning is not an optimization problem. *arXiv preprint arXiv:1910.02140*.

Ognibene, D., Fiore, V. G., and Gu, X. (2019). Addiction beyond pharmacological effects: The role of environment complexity and bounded rationality. *Neural Networks*, 116:269–278.

Patel, H. and Amlung, M. (2020). Acute and extended exposure to episodic future thinking in a treatment seeking addiction sample: A pilot study. *Journal of Substance Abuse Treatment*, 116:108046.

Pierce, R. C. and Kumaresan, V. (2006). The mesolimbic dopamine system: the final common pathway for the reinforcing effect of drugs of abuse? *Neuroscience & biobehavioral reviews*, 30(2):215–238.

Poletti, M., Logi, C., Lucetti, C., Del Dotto, P., Baldacci, F., Vergallo, A., Ulivi, M., Del Sarto, S., Rossi, G., Ceravolo, R., et al. (2013). A single-center, cross-sectional prevalence study of impulse control disorders in parkinson disease: association with dopaminergic drugs. *Journal of clinical psychopharmacology*, 33(5):691–694.

Radenbach, C., Reiter, A. M., Engert, V., Sjoerds, Z., Villringer, A., Heinze, H.-J., Deserno, L., and Schlagenhauf, F. (2015). The interaction of acute and chronic stress impairs model-based behavioral control. *Psychoneuroendocrinology*, 53:268–280.

Redish, A. D. (2004). Addiction as a computational process gone awry. *Science*, 306(5703):1944–1947.

Redish, A. D., Jensen, S., and Johnson, A. (2008). Addiction as vulnerabilities in the decision process. *Behavioral and brain sciences*, 31(4):461–487.

Rösch, S. A., Stramaccia, D. F., and Benoit, R. G. (2022). Promoting farsighted decisions via episodic future thinking: A meta-analysis. *Journal of Experimental Psychology: General*, 151(7):1606.

Rozeboom, W. W. (1958). " what is learned?"—an empirical enigma. *Psychological Review*, 65(1):22.

Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of neurophysiology*, 80(1):1–27.

Schultz, W., Apicella, P., and Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *Journal of neuroscience*, 13(3):900–913.

Schwartenbeck, P., FitzGerald, T. H., Mathys, C., Dolan, R., Wurst, F., Kronbichler, M., and Friston, K. (2015). Optimal inference with suboptimal models: addiction and active bayesian inference. *Medical hypotheses*, 84(2):109–117.

Shenhav, A., Rand, D. G., and Greene, J. D. (2017). The relationship between intertemporal choice and following the path of least resistance across choices, preferences, and beliefs. *Judgment and Decision making*, 12(1):1–18.

Sinclair, H., Lochner, C., and Stein, D. J. (2016). Behavioural addiction: a useful construct? *Current Behavioral Neuroscience Reports*, 3:43–48.

Snider, S. E., LaConte, S. M., and Bickel, W. K. (2016). Episodic future thinking: Expansion of the temporal window in individuals with alcohol dependence. *Alcoholism: clinical and experimental research*, 40(7):1558–1566.

Solway, A., Lohrenz, T., and Montague, P. R. (2017). Simulating future value in intertemporal choice. *Scientific Reports*, 7(1):43119.

Sozou, P. D. (1998). On hyperbolic discounting and uncertain hazard rates. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 265(1409):2015–2020.

Story, G. W., Vlaev, I., Seymour, B., Darzi, A., and Dolan, R. J. (2014). Does temporal discounting explain unhealthy behavior? a systematic review and reinforcement learning perspective. *Frontiers in behavioral neuroscience*, 8:76.

Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.

Szpunar, K. K. and Schacter, D. L. (2013). Get real: effects of repeated simulation and emotion on the perceived plausibility of future experiences. *Journal of Experimental Psychology: General*, 142(2):323.

van Rooij, I. and Blokpoel, M. (2020). Formalizing verbal theories: A tutorial by dialogue (preprint). psyarxiv.

Vandaele, Y. and Ahmed, S. (2021). Habit, choice, and addiction. *Neuropsychopharmacology*, 46(4):689–698.

Vikbladh, O. M., Meager, M. R., King, J., Blackmon, K., Devinsky, O., Shohamy, D., Burgess, N., and Daw, N. D. (2019). Hippocampal contributions to model-based planning and spatial memory. *Neuron*, 102(3):683–693.

Voon, V., Derbyshire, K., Rück, C., Irvine, M. A., Worbe, Y., Enander, J., Schreiber, L. R., Gillan, C., Fineberg, N. A., Sahakian, B. J., et al. (2015). Disorders of compulsivity: a common bias towards learning habits. *Molecular psychiatry*, 20(3):345–352.

Wagner, B., Mathar, D., and Peters, J. (2022). Gambling environment exposure increases temporal discounting but improves model-based control in regular slot-machine gamblers. *Computational Psychiatry*, 6(1).

Wang, X., Li, B., Zhou, X., Liao, Y., Tang, J., Liu, T., Hu, D., and Hao, W. (2012). Changes in brain gray matter in abstinent heroin addicts. *Drug and alcohol dependence*, 126(3):304–308.

Watkins, C. J. and Dayan, P. (1992). Q-learning. *Machine learning*, 8:279–292.

Yaari, M. E. (1965). Uncertain lifetime, life insurance, and the theory of the consumer. *The Review of Economic Studies*, 32(2):137–150.

Chapter 3

Visual perspective as a two-dimensional construct in episodic future thought

Isaac Kinley, Morgan Porteous, Yarden Levy, & Suzanna Becker

Introductory note

In the previous chapter, we used computational modelling to provide a framework for understanding the role of episodic future thinking in intertemporal choice. Here we take an empirical approach, searching for covariates of the effect of future thinking cues presented during a delay discounting task. A covariate of particular interest is visual perspective, a main focus of this chapter.

An interesting quirk of human memory is that we do not always recall events from the perspective of our own eyes. Sometimes, especially for distant memories, we tend to adopt a third-person perspective and see ourselves "from the outside" in a remembered scene. The same is true of imagined future events: when we imagine the future, we might do so from the perspective of our own eyes or from that of an observer. Visual perspective in episodic future thinking has been shown to modulate the emotional impact of imagined future events, and is therefore a plausible candidate for a covariate of the effect of episodic future thinking on decision making.

Most research on visual perspective in memory and future thinking assumes that first- and third-person perspectives are mutually exclusive, but this is arguably too restrictive. For example, we might switch between first- and third-person perspectives when remembering or imagining the same event. Similarly, we might adopt a perspective that is not first-person, but from which we do not see ourselves "from the outside", which is an important element of a third-person perspective.

To measure the effect of episodic future thinking on delay discounting, we used an intertemporal choice task with future thinking cues. We collected questionnaire measures of a range of variables (e.g., vividness, emotional valence) potentially relevant to this effect. Importantly, our questionnaire separately measured first- and third-person visual perspective. This allowed us to test whether visual perspective, measured as a two-dimensional rather than a one-dimensional construct, is related to other variables measured by our questionnaire and to the effect of episodic future thinking on delay discounting.

The main author (I.K.) designed the study, collected approximately half the data (with M.P.), performed the analyses, and primarily wrote the paper. It appears in Consciousness and Cognition as Kinley, I., Porteous, M., Levy, Y., & Becker, S. (2021). Visual perspective as a two-dimensional construct in episodic future thought. *Consciousness and Cognition*, *93*, 103148.

Abstract

Visual perspective (first-person vs. third-person) is a salient characteristic of memory and mental imagery with important cognitive and behavioural consequences. Most work on visual perspective treats it as a unidimensional construct. However, third-person perspective can have opposite effects on emotion and motivation, sometimes intensifying these and other times acting as a distancing mechanism, as in PTSD. For this reason among others, we propose that visual perspective in memory and mental imagery is best understood as varying along two dimensions: first, the degree to which first-person perspective predominates in the episodic imagery, and second, the degree to which the self is visually salient from a third-person perspective. We show that, in episodic future thinking, these are anticorrelated but non-redundant. These results further our basic understanding of the potent but divergent effects visual perspective has on emotion and motivation, both in everyday life and in psychiatric conditions.

3.1 Introduction

One of the earliest works discussing visual perspective in memory is Freud's 1899 essay on "screen memories"—vivid childhood memories of seemingly anodyne situations which, he argued, were replacements for or sanitized versions of distressing but autobiographically significant events (Freud, 1962). The essay draws on Victor and Catherine Henri's survey of early childhood experiences (Nicolas et al., 2013) to argue that memories retrieved from the perspective of an outside observer ("third-person" memories), as early memories often are, cannot be faithful re-renderings of the original experiences, which must have been originally experienced from a first-person perspective. Freud argued that these memories must instead have been "worked over."

The topic seems to have lain dormant for nearly a century until the seminal work of Nigro and Neisser (1983), which not only re-established the existence of third-person memories, but explored memory characteristics predictive of first- or third-person perspective. Among these were self-awareness (predictive of third-person perspective) and a focus on emotions rather than objective details (predictive of first-person perspective). Also, older memories were generally more likely to be recalled from a third-person perspective.

3.1.1 Visual perspective in memory

Visual perspective is not an immutable characteristic of memory. Robinson and Swanson (1993) note that people report switching between first- and third-person points of view within the same memory. This switching is easier for more recent and vivid memories and, while the switch from first- to third-person is associated with reduced affect, the reverse switch is not associated with increased affect. They suggest that when the sensory-affective details of a memory are lost, it tends to be retrieved from a third-person perspective; only the cognitive component of the memory trace remains, and forcing a switch to first-person perspective cannot bring forth sensory-affective detail that is no longer accessible. However, forcing a switch away from a naturally-occurring first-person perspective involves suppressing accessible detail, resulting in reduced affect (Robinson and Swanson, 1993).

Subsequent findings have aligned well with this model: first-person perspective in natural recall is associated with greater emotional intensity (McIsaac and Eich, 2002; D'Argembeau et al., 2003; McIsaac and Eich, 2004; Talarico et al., 2004), detail (McIsaac and Eich, 2002, 2004), vividness (Sutin and Robins, 2010), and sense of reliving (Berntsen and Rubin, 2006). There are also asymmetries in perspective switching (decreases after the first-to-third-person switch but no increase after the reverse switch) for sense of reliving (Berntsen and Rubin, 2006), emotional intensity (Vella and Moulds, 2014; Gu and Tse, 2016), vividness (Williams and Moulds, 2008), and episodic detail (Akhtar et al., 2017). Furthermore, retrieval from a third-person perspective alters memories, reducing vividness (Butler et al., 2016), accuracy (Marcotti and St. Jacques, 2018), emotional intensity (Sekiguchi and Nonaka, 2014) and sensory-affective detail (Bagri and Jones, 2009) at subsequent recall. First-person perspective appears, then, to be closely related to the episodicity of memory. Indeed, the first-person–third-person distinction is closely related (though not identical; Sutton, 2010) to the remember–know distinction (Crawley and French, 2005) and the proportion of memories recalled from a first-person perspective follows a reminiscence bump for events in young adulthood (Piolino et al., 2006).

Several recent studies have found that visual perspective in memory recall is a reliable individual-differences variable (Siedlecki, 2015; Verhaeghen et al., 2018), and it may be one with important connections to psychological health. Third-person perspectives are argued to be constructed from "contextual" memory representations that are dissociable from the low-level sensory representations underlying vivid intrusions (Brewin et al., 2010). As such, individuals with PTSD often deliberately adopt a third-person perspective when recalling trauma memories to avoid reliving them (McIsaac and Eich, 2004). Similarly, more avoidant survivors of trauma are more likely to adopt third-person perspectives on their trauma memories (Kenny and Bryant, 2007). Third-person perspective memories are also more common among depressed individuals (Kuyken and Howell, 2006, but cf. McFadden and Siedlecki, 2020), possibly for positive events in particular (Lemogne et al., 2006; Nelis et al., 2013) or for individuals with a tendency toward cognitive avoidance of intrusions (Williams and Moulds, 2007; Kuyken and Moulds, 2009). Finally, trait dissociation is linked to increased third-person perspective in memory recall (Williams and Moulds, 2007; Sutin and Robins, 2010; Radvansky and Svob, 2019). The fact that third-person perspective can dampen affect apparently causes it to function as a psychological distancing mechanism across psychiatric conditions.

3.1.2 Visual perspective in episodic future thinking

The personal future can also be imagined from a first- or third-person perspective through a process known as episodic future thinking (EFT; Atance and O'Neill, 2001). Just as visual perspective in memory recall has important implications for affect, visual perspective in EFT has important implications for motivation and decision making. First-person perspective tends to emphasize sensory-affective information while third-person perspective tends to highlight more abstract information that contextualizes events, and this difference has behavioural consequences (Libby and Eibach, 2011). For example, 90% of registered voters who visualized voting in the 2004 American presidential election from a third-person perspective actually followed through with their voting intentions, compared to 72% who visualized voting from a first-person perspective (Libby et al., 2007). Similarly, visualizing an academic achievement from a third-person perspective increases motivation toward it more than visualization from a first-person perspective, an effect mediated by high construal level, which casts academic achievement within the larger context of more abstract goals (Vasquez and Buehler, 2007; Trope and Liberman, 2010). Finally, visualizing the future self from a third-person perspective increases hypothetical retirement saving, an effect mediated by the visual salience of the future self (Macrae et al., 2017).

3.1.3 Explaining divergent effects of third-person perspective

There is an apparent contradiction in the divergent effects of visual perspectives on motivation and emotion: why should third-person perspective be associated with increased motivation if it is associated with dampened affect? To resolve this, McCarroll (2019) argues that third-person perspective does not always function as an emotional distancing mechanism. Instead, it can sometimes enhance self-conscious emotions such as pride and thereby motivate future-oriented decisions. In the memory literature, Sutin and Robins (2008) have introduced a model in which third-person perspective can either heighten or dampen affect: for positive memories that are consistent with the current self-concept, third-person perspective will heighten affect by bringing the remembered self into the visual foreground of a memory, thereby increasing self-focus (the "salient self" effect). Meanwhile, for negative memories that are inconsistent with the current self-concept, third-person perspective will dampen affect by creating psychological distance between the remembering and remembered self (the "dispassionate observer"; see Rice, 2010, for a discussion of this model's limitations).

Extending this model to EFT, third-person perspective should increase or decrease motivation toward future goals depending on whether the "salient self" or the "dispassionate observer" effect prevails, which in turn should depend on whether the goal in question is consistent with the current self-concept. In line with this, pursuit of health goals is hindered by third-person visualization if they are peripheral to the self-concept (Stornelli et al., 2020).

3.1.4 Against a unidimensional model of visual perspective

Much of the memory and EFT literature on visual perspective treats it as a unidimensional construct, with first- and third-person perspectives at opposite poles (e.g., Berntsen and Rubin, 2006). However, the fact that third-person perspective can have divergent effects on emotion/motivation suggests that visual perspective is not unidimensional. Here we make a case against a unidimensional model of visual perspective (see also Rice and Rubin 2009).

The definition of a third-person perspective usually contains 2 criteria: a point of view different from first-person and the visibility of the remembered or imagined self (Nigro and Neisser, 1983). These two criteria appear to correspond to the two divergent effects of third-person perspective posited by Sutin and Robins (2008): a shift in perspective away from first-person should amplify the "dispassionate observer" effect, while the vivid visibility of the self should amplify the "salient self" effect.

Moreover, these two criteria seem logically separable: the fact that a point of view is not first-person does not necessarily mean that the remembered or imagined self will be visible (Sutton, 2010; McCarroll, 2019). This proposal derives from the concept of "acentred" memories, in which the point of view is not the same as the original, but neither is the self visible in the memory scene, having been "edited out" (Wollheim, 1984). However, as far as we are aware, there are presently no empirical data on the prevalence of this phenomenon.

Additionally, there is an infinite number of spatial locations and angles from which a third-person perspective can be constructed (Callow and Roberts, 2010; Morris and Spittle, 2012), and the particular third-person vantage point from which an observer memory is retrieved relates to that memory's content (e.g., when remembering giving a presentation, one is more likely to recall the scene from in front of oneself; Rice and Rubin, 2011). Clearly, then, third-person perspective is not just one end of a unidimensional construct but is instead a category of perspectives within which meaningful variation exists.

Finally, it is possible to adopt multiple perspectives when retrieving a single memory. For example, Berntsen and Rubin (2006) provide the following description of a memory in which the perspective switches:

> I see myself dancing at a party at the university. I remember my clothes and my legs (the way they moved). Suddenly, I am "inside my own body" looking out. A guy I know a little walks by me and says as he passes: "You look good today".

Thus a single memory or imagined future episode can be described as involving both first- and third-person perspectives—i.e., as simultaneously occupying both opposite poles of visual perspective understood as a unidimensional construct. Switching between perspectives is also observed in sports imagery, where it has been shown to benefit performance (Epstein, 1980; Gordon et al., 1994; Smith et al., 1998). It has even been suggested that multiple perspectives can coexist simultaneously within a mental image in a way that is not possible in perception (Sutton, 2012). This is based in part on the observation that spontaneous verbal descriptions of space often employ route (first-person) and survey (third-person) descriptions within a single clause (Taylor and Tversky, 1992). Similarly, Sartre's view that mental images reflect rather than precede knowledge implies that multiple visual perspectives can be adopted simultaneously (Sartre, 1972, as cited in McCarroll, 2018, pp. 144–145).

3.1.5 Visual perspective as a two-dimensional construct

In light of the fact that multiple visual perspectives can co-occur within a single memory retrieval, Rice and Rubin (2009) argue that first- and third-person perspectives are independent dimensions that characterize memory, and show that the two correlate differentially with ratings of vividness. However, as discussed above, meaningful variation exists between third-person perspectives, which may not be fully captured when describing a memory or imagined future episode as involving simply more or less third-person perspective. Here we offer an alternative two-dimensional conceptualization of visual perspective.

We argue that visual perspective, in both memory and episodic future thinking, is best understood as comprising (1) the predominance of first-person perspective (henceforth "first-person predominance") and (2) the visibility of the imagined self when a third-person perspective is adopted (henceforth "third-person self-visibility"). As noted above, people can adopt multiple visual perspectives on the same event, meaning that first-person perspective can predominate to varying degrees. Furthermore, from a third-person perspective, the imagined self may be more or less visually salient.

These two dimensions articulate the space of possibilities for visual perspective: when first-person perspective predominates and third-person self-visibility is low, the overall perspective is best described as first-person. When these are both reversed, the result is a third-person perspective as conceived by Nigro and Neisser (1983). When third-person perspective predominates but self-visibility is low, the result is acentred imagery. Finally, when both third-person self-visibility and the predominance of first-person perspective are moderate to high, this describes mental imagery in which the visual perspective switches between first- and third-person.

3.1.6 The present work

The present work aimed to evaluate the utility of measuring visual perspective in terms of the two dimensions proposed above. In order to measure effects on motivation and decision making, we used an episodic future thinking paradigm. We also aimed to identify perceptual and affective correlates of the two proposed dimensions.

First (prediction 1), we expected that the two dimensions would be separable, allowing the measurement of both acentred imagery (low first-person predominance and low third-person self-visibility) and perspective switching (high first-person predominance and high third-person self-visibility). Second (prediction 2), in line with previous research, we expected first-person predominance to be positively correlated with subjective ratings of vividness and negatively correlated with temporal distance (i.e. lower first-person predominance with increasing temporal distance). Third (prediction 3), based on the model of Sutin and Robins (2008), we expected emotional intensity to be positively correlated with both first-person predominance and third-person self-visibility: first-person predominance should reduce the "dispassionate observer" effect, while third-person self-visibility should increase the "salient self" effect.

We also aimed to measure trait predictors of each of these two dimensions. One likely candidate is trait dissociation, which manifests as feelings of detachment from one's experiences and surroundings (Lyssenko et al., 2018). As noted earlier, trait dissociation is correlated with increased rates of third-person perspective in memory recall (Williams and Moulds, 2007; Sutin and Robins, 2010; Radvansky and Svob, 2019), an effect interpreted in terms of emotional distancing (i.e., the "dispassionate observer" effect). Extending this idea to EFT, we predicted that trait dissociation would be associated with low ratings of both first-person predominance and third-person self-visibility (prediction 4).

Another relevant variable is trait imagery vividness, which we expected to be positively associated with first-person predominance. Mental imagery is generated by projecting representations in long-term memory into imagery areas (Byrne et al., 2007; Pearson, 2019). This activates modality-specific sensory cortices in a top-down manner, and the strength of this activation is correlated with subjective vividness ratings (Belardinelli et al., 2009). The ability to rapidly construct vivid mental images differs between individuals (Isaac and Marks, 1994), and these differences can be understood in terms of the information-theoretic concept of channel capacity: in high trait imagery individuals, there is greater information flow from long-term memory to imagery areas (Hishitani et al., 2011). We therefore expected that individuals with greater trait imagery vividness would construct more perceptually detailed mental images of the future and, given the inherently egocentric/body-centred nature of perception, would report more first-person images of the future (prediction 5).

Finally, we explored the role of visual perspective in EFT-driven reductions of impulsivity. EFT shifts preferences away from smaller immediate rewards and toward larger delayed rewards (Peters and Büchel, 2010), and is thought to do so by simulating the value of future rewards in the present (Benoit et al., 2011).

If this is the case, first-person perspective in EFT should increase the vividness of these simulations and thereby amplify the effect on impulsivity. However, as noted earlier, visualizing the far future from a third-person perspective increases hypothetical saving behaviour (Macrae et al., 2017). Thus we expected (prediction 6) that the effect of EFT on impulsivity is amplified by both first-person predominance (through increased vividness) and third-person self-visibility (through the "salient self" effect).

To summarize, the present work aimed to examine visual perspective in episodic future thought as a two-dimensional construct, comprising first-person predominance and third-person self-visibility. We expected that first-person predominance would be associated with increased state and trait vividness, while third-person self-visibility would be associated with positive affect and negatively correlated with trait dissociation. Finally, we expected that both first-person predominance and third-person self-visibility would be independently correlated with the effect of EFT on impulsivity.

3.2 Materials and methods

McMaster undergraduate students were recruited to complete an episodic future thinking task, a delay discounting task, and several questionnaires. Participants (N = 92, 77 women, 1 other/unspecified, ages 18 – 24, median = 18) completed the experiment for introductory psychology course credit, giving written consent as approved by the local research ethics board. The experiment was programmed using jsPsych (De Leeuw, 2015) and hosted on a web server that was accessed from a browser on a laboratory computer. Participants who were deemed to be "off task" (those who wrote nothing for at least one event in the future thinking writing task explained below) were excluded from further analysis, leaving a final sample size of 84 (69 women, 1 other/unspecified, ages 18 – 24, median = 18). The median time to complete the experiment was 26.58 minutes.

3.2.1 Questionnaires

In order to assess individual-differences factors potentially related to visual perspective, participants completed computerized versions of the Plymouth Sensory Imagery Questionnaire (Psi-Q; Andrade et al., 2014) and the Dissociative Experiences Scale II (DES; Carlson and Putnam, 1993). The Psi-Q measures imagery vividness across sensory modalities including vision, sound, smell, taste, touch, bodily sensation, and emotional feeling. The pencil-and-paper DES asks about the frequency with which various experiences occur in the respondent's daily life, with possible responses ranging from 0% of the time to 100% in intervals of 10%. Our computerized version of the DES used a digital visual analog scale which allowed participants to specify responses in intervals of 1%. Responses were made by moving a slider that started in the middle of the range for each question.

3.2.2 Future thinking writing task

To generate mental images of potential future events in their lives, participants completed an episodic future thinking writing task. They were given a set of criteria for the future events they would be imagining (must be actually planned or at least realistic; must be confined to a specific place and time; must not last longer than a day; must be distinct and not have happened yet). Participants then completed a future thinking writing task using the Galton-Crovitz cuing paradigm (Galton, 1879; Crovitz and Schiffman, 1974). Cue words were drawn from the Clark and Paivio (2004) norms according to the following criteria: imagery ratings above 5.5, familiarity ratings above 5.56, and log[Lorge-Thorndike frequency + 1] above 1.46. In total, of the 2311 words in the norms, 204 met our criteria (imagery mean = 6.28, 5.53 – 6.87; familiarity mean = 1.88, 1.48 – 2.00; log[Lorge-Thorndike frequency + 1] mean = 6.29, 5.57 – 6.92. In response to cue words, participants first wrote an event title (for example, "Family Visit") which would later be displayed during the delay discounting task described below, and then elaborated on the event in a free text field. Participants were asked to construct four future events, at delays of 1 week, 1 month, 6 months, and 12 months.

3.2.3 Imagery characteristics questionnaire

To query visual perspective as well as other potentially relevant phenomenological variables, we developed a brief questionnaire probing participants' mental imagery in more detail. Prior to providing their responses, participants were given the following explanation of visual perspective:

> Some of these questions ask about visual perspective. When we imagine events, we can see them from different points of view in our mind's eye. If we see the scene from the point of view of our own eyes, this is called a "first-person" perspective. If we see it from any other point of view, this is called a "third-person" perspective. Sometimes we switch back and forth between the two.

This explanation deliberately omits the self-visibility requirement from the definition of third-person perspective. Participants answered up to 2 questions pertaining to visual perspective. The first queried first-person predominance:

> What percentage of the time did you see the scene from a first-person perspective?

The digital visual analog scale had anchors at 0%, 25%, 50%, 75%, and 100%. Participants who responded less than 100% then answered the following question querying third-person self-visibility:

> When you saw the scene from a third-person perspective, did you see yourself in it?

This scale had anchors “Not at all”, “Somewhat”, and “Very clearly” at the far left, middle, and far right points, respectively. The response to the above question for participants reporting 100% first-person predominance was coded as missing.

Prior to making any of these responses, participants were given the following instructions:

> Please do not try to change your imagination of the events based on the questions. Instead, answer them based on the mental images you already had.

These instructions were included based on the possibility that participants would, when answering the third-person self-visibility question, unintentionally bring to mind their own appearance while searching for themselves in their third-person imagery and, as a result, answer affirmatively. The goal was thus to reduce the likelihood that the question would act as a suggestion.

The rest of the items queried vividness (“Was your mental image of the event faint or vivid?”, with anchors “Faint” and “Vivid”), location familiarity (“Was the location of the event familiar to you?”, with anchors “Totally unfamiliar” and “Totally familiar”), emotional valence (“Was the general emotional tone of the event positive or negative?”, with anchors “Negative” and “Positive”), and emotional intensity (“Were the emotions associated with the event intense?”, with anchors “Not intense” and “Very intense”) (D’Argembeau and Van der Linden, 2006; Johnson et al., 1988). Participants provided this series of ratings (henceforth, “event-wise ratings”) for each imagined event.

3.2.4 Delay discounting task

To measure impulsivity and its modulation by episodic future thinking, participants completed a computerized delay discounting task, making a series of hypothetical choices between smaller immediate rewards and larger delayed rewards (Kirby et al., 1999).

From trial to trial, the value of the smaller immediate reward changed according to the adjusting amount procedure such that, when the immediate reward was chosen, its value was adjusted downward on the next trial and vice versa (Koffarnus and Bickel, 2014). The reward values began at $400 and $800 and the immediate reward was first adjusted (upward or downward) by $200. This amount decreased by half after each adjustment over the course of 5 trials so that the two rewards converged on a pair between which the participant was indifferent. From this pair, the “indifference point” was computed by taking the smaller reward as a fraction of the larger.

For example, after choosing the larger delayed reward for 4 successive trials, a participant would be presented with a choice between $775 available immediately (adjusted upward from $750) and $800 available at some delay. If they chose the immediate reward on this final trial, their indifference point would be computed by adjusting $775 downward by half the previous amount, ($25/2), and dividing the result by $800, yielding approximately 0.95. This process was repeated for delays of 1 week, 2 weeks, 1 month, 2 months, 6 months, 8 months, 12 months, and 14 months (displayed to participants in units of days).

On trials where the delay corresponded to the time of occurrence of one of the previously constructed events, the title for that event provided by the participant in the future thinking writing task was displayed beneath the delayed option (henceforth "EFT cuing"; Peters and Büchel, 2010).

Delay discounting rates were calculated according to the hyperbolic model (Mazur, 1987):

$$I_D = \frac{1}{1 + kD} \tag{3.1}$$

where I_D is the individual's indifference point at delay D and k is a free parameter quantifying the steepness of an individual's discounting function. k values were computed using nonlinear least squares estimation as implemented in R's `nls` function (R Core Team, 2018). For each participant, k values were computed separately for the cued and uncued trials. For statistical analysis, k values were log-transformed for normality (Kirby et al., 1999). The effect of EFT cuing on delay discounting, Δ_k, was computed as follows:

$$\Delta_k = \log k_{\text{uncued}} - \log k_{\text{cued}} \tag{3.2}$$

where k_{uncued} is computed by fitting the model in Eq. 3.1 to the data from an individual's uncued trials and k_{cued} is computed using data from cued trials. A more positive value of Δ_k reflects larger decreases in delay discounting during EFT cuing (i.e., less impulsive decision making).

3.2.5 Analyses

All analyses were performed using the R statistical computing environment (R Core Team, 2018). Linear mixed effects models were computed using the `lmer` function from the Analysis of Factorial Experiments (`afex`) package (Singmann et al., 2019) and p values for fixed effects were computed using degrees of freedom estimated by the Satterthwaite approximation (Satterthwaite, 1946).

Briefly, linear mixed effects models are used for data in which observations occur in clusters, and they account for similarity between observations from the same cluster (Hedeker, 2005). In the current case, there were 4 sets of event-wise ratings collected from each participant. Every mixed effects model reported here used random intercepts, meaning that for a given pair of variables, all within-subjects correlations were described using a single best-fitting regression slope. These slopes are reported as β statistics along with their associated p values.

Figures were generated using the Grammar of Graphics package (Wickham, 2016).

3.3 Results

3.3.1 Joint distribution of visual perspective dimensions

We first aimed to characterize the relationship between the two proposed dimensions of visual perspective and determine whether they are indeed separable. To this end, Fig. 3.1a illustrates the bivariate distribution along both dimensions of all imagined future events pooled from all participants. There was a strong negative association between the two dimensions, as indicated by a linear mixed effects model (dependent variable: third-person self-visibility; fixed effect: first-person perspective; random effect: participant identity; β = -0.301, p < 0.0001).

We expected to find evidence of both acentred imagery and perspective switching (prediction 1). We defined acentred imagery as low first-person predominance and low third-person self-visibility (both ratings less than 50/100, i.e., the bottom left quadrant of Fig. 3.1a). 17 imagined future events (5.06%) met these criteria and 13 participants (15.48%) reported at least one instance of acentred imagery. Similarly, we defined perspective switching as high ratings of both first-person predominance and third-person self-visibility (both ratings greater than 50/100, i.e., the top right quadrant of Fig. 3.1a). 90 events (26.79%) met these criteria and 50 participants (59.52%) reported at least one instance of perspective switching. First-person imagery was defined as ratings of first-person predominance greater than 50/100 and ratings of third-person self-visibility below 50/100, or not applicable for events with 100/100 first-person prefominance. 121 events (36.01%) met these criteria and 58 participants (69.05%) reported at least one such event. Finally, third-person imagery was defined as low (<50/100) ratings of first-person predominance and high (>50/100) ratings of third-person self-visibility. 68 events (20.24%) met these criteria and 38 participants (45.24%) reported at least one such event.

Thus, while acentred imagery was relatively uncommon, in line with prediction 1 we found evidence for its occurrence, along with the occurrence of perspective switching. This suggests that the two dimensions of visual perspective, while anticorrelated, are separable.

Exploratory analyses

Low vividness ratings appear overrepresented in the bottom left quadrant of Fig. 3.1a, likely because participants with unclear mental images would endorse neither a clear first-person perspective nor clear third-person self-visibility. Indeed, a linear mixed effects model predicting vividness ratings using the two dimensions of visual perspective (fixed effects: first-person predominance and third-person self-visibility; random effect: participant identity) found that both were independently correlated with vividness (first-person predominance: $\beta = 0.184$, $p = 0.0030$; third-person self-visibility: β = 0.383, p < 0.0001). This means that apparent cases of acentred imagery may actually have merely been very faint mental images.

Visualizing the same data on a by-participant basis, it becomes clear that there was an outlier participant reporting virtually no visual imagery (Fig. 3.1b). The above mixed effects analyses remain significant after removing this participant's data, i.e., vividness ratings were still independently predicted by both first-person predominance ($\beta = 0.142$ $p = 0.0292$) and third-person self-visibility ($\beta = 0.343$, p < 0.0001). However, to avoid exaggerating the role of imagery vividness in determining visual perspective, this participant's data were removed from subsequent analyses.

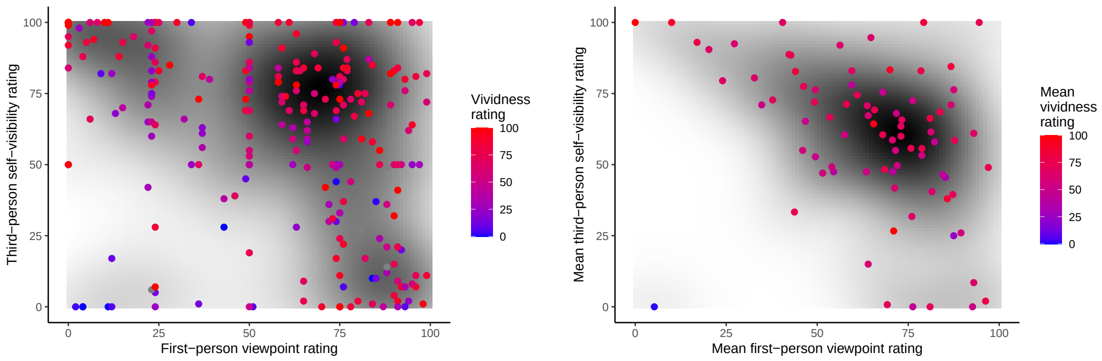

(a) (b)

Figure 3.1: Third-person self-visibility as a function of first-person predominance. Background shading reflects 2-dimensional kernel density estimates. (a) Each data point corresponds to a pair of ratings for a single event. Gray points are those for which participants did not submit vividness ratings. (b) Each data point corresponds to the mean response for a single participant.

3.3.2 State (event-wise) predictors of visual perspective dimensions

We next aimed to identify associations between each dimension of visual perspective and the self-report ratings provided for each event (i.e., vividness, location familiarity, emotional valence, and emotional intensity; see section 3.2.3). To this end, we computed 2 linear mixed effects models, first modelling first-person predominance as a function of all other event-wise ratings, then doing the same for third-person self-visibility. Temporal distance was included as a fixed effect in these models and coded as an ordinal variable.

We expected first-person predominance to be positively related to vividness ratings and negatively related to temporal distance (prediction 2). We also expected emotional intensity to be positively related to first-person predominance (prediction 3). However, the only significant negative predictor of first-person predominance was third-person self-visibility (β = -0.383, p $<$ 0.0001) and the only positive predictor was emotional valence ($\beta = 0.122$, $p = 0.0414$; higher ratings indicate more positive valence). While the relationships of first-person predominance to the other variables were in the expected directions (vividness: $\beta = 0.089$; temporal distance: $\beta = -0.148$; emotional intensity: $\beta = 0.080$), none was significant (all p values > 0.12).

Our only prediction around third-person self-visibility was that it would be associated with increased emotional intensity (prediction 3). This was borne out in a second linear mixed effects model, which found that third-person self-visibility was positively predicted by emotional intensity ($\beta = 0.186$, $p = 0.0004$) and vividness ($\beta = 0.243$, p < 0.0001), and negatively predicted by first-person predominance (β = -0.386, p < 0.0001). No other event-wise predictor was significant (all p values > 0.2).

Thus, between predictions 2 and 3, only part of prediction 3 was statistically supported: there was an association between emotional intensity and third-person self-visibility, in line with the Sutin and Robins (2008) model. We also found an unexpected association between ratings of vividness and third-person self-visibility.

Exploratory analyses

In section 3.3.1, there was an association between ratings of vividness and first-person predominance. However, there was no association between these two variables in section 3.3.2, in a model that included other event-wise ratings. This may be because, in EFT, the effect of temporal distance on vividness is mediated by location familiarity: temporally distant events are less vivid because they are imagined in unfamiliar locations (Arnold et al., 2011). Therefore, to address potential collinearity between vividness, temporal distance, and location familiarity, the linear mixed effect model predicting first-person predominance was re-computed keeping each of these variables in turn as a fixed effect while excluding the other two. I.e. the model was first re-computed excluding temporal distance and location familiarity as predictors, then excluding vividness and location familiarity, then excluding vividness and temporal distance.

Using this approach, both vividness (β = 0.129, p = 0.0179) and location familiarity ($\beta = 0.108$, $p = 0.0137$) were significant predictors of first-person predominance. The temporal distance linear contrast was negative as expected (i.e. decreasing first-person predominance with increasing temporal distance) but not significant (t = -0.565, p = 0.57). Thus both vividness and location familiarity predict first-person predominance, but not independently. This is likely because more familiar locations tend to be imagined more vividly.

3.3.3 Trait predictors of visual perspective dimensions

We next sought to test our predictions around trait variables and their relationships to the two proposed dimensions of visual perspective. We expected trait dissociation (DES score) to be negatively correlated with both first-person predominance and third-person self-visibility (prediction 4). We also expected trait imagery vividness (Psi-Q visual score) to be positively correlated with first-person predominance (prediction 5).

However, in a multivariate linear regression predicting DES score, there was no main effect of first-person predominance (β = -0.356, $p = 0.21$) or third-person self-visibility (β = -0.394, $p = 0.14$), despite the associations being in the expected directions. Similarly, in a second multivariate linear regression predicting Psi-

Q visual score, the main effect of first-person predominance was in the expected direction but not significant ($\beta = 0.032$, $p = 0.13$). However, there was a marginally significant main effect of third-person self-visibility ($\beta = 0.042$, $p = 0.0403$). Thus the expected associations between trait dissociation and trait imagery vividness and the two proposed dimensions of visual perspective were not borne out.

Exploratory analyses

Although trait dissociation was not linearly associated with either dimension of visual perspective, in separate polynomial regressions, both had inverse quadratic relationships to DES score (first-person $p = 0.0277$; third-person $p = 0.0471$). That is, those with high trait dissociation were, on average, not tied to one perspective or the other. This raises the possibility that trait dissociation is related to the tendency to switch between perspectives when visualizing a single event.

To test this hypothesis, "dual-perspective" events were identified (as in section 3.3.1) as those receiving both high first-person predominance ratings and high third-person self-visibility ratings (both ratings greater than 50/100, i.e., events in the top right quadrant of Fig. 3.1a). The distribution of the number of these events reported by participants is displayed in Fig. 3.2a. The number of dual-perspective events (coded as an ordinal variable) was associated with trait dissociation (DES score; linear contrast: t = 2.84, $p = 0.0058$; Fig. 3.2b). Thus, trait dissociation was associated with a greater tendency to switch between first- and third-person visual perspectives.

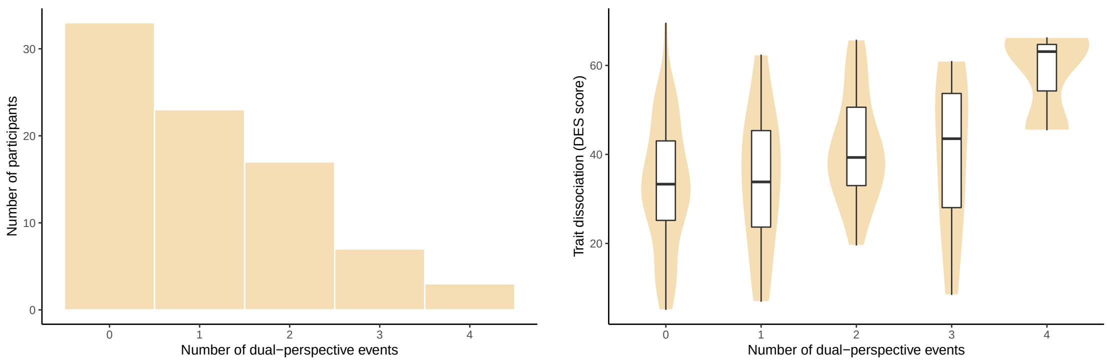

(a) (b)

Figure 3.2: "Dual-perspective events" are those rated by participants as involving high first-person predominance and high third-person self-visibility. (a) The distribution of the number of dual-perspective events tapers off toward the maximum value of 4. (b) Trait dissociation is associated with a greater number of dual-perspective events (i.e., perspective switching).

Finally, as in previous research, trait imagery vividness and trait dissociation were uncorrelated ($r(81)$ = -0.104, $p = 0.35$; Koppenhaver et al., 1997; Vannucci and Mazzoni, 2009).

3.3.4 Delay discounting

We next aimed to identify whether either dimension of visual perspective was correlated with the EFT cue effect on delay discounting, Δ_k (Eq. 3.2). As noted earlier, Δ_k was computed as the logarithm of the k value (the slope parameter of the hyperbolic discounting curve; Eq. 3.1) computed from uncued trials minus the logarithm of the k value computed from cued trials (Eq. 3.2; Fig. 3.3a). Thus positive Δ_k values indicate shifts toward greater intertemporal patience as a result of EFT cuing. The distribution of this statistic contained outliers (Fig. 3.3b), so the following analyses below were completed after removing boxplot outliers (7 total). A significant main effect of the EFT cues on delay discounting rates remained after removing these outliers (t(75) = 3.266, $p = 0.0016$).

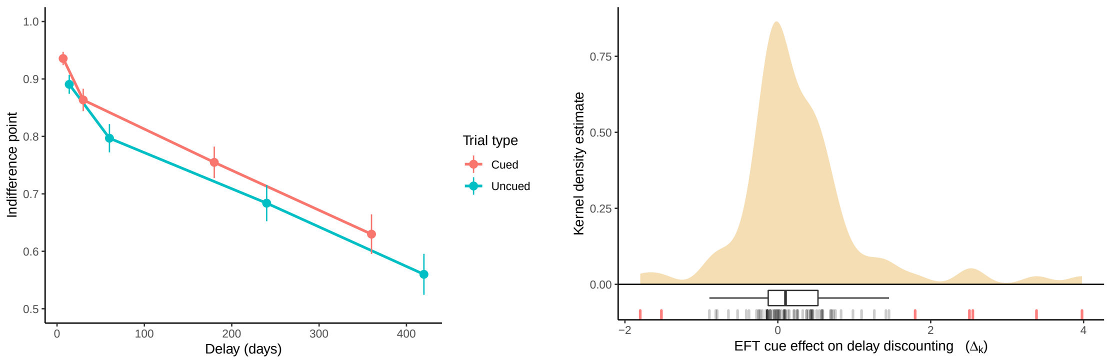

(a) (b)

Figure 3.3: EFT cue effect on delay discounting. (a) Mean empirical discount functions computed for cued and uncued trials. Error bars reflect the mean ± one standard error of the mean. (b) Distribution of the EFT cue effect on delay discounting, as computed using Eq. 3.2. The distribution contains outliers (displayed in red), particularly at the higher end.

Event-wise predictors of EFT effect on delay discounting

We expected (prediction 6) that the EFT cue effect on delay discounting would be amplified by both first-person predominance (through increased vividness) and third-person self-visibility (through the "salient self" effect). However, in a multivariate linear regression predicting Δ_k using each average event-wise rating (i.e., mean vividness rating, mean first-person predominance rating, mean third-person self-visibility rating, etc.), none was significant (all p values > 0.3).

Exploratory analyses

Imagery vividness has been found to be associated with reward sensitivity and delay discounting (Linke and Wessa, 2017; Parthasarathi et al., 2017). Therefore we checked for an association between baseline delay discounting (k values computed using uncued trials, log-transformed) and trait imagery vividness (Psi-Q visual score). However, none was found ($r(74)$ = -0.093, p = 0.42). Unexpectedly, though, higher trait dissociation predicted steeper baseline delay discounting ($r(74) = 0.296$, $p = 0.0095$; Fig. 3.4). This effect was also significant when Δ_k outliers were included ($r(81) = 0.330$, $p = 0.0023$).

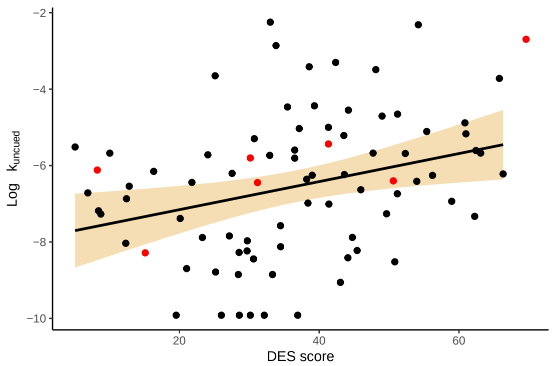

Figure 3.4: Delay discounting rate (log k, computed using uncued trials) as a function of trait dissociation (DES score). Data from participants who were EFT cue effect (Δ_k) boxplot outliers are shown in red. The trendline was computed ignoring these. Shaded error region reflects a 95% confidence interval.

3.4 Discussion

The present work tested a novel conceptualization of visual perspective as comprising the two dimensions of first-person predominance and third-person self-visibility. Several lines of evidence support this model. First, the two dimensions, while anticorrelated, were separable. That is, we found evidence for both acentred imagery (low scores on both dimensions) and perspective switching (high scores on both dimensions). Second, the two dimensions were non-redundant, with both independently predicting ratings of vividness. Finally, in line with the model of Sutin and Robins (2008), third-person self-visibility was uniquely correlated with emotional intensity. We next discuss these findings in detail.

3.4.1 Prevalence of acentred imagery

While the present study found evidence of acentred imagery, the phenomenon appears to be relatively rare. Moreover, the only participant consistently meeting this study's definition of acentred imagery may have bordered on aphantasia. Thus some cases of acentred imagery may have been a byproduct of the fact that it was defined and measured implicitly, as a lack of both first-person predominance and third-person self-visibility. That is, a participant with no mental image of an event at all could affirm neither a clear first-person perspective nor clear self-visibility. According to this study's definition, this would be mistakenly labelled as an acentred perspective. Future research aiming to measure acentred imagery or memory should do so directly by positively querying characteristics of acentred mental images. For example, participants could be directly asked how often they experience memories or mental images of the personal future which are not from the their own point of view and in which they themselves do not visually appear.

Many events in our sample that were imagined mainly but not entirely from a first-person perspective were rated as also not involving high third-person self-visibility (i.e., the bottom edge of the bottom right quadrant of Fig. 3.1a). Thus, at least some of the time, participants were not imagining these events from a first-person perspective, nor were they imagining the appearances of their future selves. It could be argued that these cases warrant being called "acentred". However, both deliberate memory recall and EFT involve two stages: event construction and elaboration (Addis et al., 2007). Only the elaboration phase is associated with imagery-related activity in the precuneus and retrosplenial cortex, both major nodes of the default mode network (Raichle, 2015). In constrast, the construction phase is associated less with imagery and more with searching memory and binding episodic details into a single trace (Addis et al., 2007). The apparent cases of acentred imagery reported for events that were imagined mainly from a first-person perspective may actually have occurred during the construction phase, when the mental image was still taking shape and was too vague to involve high self-visibility or to definitively be assigned a first-person perspective. In this case, once again, apparent cases of acentred imagery would be artifacts of the implicit definition used in this study.

Still, it is worth noting that Wollheim (1984) originally described acentred perspective as "unstable":

> [I]f I am not amongst the dramatis personae of the memory—and I cannot be if I acentrally remember it—this must be because I have edited myself out of it. But not far out of it: I wait in the wings. The situation is unstable, and at any moment I may, I am likely to, reassume in the representation of the past event the very part that I played in the event itself. And when this happens, acentred event-memory converts itself into centred event-memory.

This conversion may simply be the transition from construction to elaboration. Thus, consistent with the results of the present work, it may be that acentred perspective largely exists only in the vague imagery that occurs as one begins to remember or imagine an episode. Indeed, recent neuroscientific work has found that the neural correlates of memory retrieval differ as a function of visual perspective mainly during the elaboration phase, suggesting that memories take time to develop a distinct perspective (Iriye and Jacques, 2020; Hebscher et al., 2020).

3.4.2 Correlates of visual perspective dimensions

In addition to replicating the widely-reported association between vividness and first-person perspective, the present work also found a positive correlation between vividness and third-person self-visibility. This shows that this latter dimension is not redundant and adds an important clarification to the relationship between vividness and visual perspective: self-visibility by definition requires visual detail, thus, third-person mental images can be vivid when the imagined self is vivid.

Furthermore, the relationship between third-person self-visibility and emotional intensity aligns well with the model of visual perspective in memory proposed by Sutin and Robins (2008), which states that third-person perspective heightens affect when the episode is consistent with the present self-concept by increasing the salience of the remembered (in this case, imagined) self (the "salient self" effect). That model also predicts that third-person perspective will dampen affect when the episode is inconsistent with the present self-concept (the "dispassionate observer" effect). We had thus expected to find a positive correlation between emotional intensity and first-person predominance (since first-person predominance should negate any affective dampening arising from the "dispassionate observer" effect). This prediction was not borne out, possibly because, as explained below, very few imagined future events in the present study would have been inconsistent with participants' self-concepts.

In general, mental images of the future seem to be organized around personal goals (D'Argembeau and Mathy, 2011). Autobiographical knowledge is thought to constitute a hierarchy, with personal semantic knowledge and goals at the top and episodic details at the bottom (Conway and Pleydell-Pearce, 2000; Conway et al., 2019). When episodic future thoughts are intentionally constructed, this hierarchy is traversed downward, a process that is accelerated when specific personal goals are used as cues (Cole and Kvavilashvili, 2019; D'Argembeau and Mathy, 2011). Thus it is likely that participants' imagined future events generally drew on personal semantic information and were therefore consistent with their (then-) present self-concepts. In this case the "dispassionate observer" effect may not have occurred to a statistically detectable extent. In may be fruitful in future research to ask participants to imagine events that both are and are not consistent with their self concepts, in order to measure both the "dispassionate observer" and "salient self" effects in terms of first-person predominance and third-person self-visibility.

3.4.3 Dissociation and perspective switching

In contrast to previous work showing a correlation between trait dissociation and third-person perspective in memory (Sutin and Robins, 2010; Radvansky and Svob, 2019), we found that in EFT, trait dissociation was associated with a greater tendency to switch between perspectives rather than to adopt one perspective or the other per se. There are at least three potential explanations for this discrepancy: first, given the distinctiveness of third-person memory, it could be that participants in previous work reported third-person perspective in memory recall when in fact they switched between perspectives. Alternatively, it could be that the relationship between dissociation and visual perspective is simply different between memory and EFT. This would be similar to the case of autism, which is associated with lower rates of first-person perspective in episodic memory but not EFT (Lind and Bowler, 2010). Finally, dissociation is linked to distractibility and fantasy-proneness, which is argued to lead to the re-experiencing of multiple, sometimes incongruent, emotions during memory recall (Giesbrecht et al., 2008; Sutin and

Stockdale, 2011). These tendencies may similarly be related to switching between visual perspectives. Further research querying the frequency of perspective switching in both memory and EFT is required to adjudicate between these possibilities.

3.4.4 Delay discounting

In examining the impact of visual perspective on decision making, we expected the effect of EFT on delay discounting to be correlated with both first-person predominance (through increased vividness) and third-person self-visibility (through the "salient self" effect; Sutin and Robins, 2008). However, this prediction was not borne out, possibly due to the brevity of the delay discounting task used. While the 5-trial task used in this study is a valid measure of delay discounting (Koffarnus and Bickel, 2014), studies that have identified covariates of the EFT effect have tended to use more than 200 trials total (Peters and Büchel, 2010; O'Donnell et al., 2017), compared to the present study's 40.

Furthermore, the most temporally distant imagined events in the present study were a year into the future, while research finding that third-person visualization increases saving behaviour had participants visualize events 40 years in the future (Macrae et al., 2017). McCarroll (2019) notes that when third-person imagery increases emotional intensity, it is usually for abstract emotions like pride that are more closely identified with goals in the far future than goals in the near future. Thus, the effect of EFT on delay discounting may be moderated by visual perspective only when making decisions about the far future.

Finally, the event-wise ratings used as correlates of the EFT cue effect were collected after the future thinking writing task and prior to the delay discounting task. In contrast, Peters and Büchel (2010) queried the imagery evoked by cues in the delay discounting task to identify correlates of the cue effect. In the present study, a participant may have experienced vivid imagery during the writing task but no imagery during the delay discounting task, meaning that the ratings collected during the writing task would be irrelevant to the EFT cue effect.

The present study incidentally identified a correlation between delay discounting and dissociation which is, to our knowledge, novel. There does not appear to be an obvious conceptual link between the constructs of impuslivity and dissociation, but delay discounting is a transdiagnostic process across numerous psychiatric disorders (Amlung et al., 2019). It is possible that dissociation would not predict steep discounting when controlling for, e.g., potential comorbid depression or anxiety. On the other hand, it could be that dissociation is yet another clinical construct related to delay discounting. Dissociation is not generally considered to involve altered reward valuation, but it could be that the feeling of detachment from the present that manifests in dissociation extends to a feeling of detachment from the personal future. Similarly, the depersonalization characteristic of dissociation could interfere with the complex metacognitive judgments required when making decisions about the future (Hoerl and McCormack, 2016; Bulley and Schacter, 2020). While these explanations are speculative, if the present findings turn out to be robust, the possibility of a relationship between dissociation and delay discounting may be a fruitful avenue for future research.

3.5 Conclusions

The present work provides evidence for a novel conceptualization of visual perspective in episodic imagery as a two-dimensional construct, comprising first-person predominance and third-person self-visibility. While not all of our specific predictions regarding trait and state correlates of these were borne out, a two-dimensional understanding of visual perspective is nonetheless useful: both dimensions independently predicted vividness ratings, and third-person self-visibility was uniquely associated with emotional intensity. Furthermore, this understanding can clarify the divergent effects of third-person perspective on emotion/motivation: detachment from a first-person perspective should index the "dispassionate observer" effect, which dampens affect, while third-person self-visibility should index the "salient self" effect, which intensifies affect (Sutin and Robins, 2008). While the present study found statistical support only for the latter assertion, the former is well-supported in the literature (McIsaac and Eich, 2002; D'Argembeau et al., 2003; McIsaac and Eich, 2004; Talarico et al., 2004; Sekiguchi and Nonaka, 2014).

Future research may find it fruitful to extend this model to studies of memory and to search for distinct cognitive and neural underpinnings of each of these dimensions. For example, first-person predominance may be closely related to embodiment and bodily representations in, e.g., the insula (St. Jacques, 2019), while third-person self-visibility may rely on propositional self-beliefs and self-referential thought supported by the cortical midline structures (Libby et al., 2014; Northoff and Bermpohl, 2004). A more mechanistic understanding of visual perspective as a two-dimensional construct would go a long way toward explaining its role in psychopathology and identifying how it can be best harnessed as a tool to improve motivation and decision making.

Bibliography

Addis, D. R., Wong, A. T., and Schacter, D. L. (2007). Remembering the past and imagining the future: common and distinct neural substrates during event construction and elaboration. *Neuropsychologia*, 45(7):1363–1377.

Akhtar, S., Justice, L. V., Loveday, C., and Conway, M. A. (2017). Switching memory perspective. *Consciousness and cognition*, 56:50–57.

Amlung, M., Marsden, E., Holshausen, K., Morris, V., Patel, H., Vedelago, L., Naish, K. R., Reed, D. D., and McCabe, R. E. (2019). Delay discounting as a transdiagnostic process in psychiatric disorders: a meta-analysis. *JAMA psychiatry*, 76(11):1176–1186.

Andrade, J., May, J., Deeprose, C., Baugh, S.-J., and Ganis, G. (2014). Assessing vividness of mental imagery: the plymouth sensory imagery questionnaire. *British Journal of Psychology*, 105(4):547–563.

Arnold, K. M., McDermott, K. B., and Szpunar, K. K. (2011). Imagining the near and far future: The role of location familiarity. *Memory & Cognition*, 39(6):954–967.

Atance, C. M. and O'Neill, D. K. (2001). Episodic future thinking. *Trends in cognitive sciences*, 5(12):533–539.

Bagri, G. and Jones, G. V. (2009). Category-specific enhancement of retrieval due to field perspective. *Memory*, 17(3):337–345.

Belardinelli, M. O., Palmiero, M., Sestieri, C., Nardo, D., Di Matteo, R., Londei, A., D'Ausilio, A., Ferretti, A., Del Gratta, C., and Romani, G. L. (2009). An fmri investigation on image generation in different sensory modalities: the influence of vividness. *Acta psychologica*, 132(2):190–200.

Benoit, R. G., Gilbert, S. J., and Burgess, P. W. (2011). A neural mechanism mediating the impact of episodic prospection on farsighted decisions. *Journal of Neuroscience*, 31(18):6771–6779.

Berntsen, D. and Rubin, D. C. (2006). Emotion and vantage point in autobiographical. *Cognition and Emotion*, 20(8):1193–1215.

Brewin, C. R., Gregory, J. D., Lipton, M., and Burgess, N. (2010). Intrusive images in psychological disorders: characteristics, neural mechanisms, and treatment implications. *Psychological review*, 117(1):210.

Bulley, A. and Schacter, D. L. (2020). Deliberating trade-offs with the future. *Nature Human Behaviour*, 4(3):238–247.

Butler, A. C., Rice, H. J., Wooldridge, C. L., and Rubin, D. C. (2016). Visual imagery in autobiographical memory: The role of repeated retrieval in shifting perspective. *Consciousness and cognition*, 42:237–253.

Byrne, P., Becker, S., and Burgess, N. (2007). Remembering the past and imagining the future: a neural model of spatial memory and imagery. *Psychological review*, 114(2):340.

Callow, N. and Roberts, R. (2010). Imagery research: An investigation of three issues. *Psychology of Sport and Exercise*, 11(4):325–329.

Carlson, E. B. and Putnam, F. W. (1993). An update on the dissociative experiences scale. *Dissociation: progress in the dissociative disorders*.

Clark, J. M. and Paivio, A. (2004). Extensions of the paivio, yuille, and madigan (1968) norms. *Behavior Research Methods, Instruments, & Computers*, 36(3):371–383.

Cole, S. and Kvavilashvili, L. (2019). Spontaneous and deliberate future thinking: a dual process account. *Psychological Research*, pages 1–16.

Conway, M. A., Justice, L. V., and D'Argembeau, A. (2019). The self-memory system revisited. *The organization and structure of autobiographical memory*, page 28.

Conway, M. A. and Pleydell-Pearce, C. W. (2000). The construction of autobiographical memories in the self-memory system. *Psychological review*, 107(2):261.

Crawley, S. and French, C. (2005). Field and observer viewpoint in remember-know memories of personal childhood events. *Memory*, 13(7):673–681.

Crovitz, H. F. and Schiffman, H. (1974). Frequency of episodic memories as a function of their age. *Bulletin of the Psychonomic Society*, 4(5):517–518.

D'Argembeau, A., Comblain, C., and Van der Linden, M. (2003). Phenomenal characteristics of autobiographical memories for positive, negative, and neutral events. *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition*, 17(3):281–294.

D'Argembeau, A. and Mathy, A. (2011). Tracking the construction of episodic future thoughts. *Journal of experimental psychology: General*, 140(2):258.

De Leeuw, J. R. (2015). jspsych: A javascript library for creating behavioral experiments in a web browser. *Behavior research methods*, 47(1):1–12.

D'Argembeau, A. and Van der Linden, M. (2006). Individual differences in the phenomenology of mental time travel: The effect of vivid visual imagery and emotion regulation strategies. *Consciousness and cognition*, 15(2):342–350.

Epstein, M. L. (1980). The relationship of mental imagery and mental rehearsal to performance of a motor task. *Journal of Sport and Exercise Psychology*, 2(3):211–220.

Freud, S. (1962). Screen memories. In *The Standard Edition of the Complete Psychological Works of Sigmund Freud, Volume III (1893-1899): Early Psycho-Analytic Publications*, pages 299–322. Hogarth Press.

Galton, F. (1879). Psychometric experiments. *Brain*, 2(2):149–162.

Giesbrecht, T., Lynn, S. J., Lilienfeld, S. O., and Merckelbach, H. (2008). Cognitive processes in dissociation: an analysis of core theoretical assumptions. *Psychological bulletin*, 134(5):617.

Gordon, S., Weinberg, R., and Jackson, A. (1994). Effect of internal and external imagery on cricket performance. *Journal of Sport Behavior*, 17(1):60.

Gu, X. and Tse, C.-S. (2016). Narrative perspective shift at retrieval: The psychological-distance-mediated-effect on emotional intensity of positive and negative autobiographical memory. *Consciousness and cognition*, 45:159–173.

Hebscher, M., Ibrahim, C., and Gilboa, A. (2020). Precuneus stimulation alters the neural dynamics of autobiographical memory retrieval. *Neuroimage*, 210:116575.

Hedeker, D. (2005). Generalized linear mixed models. *Encyclopedia of statistics in behavioral science.*

Hishitani, S., Miyazaki, T., and Motoyama, H. (2011). Some mechanisms responsible for the vividness of mental imagery: Suppressor, closer, and other functions. *Journal of mental imagery*, 35(1&2):5–32.

Hoerl, C. and McCormack, T. (2016). Making decisions about the future. *Seeing the future: Theoretical perspectives on future-oriented mental time travel*, page 241.

Iriye, H. and Jacques, P. L. S. (2020). How visual perspective influences the spatiotemporal dynamics of autobiographical memory retrieval. *Cortex*, 129:464–475.

Isaac, A. R. and Marks, D. F. (1994). Individual differences in mental imagery experience: developmental changes and specialization. *British Journal of Psychology*, 85(4):479–500.

Johnson, M. K., Foley, M. A., Suengas, A. G., and Raye, C. L. (1988). Phenomenal characteristics of memories for perceived and imagined autobiographical events. *Journal of Experimental Psychology: General*, 117(4):371.

Kenny, L. M. and Bryant, R. A. (2007). Keeping memories at an arm's length: Vantage point of trauma memories. *Behaviour research and therapy*, 45(8):1915–1920.

Kirby, K. N., Petry, N. M., and Bickel, W. K. (1999). Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. *Journal of Experimental psychology: general*, 128(1):78.

Koffarnus, M. N. and Bickel, W. K. (2014). A 5-trial adjusting delay discounting task: Accurate discount rates in less than one minute. *Experimental and clinical psychopharmacology*, 22(3):222.

Koppenhaver, J. M., Kumar, V., and Pekala, R. J. (1997). Dissociativity, imagery vividness, and reality monitoring. *Dissociation: Progress in the Dissociative Disorders.*

Kuyken, W. and Howell, R. (2006). Facets of autobiographical memory in adolescents with major depressive disorder and never-depressed controls. *Cognition and Emotion*, 20(3-4):466–487.

Kuyken, W. and Moulds, M. L. (2009). Remembering as an observer: How is autobiographical memory retrieval vantage perspective linked to depression? *Memory*, 17(6):624–634.

Lemogne, C., Piolino, P., Friszer, S., Claret, A., Girault, N., Jouvent, R., Allilaire, J.-F., and Fossati, P. (2006). Episodic autobiographical memory in depression: Specificity, autonoetic consciousness, and self-perspective. *Consciousness and cognition*, 15(2):258–268.

Libby, L. K. and Eibach, R. P. (2011). Visual perspective in mental imagery: A representational tool that functions in judgment, emotion, and self-insight. In *Advances in experimental social psychology*, volume 44, pages 185–245. Elsevier.

Libby, L. K., Shaeffer, E. M., Eibach, R. P., and Slemmer, J. A. (2007). Picture yourself at the polls: Visual perspective in mental imagery affects self-perception and behavior. *Psychological Science*, 18(3):199–203.

Libby, L. K., Valenti, G., Hines, K. A., and Eibach, R. P. (2014). Using imagery perspective to access two distinct forms of self-knowledge: Associative evaluations versus propositional self-beliefs. *Journal of Experimental Psychology: General*, 143(2):492.

Lind, S. E. and Bowler, D. M. (2010). Episodic memory and episodic future thinking in adults with autism. *Journal of abnormal psychology*, 119(4):896.

Linke, J. and Wessa, M. (2017). Mental imagery training increases wanting of rewards and reward sensitivity and reduces depressive symptoms. *Behavior Therapy*, 48(5):695–706.

Lyssenko, L., Schmahl, C., Bockhacker, L., Vonderlin, R., Bohus, M., and Kleindienst, N. (2018). Dissociation in psychiatric disorders: a meta-analysis of studies using the dissociative experiences scale. *American Journal of Psychiatry*, 175(1):37–46.

Macrae, C. N., Mitchell, J. P., Golubickis, M., Ho, N. S., Sherlock, R., Parlongo, R., Simpson, O. C., and Christian, B. M. (2017). Saving for your future self: The role of imaginary experiences. *Self and Identity*, 16(4):384–398.

Marcotti, P. and St. Jacques, P. L. (2018). Shifting visual perspective during memory retrieval reduces the accuracy of subsequent memories. *Memory*, 26(3):330–341.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. *Commons, ML.; Mazur, JE.; Nevin, JA*, pages 55–73.

McCarroll, C. (2018). *Remembering from the outside: Personal memory and the perspectival mind*. Oxford University Press.

McCarroll, C. J. (2019). Navigating intertemporal choices: Mental time travel, perspectival imagery, and prudent decision-making. *Psychology of Consciousness: Theory, Research, and Practice*, 6(2):200.

McFadden, E. and Siedlecki, K. L. (2020). Do depressive symptoms and subjective well-being influence the valence or visual perspective of autobiographical memories in young adults? *Memory*, pages 1–10.

McIsaac, H. K. and Eich, E. (2002). Vantage point in episodic memory. *Psychonomic bulletin & review*, 9(1):146–150.

McIsaac, H. K. and Eich, E. (2004). Vantage point in traumatic memory. *Psychological Science*, 15(4):248–253.

Morris, T. and Spittle, M. (2012). A default hypothesis of the development of internal and external imagery perspectives. *Journal of mental imagery*, 36(1 & 2):1–30.

Nelis, S., Debeer, E., Holmes, E. A., and Raes, F. (2013). Dysphoric students show higher use of the observer perspective in their retrieval of positive versus negative autobiographical memories. *Memory*, 21(4):423–430.

Nicolas, S., Gounden, Y., and Piolino, P. (2013). Victor and catherine henri on earliest recollections. *LAnnee psychologique*, 113(3):349–374.

Nigro, G. and Neisser, U. (1983). Point of view in personal memories. *Cognitive psychology*, 15(4):467–482.

Northoff, G. and Bermpohl, F. (2004). Cortical midline structures and the self. *Trends in cognitive sciences*, 8(3):102–107.

O'Donnell, S., Daniel, T. O., and Epstein, L. H. (2017). Does goal relevant episodic future thinking amplify the effect on delay discounting? *Consciousness and cognition*, 51:10–16.

Parthasarathi, T., McConnell, M. H., Luery, J., and Kable, J. W. (2017). The vivid present: Visualization abilities are associated with steep discounting of future rewards. *Frontiers in psychology*, 8:289.

Pearson, J. (2019). The human imagination: the cognitive neuroscience of visual mental imagery. *Nature Reviews Neuroscience*, 20(10):624–634.

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

Piolino, P., Desgranges, B., Clarys, D., Guillery-Girard, B., Taconnat, L., Isingrini, M., and Eustache, F. (2006). Autobiographical memory, autonoetic consciousness, and self-perspective in aging. *Psychology and aging*, 21(3):510.

R Core Team (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Radvansky, G. A. and Svob, C. (2019). Observer memories may not be for everyone. *Memory*, 27(5):647–659.

Raichle, M. E. (2015). The brain's default mode network. *Annual review of neuroscience*, 38:433–447.

Rice, H. J. (2010). Seeing where we're at: A review of visual perspective and memory retrieval. In Mace, J. H., editor, *The Act of Remembering: Toward an Understanding of How We Recall the Past*, chapter 10, pages 228–258. Wiley Online Library.

Rice, H. J. and Rubin, D. C. (2009). I can see it both ways: First-and third-person visual perspectives at retrieval. *Consciousness and cognition*, 18(4):877–890.

Rice, H. J. and Rubin, D. C. (2011). Remembering from any angle: The flexibility of visual perspective during retrieval. *Consciousness and Cognition*, 20(3):568–577.

Robinson, J. A. and Swanson, K. L. (1993). Field and observer modes of remembering. *Memory*, 1(3):169–184.

Sartre, J.-P. (1940/1972). *The Psychology of Imagination*. Methuen & Co.

Satterthwaite, F. E. (1946). An approximate distribution of estimates of variance components. *Biometrics bulletin*, 2(6):110–114.

Sekiguchi, T. and Nonaka, S. (2014). The long-term effect of perspective change on the emotional intensity of autobiographical memories. *Cognition & emotion*, 28(2):375–383.

Siedlecki, K. L. (2015). Visual perspective in autobiographical memories: Reliability, consistency, and relationship to objective memory performance. *Memory*, 23(2):306–316.

Singmann, H., Bolker, B., Westfall, J., and Aust, F. (2019). *afex: Analysis of Factorial Experiments*. R package version 0.23-0.

Smith, D., Collins, D., and Hale, B. (1998). Imagery perspectives and karate performance. *Journal of Sports Sciences*, 16(1):103–104.

St. Jacques, P. L. (2019). A new perspective on visual perspective in memory. *Current Directions in Psychological Science*, 28(5):450–455.

Stornelli, J., Pereira, B., and Vann, R. J. (2020). Big picture, bad outcomes: When visual perspectives harm health goal pursuit. *Journal of Consumer Psychology*, 30(2):368–378.

Sutin, A. R. and Robins, R. W. (2008). When the "i" looks at the "me": Autobiographical memory, visual perspective, and the self. *Consciousness and cognition*, 17(4):1386–1397.

Sutin, A. R. and Robins, R. W. (2010). Correlates and phenomenology of first and third person memories. *Memory*, 18(6):625–637.

Sutin, A. R. and Stockdale, G. D. (2011). Trait dissociation and the subjective affective, motivational, and phenomenological experience of self-defining memories. *Journal of personality*, 79(5):939–964.

Sutton, J. (2010). Observer perspective and acentred memory: Some puzzles about point of view in personal memory. *Philosophical Studies*, 148(1):27–37.

Sutton, J. (2012). Memory before the game: switching perspectives in imagining and remembering sport and movement. *Journal of mental imagery*, 36(1-2):85–95.

Talarico, J. M., LaBar, K. S., and Rubin, D. C. (2004). Emotional intensity predicts autobiographical memory experience. *Memory & cognition*, 32(7):1118–1132.

Taylor, H. A. and Tversky, B. (1992). Descriptions and depictions of environments. *Memory & cognition*, 20(5):483–496.

Trope, Y. and Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological review*, 117(2):440.

Vannucci, M. and Mazzoni, G. (2009). Individual differences in object and spatial imagery: Personality correlates. *Personality and Individual Differences*, 46(4):402–405.

Vasquez, N. A. and Buehler, R. (2007). Seeing future success: Does imagery perspective influence achievement motivation? *Personality and social psychology bulletin*, 33(10):1392–1405.

Vella, N. C. and Moulds, M. L. (2014). The impact of shifting vantage perspective when recalling and imagining positive events. *Memory*, 22(3):256–264.

Verhaeghen, P., Aikman, S. N., Doyle-Portillo, S., Bell, C. R., and Simmons, N. (2018). When i saw me standing there: first-person and third-person memories and future projections, and how they relate to the self. *Journal of Cognitive Psychology*, 30(4):438–452.

Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York.

Williams, A. D. and Moulds, M. L. (2007). Cognitive avoidance of intrusive memories: Recall vantage perspective and associations with depression. *Behaviour research and therapy*, 45(6):1141–1153.

Williams, A. D. and Moulds, M. L. (2008). Manipulating recall vantage perspective of intrusive memories in dysphoria. *Memory*, 16(7):742–750.

Wollheim, R. (1984). *The thread of life*. Yale University Press.

Chapter 4

Autonoetic consciousness may not explain the effect of episodic future thinking on delay discounting

Isaac Kinley, Morgan Porteous, Yarden Levy, & Suzanna Becker

Introductory note

In the previous chapter, we took a correlational approach to search for covariates of the future thinking cue effect. In this chapter, we use an experimental manipulation to induce a greater degree of variation in factors that might explain this effect, thereby increasing our statistical power.

In most experiments measuring the effect of episodic future thinking on delay discounting, an episodic future thinking condition is compared to a recent episodic memory condition. However, this prevents us from understanding whether episodic future thinking influences decision making because of the special qualities that constitute its "episodicity", or simply because it orients participants' attention to the future. For example, when the framing of a decision is changed from "*$100 in a year, or $50 now*" to "*$100 in a year and $0 now, or $50 now and $0 in a year*", people are more likely to select the delayed option. This is thought to occur because attention is oriented toward the future opportunity cost of choosing the immediate reward. Could such a simple mechanism similarly lie behind the effect episodic future thinking on delay discounting?

The factor that differentiates episodic from semantic memory is thought to be *autonoetic consciousness*, or the feeling of mentally re-experiencing a past event. Episodic future thinking similarly evokes a sense of mentally experiencing an event in one's future, which differentiates it from *semantic future thinking*. In this chapter, we operationalize semantic future thinking as calling to mind non-self-specific semantic knowledge about the future (for example, knowing about an upcoming holiday one does not celebrate). To compare the effect of episodic and semantic future thinking on delay discounting, participants were assigned either to an episodic or a semantic future thinking condition and completed an intertemporal choice task with future thinking cues. Because participants were cued to think about the future in both conditions, this experiment enables us to test whether *episodic* future thinking is necessary for the cues to have an effect. Moreover, we can test whether any factors differentiating between episodic and semantic future thinking mediate any unique effect of episodic future thinking on delay discounting.

The main author (I.K.) designed the study, collected the data, performed the analyses, and primarily wrote the paper. At the time of writing, it has been submitted for publication in Frontiers in Psychology: Kinley, I., Porteous, M., Levy, Y., & Becker, S. (2023). *Autonoetic consciousness may not explain the effect of episodic future thinking on impulsivity.* Manuscript submitted for publication in Frontiers in Psychology.

Abstract

Much recent work demonstrates that episodic future thinking (EFT) can reduce delay discounting, but the source of this effect remains poorly understood. The phenomenological hallmark of EFT is "autonoetic consciousness," which is the sense of "pre-experiencing" mentally simulated future events, analogous to the "re-experiencing" of a remembered episode. Here we sought to test whether autonoetic consciousness mediates the effect of EFT on delay discounting, given that mental simulations of delayed rewards are thought to play a important role in intertemporal choice, and self-reported autonoetic consciousness should index the success of such simulations. Across 2 experiments, participants imagined either events in their personal futures (EFT group) or non-self-specific future events (semantic future thinking group) and saw cues to these during a delay discounting task. The episodic group exhibited greater cue-driven reductions in delay discounting, but this effect was not mediated by their higher self-reported autonoetic consciousness. We speculate that cue-driven reductions in delay discounting may have resulted from semantic rather than episodic autobiographical representations evoked by self-relevant imagined future events.

4.1 Introduction

4.1.1 Delay discounting and episodic future thinking

"Delay discounting" or "temporal discounting" refers to the universal tendency to discount/devalue future rewards as a function of their delay (Odum, 2011a). The steepness with which a future reward's subjective value declines with increasing temporal distance is a relatively stable trait variable quantifying the ability to delay gratification, an important component of impulsivity (Odum, 2011b; Dalley and Robbins, 2017). Aberrant rates of temporal discounting have been identified as a transdiagnostic factor across numerous psychiatric conditions (Amlung et al., 2019). Steep discounting is also associated with non-clinical consequences such as poor academic performance (Kirby et al., 2005) and relational infidelity (Reimers et al., 2009).

A reliable method of shifting preference toward distal rewards is to evoke episodic future thinking (EFT), the mental simulation of the personal future (Atance and O'Neill, 2001), which relies on the same neural networks and perceptual content as episodic memory (Schacter and Addis, 2007). The standard experimental paradigm is to have decision-makers construct plausible future scenarios in their lives and then to provide cues to these during a delay discounting task (Peters and Büchel, 2010; Benoit et al., 2011). For decisions where these cues are displayed, participants exhibit shallower delay discounting (henceforth, the "cue effect"). Most often, the control condition in this paradigm displays cues to recent episodic memories, with the maximum age of these memories ranging from 24 hours (Daniel et al., 2016) to 12 days (Stein et al., 2018). As pointed out by Hollis-Hansen et al. (2019), more recent memories may be more likely to evoke spontaneous future thoughts, which could explain why studies using a range of memory ages have produced variable effect sizes (Rung and Madden, 2018).

A deeper issue with using an episodic memory control condition is that it compares two forms of episodic imagery that simply differ in temporal orientation, yet EFT is defined not only by its orientation toward the future but also by its *episodicity* (Perrin and Rousset, 2014); indeed, the concept of EFT came from an extension of the episodic–semantic distinction from memory to prospection (Atance and O'Neill, 2001). Given that episodicity is central to the definition of EFT, it may also be central to the role of EFT in reducing delay discounting. Yet studies using an episodic memory control condition cannot address whether EFT reduces discounting because it is episodic or because it is simply about the future. Could a non-episodic form of future thought produce the same cue effect as EFT? Before exploring this question, we must first consider what a non-episodic form of future thought might be.

4.1.2 Episodic and semantic future thinking

Tulving (1972) first posited episodic memory as a functional system distinguished by the fact that it processes self-specific and spatiotemporally constrained information. However, as noted by Perrin and Rousset (2014), his subsequent work placed less emphasis on the information processed by the episodic memory system and more on the recollective *experience* arising from the operation of that system. This experience, termed "autonoetic consciousness," characterizes the difference between merely knowing that an event occurred and mentally re-experiencing it during recall (Tulving, 1985).

As mentioned above, the concept of episodic future thinking comes from applying the episodic–semantic distinction to prospection (Atance and O'Neill, 2001): EFT involves the same scene imagery and autonoetic consciousness as episodic memory, and mentally simulating future experiences gives rise to a sense of "pre-experiencing" the future that is analogous to the sense of "re-experiencing" a remembered event (Addis, 2020; Atance and O'Neill, 2001). However, the correspondence between semantic future thinking and semantic memory is not as clear as the correspondence between episodic future thinking and episodic memory. According to Tulving (1985), while episodic memory is associated with "autonoetic" (self-knowing) consciousness, semantic memory is associated with "noetic" (knowing) consciousness. In the remember–know paradigm, the latter is expressed through self-reports of knowing that an event occurred but being unable to mentally re-experience it (Umanath and Coane, 2020). The future-oriented equivalent of knowing an event occurred despite being unable to recollect it would seem to be expecting an event to occur despite being unable to imagine it. However, whereas recognized-but-not-recollected events are commonplace, expected-but-unimaginable ones are not: if prompted, we are able to imagine even unfamiliar and unlikely events, let alone ones we expect to happen (Sasse et al., 2015; Szpunar, 2010). Atance and O'Neill (2001) instead define semantic future thought as semantic knowledge about the future. In particular, self-specificity is central to their distinction between episodic and semantic future thought: normative, script-like representations of future events, while spatiotemporally specific, are semantic to the extent that they do not account for self-specific constraints. For example, a script-like representation of a meal at a restaurant might be semantic in that it does not account for self-specific information about one's food allergies.

The isomorphism between semantic memory and semantic future thinking (SFT) is therefore not perfect, because there is no exact future analogue to recognized-but-not-recollected semantic event memories. Furthermore, whereas semantic memory can be atemporal (for example, remembering the Pythagorean theorem), SFT definitionally has a temporal orientation. SFT is therefore perhaps best thought of not as the future equivalent of semantic memory but as the subset of semantic knowledge about the future that is non-self-specific. Regardless, what is common to the episodic–semantic distinctions in memory and future thought is the criterion of autonoetic consciousness, measured in the EFT literature through self-reports of a subjective sense of "pre-experiencing" (i.e., self-reports of mentally simulating the experience of) an imagined future event (D'Argembeau and Van der Linden, 2012; Rasmussen and Berntsen, 2014; Duval et al., 2012).

One recent study found that EFT has a greater effect on delay discounting than SFT, defined, as above, as semantic knowledge about the future (Chiou and Wu, 2017). Similarly, Rung and Madden (2019) found that imagining a future time (but not a particular future event in one's personal future) had no effect on delay discounting. If autonoetic consciousness is what differentiates EFT from SFT, it may also be responsible for the greater effect of EFT on delay discounting. We now turn to this possibility and explore alternatives.

4.1.3 Possible roles of autonoetic consciousness in the EFT cue effect

Among formal models of decision making from reinforcement learning, "model-based" algorithms are those that use "simulated experience" to guide decisions (Sutton and Barto, 2018), a process that is analogous to episodic future thinking (Kinley et al., 2022). Indeed, Benoit et al. (2011) argue that EFT, elicited by cues presented during intertemporal choice, reduces delay discounting by simulating the experience of the undiscounted value of a future reward in the present. This language of simulated experience implies that the cue effect will occur only for future events that are mentally "pre-experienced", i.e., that evoke autonoetic consciousness. On this view, autonoetic consciousness is essential to the effect of EFT on delay discounting.

A related alternative to this proposal is that the effect of EFT on delay discounting could be explained not by autonoetic consciousness *per se* but by some associated characteristic such as vivid imagery or first-person perspective. These could amplify known construal level-related effects on delay discounting: manipulations that decrease the psychological distance of distal rewards by reducing their construal level (i.e., the degree of abstractness with which they are represented; Trope and Liberman, 2010) tend to reduce delay discounting (Kim et al., 2013). One such manipulation is to provide many details of a future reward (Kim et al., 2013; Lempert and Phelps, 2016), and vivid and detailed imagery of a future event could have a similar effect. Likewise, participants visualize future events from a range of visual perspectives (Kinley et al., 2021), and first-person perspective is associated with lower construal level (Tausen et al., 2020). Imagining a future event from a first-person perspective could reduce the construal level of a reward that would be obtained at that time, thereby attenuating discounting.

Vivid imagery could also amplify known attentional effects on delay discounting. A cue-evoked vivid mental image might orient attention toward the future, increasing awareness of the opportunity cost of choosing the immediate reward. Such a shift in temporal attention is thought to underlie the "explicit-zero" effect, in which delay discounting is reduced through reminders the null outcome associated with each choice, e.g., "$50 now and $0 in 3 months or $0 now and $100 in 3 months" (Radu et al., 2011). Even if a vivid mental image served only as a distraction, this could still attenuate discounting: distracted participants may be more likely to rely on heuristic decision-making which, in the context of intertemporal choice, means attending to reward magnitude more than reward delay (Ebert, 2001; Ebert and Prelec, 2007; Amasino et al., 2019) and thus discounting future rewards less.

Autonoetic consciousness may be essential to the effect of EFT on delay discounting, or it may be incidental, and the EFT cue effect may be reducible to known construal level or attentional effects. The present work was meant to explore these possibilities. We aimed to first investigate whether episodic future thinking has a greater effect on delay discounting than semantic future thinking, and then to investigate whether any difference is accounted for by autonoetic consciousness or related characteristics of mental imagery such as vividness or visual perspective.

4.2 Experiment 1

4.2.1 Materials and methods

Participants (N = 172, 143 women, ages 17 – 25, median = 18) were recruited from McMaster University's undergraduate population and completed the experiment online for credit in a psychology course, providing consent as approved by the local research ethics board. The experiment was programmed using jsPsych (De Leeuw, 2015) and hosted on a web server that participants accessed from browsers on their own computers.

Future thinking writing task

Participants were assigned by random number generator to an episodic or semantic future thinking condition with equal probability. 93 participants were assigned to the EFT condition and 76 were assigned to the SFT condition. For the purpose of providing cues during a delay discounting task to future events whose delay would match the delay of the distal reward, participants in both groups were asked to think of a future event that would occur approximately 90 days from the date of their participation.

Participants in the EFT condition were given the following instructions:

> Please think of an event that might happen to you on or around [*date 90 days in the future, dynamically generated*]. This should be an event that is specific to you. It should happen on a specific date and in a specific place. It should also not be something that you have experienced many times. You may use your calendar to come up with something. If you do not have anything planned, you can imagine something that could realistically happen.

They were given the following examples and explanations of events ill-suited to the study:

- A weekend music festival (does not happen on one specific date)
- A national holiday (not specific to you, does not happen in a specific place)
- Going grocery shopping (something you have probably experienced many times)

Finally, they were given the following examples of events well-suited to the study:

- Running into an old friend
- Moving into a new apartment

- Meeting up with someone to buy a piece of furniture

Participants in the SFT group received a definition of their task that essentially mirrored that given to the EFT group. The SFT group's instructions were the following:

> Please think of an event that might happen on or around [*date 90 days in the future, dynamically generated*]. This should be an event that is not specific to you. It should happen on a specific date but does not need to happen in a specific place. You may search the web to come up with something.

The SFT group was given the following examples of events ill-suited to the study:

- A weekend music festival (does not happen on one specific date)
- Running into an old friend (specific to you)

Finally, they were given the following examples of events well-suited to the study:

- A national holiday
- A sports event
- An election

Participants provided a title for their events and then wrote a detailed description of them in a free text field. The goal of the manipulation was for the conditions to be as similar as possible during the delay discounting task (both groups would see cues to future events with delays matching the delay of a distal reward, as described below) while still being differentially likely to evoke autonoetic consciousness: the episodic group was expected to actively imagine an event in their personal future while the semantic group was expected to merely draw on semantic knowledge that an event would occur. As explained below, we used a manipulation check to verify this.

Imagery characteristics questionnaire

After completing the writing task, participants responded via digital visual analog scale to a series of questions querying various characteristics of their mental imagery. The visual analog scales had written anchors at either side and numerical anchors at every 10% increment. A digital slider began in the middle of the range (50%) and could be placed at locations from 0% to 100% in increments of 1%.

The questions were the following:

- Did you have a mental image of the event? (with anchors "No image at all" and "Image as clear and vivid as real life"; Andrade et al., 2014)

- Was the general emotional tone of the event positive or negative? (with anchors "Negative" and "Positive"; Johnson et al. 1988)
- Were the emotions associated with the event intense? (with anchors "Not intense" and "Very intense"; Johnson et al. 1988)
- What percentage of the time did you see the scene from a first-person perspective? (with anchors "0%" and "100%")
- When you imagined the event, did it feel like you were "pre-experiencing" it? (with anchors "Not at all" and "Very much"; Johnson et al. 1988; D'Argembeau and Van der Linden 2012)

Delay discounting task

Participants completed a computerized delay discounting task following the adjusting-amount procedure for monetary amounts starting at $400 and $800 with 5 trials per delay (Koffarnus and Bickel, 2014). Delays were 1 week, 1 month, 3 months, 6 months, and 1 year (displayed to participants in units of days). Participants completed one block of trials at a time, with each block corresponding to a single delay. Blocks were presented in order of increasing delay.

A participant's "indifference point" at a given delay is the value of an immediately available reward expressed as a proportion of a larger reward available at that delay, such that the two are of equal subjective value. For example, if a participant's indifference point at 5 months is 0.6, they would be indifferent between $60 available immediately and $100 available in 5 months.

The adjusting amount procedure computes indifference points as follows: when the immediate reward is chosen, its value is adjusted downward on the next trial and vice versa, while the value of the delayed reward is held constant. The amount by which the immediate reward is adjusted begins at one quarter of the delayed reward value and is halved with each subsequent trial. A participant's indifference point at a given delay was calculated as the relative value to which the immediate reward would be adjusted based on the final trial at that delay.

On the block of trials where the delay was 90 days (i.e., the temporal distance of the future event from the writing task), the participant's chosen title for their event was displayed beneath the delayed option (Peters and Büchel, 2010). No cues were displayed during any other block of trials. Participants were given no instructions about what to do when the cues appeared, instead simply being told to select which of the two given options they would prefer on each trial.

To compute the effect of this cuing on delay discounting, baseline rates of discounting were first determined. Following Franck et al. (2015), 5 candidate functions were fit to each participant's indifference points:

$$
\begin{aligned}
I_D &= \frac{1}{1+kD} \quad \text{(Mazur, 1987)} \\
I_D &= e^{-kD} \quad \text{(Samuelson, 1937)} \\
I_D &= \frac{1}{(1+kD)^s} \quad \text{(Myerson and Green, 1995)} \\
I_D &= \frac{1}{1+kD^s} \quad \text{(Rachlin, 2006)} \\
I_D &= \beta\delta^D \quad \text{(Laibson, 1997)}
\end{aligned}
$$

where I_D is the individual's indifference point at delay D. Each model was fit using R's `nls` function (R Core Team, 2018) and the model with the lowest value of the Bayesian Information Criterion (BIC; Schwarz, 1978) was selected. This approach was taken because we were interested only in accurately characterizing each participant's baseline discounting and were theoretically noncommittal as to the specific functional form this discounting took. The value of the BIC is, all else being equal, lower for models that match the data closely and higher for models with more parameters, and thus can be used to select models that strike an appropriate balance between goodness of fit and parsimony.

The cuing effect Δ_I on delay discounting was computed by subtracting the actual cued indifference point, derived from the titrating procedure in the 90-day delay block, from the expected indifference point at 90 days, according to the model that best fit the indifference points at uncued delays. E.g.:

$$
\begin{aligned}
\Delta_I &= I_{90} - \hat{I}_{90} \\
&= I_{90} - \frac{1}{1+k\times 90}
\end{aligned}
$$

assuming that hyperbolic discounting (Mazur, 1987) best fit the participant's uncued indifference points. Alternatively, Δ_I would be computed as $I_{90} - \beta\delta^{90}$ if the discount function of Laibson (1997) provided the best fit. A larger value of Δ_I indicates a positive cuing effect, i.e., that the future thinking cues attenuated temporal discounting. This approach borrows elements from existing research: measures of single indifference points (like I_{90} in the present work) have been shown to be sufficient to replicate known effects in the intertemporal choice literature such as greater discounting of gains vs losses (Hardisty and Weber, 2009) and a correlation between discount rates and future self continuity (Sokol and Serper, 2019). Similarly, difference scores between measures of baseline and cued discounting (like Δ_I in the present work) have previously been used to measure the behavioural effects of episodic future thinking (Benoit et al., 2011; Peters and Büchel, 2010).

Manual data rejection

Raters (authors IK, MP, & YL) provided binary ratings of the written responses as having properly followed the instructions or not. In particular, participants in the EFT group were expected to use personal pronouns ("I", "me", "my") when describing their imagined future events while participants in the semantic future thinking group were not. Data were rejected on the basis of majority rule from the 3 raters. While the raters were necessarily aware of the experimental condition when providing ratings, they did not have access to any dependent measures (e.g., delay discounting data, imagery ratings).

Identifying non-systematic delay discounting data

Following Johnson and Bickel (2008), we identified delay discounting data as "non-systematic" when either (1) any indifference point was greater than the previous (excluding the indifference point from cued trials) by more than 0.2 (i.e., non-monotonicity) or (2) the final indifference point was not lower than the first by at least 0.1 (i.e., delay insensitivity).

4.2.2 Results

Data rejection

The rates of agreement between the 3 raters were as follows: IK – MP: 93.02%; IK – YL: 92.44%; MP – YL: 95.93%. A majority rule system rejected data from 20 out of 172 participants (11.6%). Delay discounting data from 53 of the remaining 152 participants (34.9%) were identified as non-systematic. However, only 6 of these were non-systematic according to criterion 1 (non-monotonicity), meaning that 47 met only criterion 2 (delay insensitivity). This criterion was originally developed using a delay discounting task with a maximum delay of 25 years (Johnson and Bickel, 2008) compared to the current study's 1 year. Given that insensitivity to a 1-year delay does not indicate non-systematic discounting as strongly as insensitivity to a 25-year delay, only data meeting criterion 1 were removed. This left a final sample size of 146 (121 women, ages 18 – 25, median = 18). 83 participants were assigned to the EFT condition and 63 were assigned to the SFT condition.

Imagery questionnaire

Table 4.1 summarizes the correlation structure of the brief imagery characteristics questionnaire. There were particularly strong correlations between autonoetic consciousness and both imagery vividness and first-person perspective.

Table 4.2 summarizes the mean responses to each item in the phenomenological questionnaire. Participants in the EFT group had significantly higher word counts (EFT mean: 64.10; SFT mean: 30.19; Kruskall-Wallis $\chi^2(1) = 22.06$, $p < 0.0001$), self-reported autonoetic consciousness (sense of "pre-experiencing", $t(144) = 3.61$, $p = 0.0004$) and imagery vividness ($t(144) = 2.93$, $p = 0.0039$). The groups did not differ significantly in self-reported first-person perspective ($t(144) = 1.52$, $p = 0.13$), emotional intensity ($t(144) = 0.68$, $p = 0.50$), or emotional valence ($t(144) = 0.56$, $p = 0.58$).

Variables	1.	2.	3.	4.	5.
1. Imagery vividness	–				
2. Emotional tone	0.18*	–			
3. Emotional intensity	0.23**	-0.13	–		
4. First-person perspective	0.23**	0.07	0.07	–	
5. Autonoetic consciousness	0.29***	0.18*	0.16*	0.035***	–

Table 4.1: Bivariate correlations between items on the imagery characteristics questionnaire. Each r value is computed on 150 degrees of freedom. *p < .05; **p < .01; ***p < .001.

Variable	EFT group	SFT group	p
Imagery vividness	82.2 (16.6)	72.9 (21.4)	0.0039
Emotional tone	80.6 (24.7)	78.2 (24.7)	0.58
Emotional intensity	57.2 (26.9)	54.2 (26.1)	0.50
First-person perspective	64.8 (31.7)	56.7 (31.9)	0.13
Autonoetic consciousness	71.6 (25)	55.4 (29.1)	0.0004

Table 4.2: Means and standard deviations for each item in the phenomenological questionnaire, for each experimental conditions. Each p value is the result of a t test with 144 degrees of freedom.

Cue effects on delay discounting

The median R^2 value of the best-fitting discounting curves was 0.90, indicating a degree of model fit comparable to previous research (Miglin et al., 2017). There was no difference between conditions in the cue effect on delay discounting (Δ_I; t(144) = -0.86, p = 0.39; Fig. 4.1). However, visually inspecting the distributions of the cue effect shows that both conditions contained outliers which may have obscured the overall trend. Upon removing the boxplot outliers within each condition (16 total, 11% of the sample), a significant difference between groups appeared (t(128) = -2.08, p = 0.0396; Fig. 4.2).

We next tested whether this difference in cue effects between groups could be explained by the difference in self-reported autonoetic consciousness. A mediation analysis was performed using structural equation modelling as implemented in the `lavaan` package Rosseel (2012) to test whether autonoetic consciousness ratings mediated the effect of experimental condition on Δ_I. While the total and direct effects were significant (p = 0.0361 and p = 0.0168, respectively), the indirect effect was not ($\hat{ab}$ = -0.006, p = 0.20). Moreover, even a simple bivariate correlation between the magnitude of the cue effect and self-reported autonoetic consciousness was not significant (r(128) = -0.07, p = 0.401). Thus we found no evidence that autonoetic consciousness explained the effect of EFT on delay discounting.

As outlined in the introduction, it is also possible that the cue effect is explained by vividness or first-person perspective. To explore this possibility, we ran the same mediation analysis as above using these self-report variables rather than autonoetic consciousness. Neither first-person perspective nor vividness were found to be significant moderators (p = 0.41 and p = 0.13, respectively). Thus we also found no evidence that either vividness or first-person perspective explained the effect of EFT on delay discounting.

4.2.3 Interim discussion

The EFT group exhibited both greater autonoetic consciousness and a larger cue effect, as expected, but the difference in cue effect magnitudes was not explained by the difference in autonoetic consciousness. However, some caution may be warranted when interpreting these results, as the basic finding of a group difference in the magnitude of the cue effect relied on post-hoc outlier exclusion. The relatively brief delay discounting task used may have produced noisy estimates of indifference points and thus noisy estimates of the magnitude of the cue effect. This could have concealed group differences in the cue effect prior to outlier exclusion and perhaps concealed mediators of the cue effect even after outlier exclusion. Thus the null results in experiment 1 may have been due to noisy estimates of indifference points rather than the genuine absence of any mediator of the cue effect.

Much research on delay discounting assumes that intertemporal choices are stochastic rather than strictly rule-based. For example, the probability of choosing the immediate reward can be modelled using logistic regression, where the indifference point is defined as the relative immediate reward value at which this probability is 0.5 (Wileyto et al., 2004). Furthermore, some degree of non-monotonicity is generally tolerated in the empirical discounting function, implying that a certain amount of decision noise is unavoidable (Johnson and Bickel, 2008).

In general, estimates of k produced using a titrating procedure, as in experiment 1, are strongly correlated with those produced using other methods (Koffarnus and Bickel, 2014). However, the present experimental design relies on accurate measurements in only 5 trials of the single indifference point corresponding the the cued delay (rather than a k value, which generalizes across multiple indifference points; the present design was chosen in order to assess event-by-event associations between imagery characteristics and decision-making behaviour), and it seems possible that the titrating method could be more vulnerable to decision noise when estimating single indifference points. For example, if a participant selects an immediate reward with relative value of 0.75, under the titrating method the measured indifference point at that delay is constrained to be between 0.5 and 0.75. However, if the participant's "true" indifference point (conceived of probabilistically) were 0.8, there could still be a non-zero chance of choosing an immediate option with a relative value of 0.75. Then, for this participant, the titrating method would underestimate the true indifference point.

This problem would be more pronounced for participants with more stochastic decision behaviour—i.e., participants whose preference for the immediate reward changes less sharply around their indifference point. To validate these intuitions, a simulation was run to compare the titrating procedure to a logistic regression-based method of measuring indifference points.

4.3 Experiment 2

4.3.1 Materials and methods

Two model decision-making agents were created, both of whose decision behaviour was governed by the following sigmoid function:

$$p_{\text{imm}} = \frac{1}{1 + e^{b(I-v)}}$$

where p_{imm} is the probability of choosing the immediate reward, v is the value of the immediate reward relative to the value of the delayed reward, I is the indifference point (i.e., the value of v at which $p_{\text{imm}} = 0.5$), and b is a free parameter controlling the steepness of the curve around the indifference point. For a "minimally stochastic" agent, the value of b was set to 100, while for a "moderately stochastic" agent b was set to 20. The resulting functions are plotted in Fig. 4.3.

Two methods were used to estimate indifference points. The first was the titrating procedure, which proceeded exactly as in experiment 1, with the simulated agent being presented with an adaptive series of 5 choices. The second method presented the simulated agent with 25 choices in which the relative value of the immediate reward took on linearly spaced values from 0 to 1, inclusive of endpoints (i.e., 0, 0.0417, 0.0833, ..., 1), and computed the indifference point from the best fitting logistic curve.

During a single simulation, an indifference point was drawn at random from a uniform distribution between 0.25 and 0.75 (these endpoints were chosen so that p_{imm} was always nearly 1 for $v = 1$ and nearly 0 for $v = 0$). Both methods were then used to estimate this indifference point for a given decision agent and the errors in these estimates were collected. This was repeated 1000 times for each agent type, resulting in 1000 error measurements for each conjunction of agent type and measurement method.

4.3.2 Results

The absolute errors for each agent type and estimation method are shown in Fig. 4.4. For the minimally stochastic agent, the titrating method produced smaller (square root-transformed absolute) errors than the logistic regression method ($t(1998) = -2.7171, p < 0.01$), while for the logistic regression method this pattern was reversed ($t(1998) = 9.9545, p < 0.0001$.

4.3.3 Interim discussion

This simulation illustrates that under conditions of moderate stochasticity, the more time-consuming method of estimating indifference points using logistic regression outperforms the titrating method. While the titrating method is a better estimator for minimally stochastic decisions, the size of its advantage is negligible (Cohen's $d = -0.1215$).

It could be argued that participants exhibiting significantly non-deterministic choice behaviour ought to be excluded. Even if this is the case, such non-deterministic behaviour is only detectable using the logistic regression method which, by sampling a series of relative immediate reward values, gives participants a chance to make inconsistent decisions (for example, choosing an immediate option with a relative value of 0.75 but not choosing an immediate option with a relative value of 0.8).

Thus, for the present study, the logistic regression method of estimating indifference points is clearly advantageous: it allows us to assess the stochasticity of participants' decisions and to potentially exclude non-deterministic decision makers. If we choose not to exclude these participants, it offers us a more statistically efficient method of estimating their indifference points. In either case, the result should be lower-error estimates of indifference points, increasing our power to detect correlates of the cue effect.

4.4 Experiment 3

On the basis of the simulation described above, we repeated experiment 1, this time estimating indifference points using logistic regression rather than the titrating method.

4.4.1 Methods

Participants (N = 279, 199 women, 1 other/prefer not to say, ages 17 – 34, median = 19) were recruited from McMaster University's undergraduate population and completed the experiment online for credit in a psychology course, giving consent as approved by the local research ethics board. The future thinking writing task and imagery characteristics questionnaire were identical to those in experiment 1.

Delay discounting task

As in experiment 1, participants made a series of choices between immediate rewards and rewards at delays of 1 week, 1 month, 3 months, 6 months, and 1 year (displayed to the participant in units of days). The amount of the delayed reward was an integer amount randomly selected from a uniform distribution between $80 and $100, inclusive. The amount of the delayed reward was randomized in this way to avoid having participants habituate to a fixed delayed reward amount. The value of the immediate reward took on 25 linearly spaced proportions of the delayed reward, from 0% to 100% (inclusive) rounded to the nearest integer dollar value, at each delay. The order in which the delays and immediate reward values were presented was randomized for each participant using jsPsych's `randomization.shuffle` function (De Leeuw, 2015). At the 90-day delay, the title of a participant's imagined future event was displayed under the delayed reward option.

Indifference points were computed by fitting logistic regression models using R's `glm` function (R Core Team, 2018) to the binary choice data at each delay and solving for the points at which the best-fitting logistic functions produced a value

Variables	1.	2.	3.	4.	5.
1. Imagery vividness	–				
2. Emotional tone	0.07	–			
3. Emotional intensity	0.19**	-0.01	–		
4. First-person perspective	0.32***	-0.06	0.05	–	
5. Autonoetic consciousness	0.47***	0.00	0.22***	0.31***	–

Table 4.3: Bivariate correlations between items on the imagery characteristics questionnaire. Each r value is computed on 220 degrees of freedom. *p < .05; **p < .01; ***p < .001.

of 0.5. As in experiment 1, for each participant, 5 candidate models were fit to the indifference points from uncued trials and the best fitting of these (i.e., the one with the lowest BIC value) was used to calculate $\hat{I}_{90}$, the predicted uncued indifference point at a delay of 90 days. This predicted indifference point was subtracted from the actual indifference point at the cued delay of 90 days to compute Δ_I, the cue effect on delay discounting.

4.4.2 Results

Data rejection

The rates of agreement between the 3 raters were as follows: IK – MP: 81.72%; IK – YL: 83.51%; MP – YL: 91.76%. A majority rule system resulted in the rejection of data from 32 out of 279 (11.5%) participants. As in experiment 1, data from participants identified as non-systemic discounters according to criterion 1 (non-monotonicity; Johnson and Bickel, 2008) were removed from the sample (25 out of the remaining 247, or 10.1%). This left a final sample size of 222 (199 women, 1 other/prefer not to say, ages 18 – 34, median = 19). 139 participants were assigned to the EFT condition and 83 were assigned to the SFT condition.

Questionnaire

Table 4.3 summarizes the correlation structure of the brief imagery characteristics questionnaire. As in experiment 1, autonoetic consciousness correlated strongly with both imagery vividness and first-person perspective.

Table 4.4 summarizes the mean responses to each item in the phenomenological questionnaire. As in experiment 1, participants in the EFT group had higher word counts (EFT mean: 56.61; SFT mean: 34.02; Kruskall-Wallis $\chi^2(1) = 14.39$, p = 0.0001). Participants in the EFT group also exhibited higher ratings of autonoetic consciousness ($t(220) = 2.63$, p = 0.0092), imagery vividness ($t(220) = 4.26$, p < 0.0001), and first-person perspective ($t(220) = 2.45$, p = 0.0150), but not emotional intensity ($t(220) = 0.50$, p = 0.62) or emotional valence ($t(220) = 1.52$, p = 0.13).

Variable	EFT group	SFT group	p
Imagery vividness	82.4 (17)	69.8 (27)	< 0.0001
Emotional tone	85 (20.4)	80.3 (25.2)	0.13
Emotional intensity	55.6 (28.8)	53.6 (29.5)	0.62
First-person perspective	66.8 (30.8)	55.8 (35.1)	0.0150
Autonoetic consciousness	68.4 (26.7)	58.2 (29.5)	0.0092

Table 4.4: Means and standard deviations for each item in the phenomenological questionnaire, for each experimental condition. Each p value is the result of a t test with 220 degrees of freedom.

Delay discounting and cue effect

As in experiment 1, the hyperbolic discounting curves showed comparable goodness-of-fit to previous research, with a median R^2 value of 0.94 (Miglin et al., 2017). The stochasticity of participants' decision making was assessed according to the presence of a clear border dividing relative immediate reward values for which the delayed reward was chosen from values for which the immediate reward was chosen. I.e., if there was a relative immediate reward value above which participants always chose the immediate reward, this decision making was labelled non-stochastic. If no such border existed, the decision making was labelled stochastic.

Most participants (200, or 90%) displayed stochastic decision making, defined in this way, for at least one of the five delays at which indifference points were assessed, and the modal participant (56, or 25%) displayed stochastic decision making for four out of five delays. Given the pervasiveness of stochastic decision making, we opted not to use it as a basis for data exclusion. This pervasiveness also justifies the use of logistic regression to estimate indifference points since, as demonstrated in experiment 2, this method is more suitable than titration for stochastic decision making.

Indeed, in contrast to experiment 1, participants in the EFT group exhibited a larger cue effect prior to any outlier exclusion (t(220) = -2.00, p = 0.0469; Fig. 4.5). Nonetheless, there were still clear outliers. Boxplot outliers were again removed (29 total, 13% of the sample) to enable parametric methods as used in experiment 1. Upon removing outliers, the difference between groups became more pronounced (t(191) = -3.34, p = 0.0010; Fig. 4.6).

We again performed a mediation analysis using structural equation modelling as implemented by `lavaan` (Rosseel, 2012) to determine whether the difference between groups in the magnitude of the cue effect could be explained by the difference in self-reported autonoetic consciousness. As in experiment 1, the effect of experimental condition on Δ_I was not significantly mediated by autonoetic consciousness (indirect effect $\hat{ab}$ = -0.003, p = 0.35). Also, as in experiment 1, the bivariate correlation between autonoetic consciousness and Δ_I failed to reach significance (r(191) = -0.018, p = 0.809). Thus, as in experiment 1, we found no evidence that autonoetic consciousness explained the effect of EFT on delay discounting.

We again explored the possibility that the cue effect is explained by first-person perspective or vividness. In another pair of mediation analyses, neither first-person perspective (p = 0.74) nor vividness (0.72) significantly mediated the effect of experimental condition on Δ_I. Thus, as in experiment 1, we found no evidence that first-person perspective or vividness explained the effect of EFT on delay discounting.

4.4.3 Interim discussion

The statistical power of experiment 3 was higher than that of experiment 1, owing to both a considerably larger sample size and a method of estimating the cue effect that, as demonstrated in experiment 2, is more tolerant of stochastic decision making. Indeed, unlike in experiment 1, the basic finding of a difference between groups in the cue effect existed prior to outlier exclusion in experiment 3. Nonetheless, the results of experiment 3 were similar to those of experiment 1, lending confidence that these results are robust.

4.5 General discussion

Both experiments showed the same pattern of results: the EFT group exhibited both greater self-reported autonoetic consciousness and a greater cue effect than the SFT group, however, the difference between groups in the cue effect was not accounted for by the difference in autonoetic consciousness. Thus, the present findings do not support the assertion that autonoetic consciousness explains the effect of EFT on delay discounting. We also did not find evidence that either vividness or first-person perspective, two other factors that differentiated episodic from semantic future thinking, explained the difference in cue effect between groups. Thus the present findings also do not support the claim that EFT-based reductions of delay discounting are reducible to known attentional or construal level-based effects.

An important potential limitation of these experiments concerns the survey item used to measure autonoetic consciousness: despite its ubiquity in the literature, the term “pre-experience” is rather strong language for describing normal prospection, and is perhaps even suggestive of the vivid and intrusive “flash-forwards” seen in some psychiatric conditions (Engelhard et al., 2011; Holmes et al., 2007). Participants in the EFT group may have been unsure whether their mental imagery met the threshold to be described as “pre-experiencing”. Nonetheless, there is reason to believe the item was a valid measure of autonoetic consciousness. First, responses to it were correlated with responses to the items measuring vividness and first-person perspective (Tables 4.1 and 4.3), both of which are closely linked to autonoetic consciousness in memory (van Schie et al., 2019; Zaman and Russell, 2022). There is a similarly close link in EFT between vividness, first-person perspective, and autonoetic consciousness: when a mental simulation of the future is vivid (i.e., similar to actual experience; Tooming and Miyazono, 2021) and adopts the perspective from which an event would actually be experienced, this should enhance the sense that one is mentally “pre-experiencing” that event. Indeed, previous research has found that vividness and first-person perspective predict autonoetic consciousness, measured, as in the present work, as the self-reported feeling of "pre-experiencing" imagined future events (D'Argembeau and Van der Linden, 2012). Moreover, autonoetic consciousness is what differentiates episodic from semantic future thinking (Atance and O'Neill, 2001), and responses to the "pre-experiencing" item in the present work differed significantly between the EFT and SFT groups. Along with the aforementioned pattern of inter-item correlations, this suggest that our measure of autonoetic consciousness was valid.

Another potential limitation arises due to the population being studied. Participants were undergraduate students mostly in "emerging adulthood", a period during which executive function is still developing (De Luca et al., 2003). EFT draws on executive function (Cole et al., 2013), which raises questions about how well the present results can be expected to generalize. For example, it could be that after emerging adulthood, more fully developed executive function enables a high degree of autonoetic consciousness during EFT, which enhances the effect of EFT on delay discounting. However, this seems unlikely—a recent study found no difference between adults (with an average age of 34) and adolescents (ages 12–14) in the degree to which EFT cues influenced either delay discounting or other factors such as time perception Burns et al. (2022). This suggests that the effect of the cue paradigm used in the present work does not fundamentally change even between adolescence and adulthood, let alone between emerging and early/middle adulthood, and lends confidence in the generality of the present results. An apparent exception here may be older adulthood: among older adults, the EFT cue effect on delay discounting is correlated with cognitive control (Sasse et al., 2017), and general cognitive decline is also a likely explanation for diminished autonoetic consciousness during episodic future thinking (Rendell et al., 2012). Had older adults been participants in the current paradigm, autonoetic consciousness may have been found to mediate the cue effect. However, this would be due to underlying variability in general cognitive function, a confounding factor correlated with both the cue effect and autonoetic consciousness, rather than to a fundamentally different process underlying the cue effect that was not found in the present results. Thus, despite the inherent limitations in using an undergraduate study population, the present results are likely to generalize well across age groups.

It is important to note that, in the present studies, participants in the EFT condition imagined future *events* but did not imagine receiving future *rewards*. As discussed in the introduction, Benoit et al. (2011) theorize that EFT reduces delay discounting by simulating undiscounted future reward values, i.e., through the simulated experience of future rewards. Perhaps when participants are asked to imagine experiencing future rewards, their degree of pre-experiencing (i.e., of autonoetic consciousness) determines the degree to which they value said rewards and are willing to wait for them. The present studies cannot rule this possibility out and, indeed, it seems quite plausible. However, this pre-experienced-future-value model cannot be an exhaustive account of the role of EFT in intertemporal choice—the fact remains that the present studies found an effect of EFT on delay discounting that was not mediated by autonoetic consciousness. What, then, explains this effect?

Operationally, the experimental groups were defined in terms of the self-relevance of the future events in question. This manipulation succeeded in its intended purpose of inducing a difference in episodicity, operationalized as autonoetic consciousness ratings, but self-relevance and episodicity are not synonymous; rather, self-relevance is a necessary, but not sufficient, condition for episodicity (Perrin and Rousset, 2014). That is, semantic memory (and, by extension, semantic future thought) can have self-referential content (Klein, 2013). An intriguing possibility is that the difference between our experimental groups in self-referential content, rather than episodicity *per se*, accounts for the difference in cue effects. Indeed, Burns et al. (2022) found that imagining events in other people's futures does not reduce delay discounting, despite being identical to a more efficacious EFT condition except with respect to the self-relevance of the imagined events. We now explore how self-relevance might have been critical to the effect of EFT on delay discounting found in the current studies.

The philosopher Derek Parfit (Parfit, 1982, 1984) famously argued that, when one is not psychologically connected to one's future self, it is not irrational to act preferentially in the interests of the present self. Indeed, there is a great deal of empirical evidence for an association between delay discounting and a sense of disconnection from the future self (i.e., a lack of "future self-continuity"; Hershfield, 2011). Szpunar and Tulving (2011) argue that connectedness to one's future results from autonoetic consciousness, which Tulving (1985) similarly described as mediating the awareness of one's existence as extending from the past into the future (i.e., "diachronicity"; Klein, 2014). Thus, Hoerl and McCormack (2016) suggest that EFT may reduce delay discounting by increasing the sense of connection to one's future self. However, recent evidence has not borne this out. For example, Burns et al. (2022) did not find that EFT increases future self-continuity, and McCue et al. (2019) found that the two factors independently predict reduced delay discounting. Moreover, individuals with episodic amnesia exhibit severely impaired EFT abilities but do not exhibit steeper delay discounting compared to healthy controls (Kwan et al., 2012, 2013). Craver et al. (2014) argue these cases demonstrate that autonoetic consciousness is not a prerequisite for having a temporally extended sense of self and self-interest.

A more nuanced conception of the relationship between autonoesis and self-continuity was proposed by Prebble et al. (2013). They distinguish between the *phenomenological* continuity afforded by autonoetic consciousness and the *semantic* continuity that comes from using one's autobiographical knowledge to construct a personal chronology (i.e., a life story). It may be this latter semantic continuity that explains the effect of EFT on delay discounting. That is, thinking about an event in the personal future could activate representations of one's goals and broad life trajectory (representations that are non-episodic; Conway et al., 2019), thereby increasing semantic self-continuity and orienting decision making toward the future.

Although self-continuity as studied by Hershfield (2011) is not increased by EFT (Burns et al., 2022) and thus cannot mediate the effect of EFT on delay discounting, semantic continuity as described by Prebble et al. (2013) seems to be a different construct. Hershfield (2011) conceives of self-continuity as "overlap" with one's future self, and this is closely related to how similar one feels to their future self (Sokol and Serper, 2019). However, we do not always intend to be remain similar to our present selves indefinitely and we may even aspire to change in significant ways, such as by adopting new values (Callard, 2018). In calling to mind intended lives that potentially entail personal change, episodic future thinking may not make us feel more similar to our future selves but might nevertheless motivate us to act in our future selves' interests.

On this view, the critical difference between the EFT and SFT groups in the present study was the self-relevance of the imagined future events: only in the EFT group did participants imagine events that formed parts of their personal chronologies, which is the basis of semantic self-continuity (Prebble et al., 2013). Autonoetic consciousness may have accompanied this process, but only as an epiphenomenon; as discussed above, the representations reducing delay discounting would have been essentially semantic (Conway et al., 2019)[1]. Indeed, the cuing paradigm used in the present experiments can reduce delay discounting even among individuals with episodic amnesia (Kwan et al., 2015, but c.f. Palombo et al., 2015), who must be making use of semantic representations of their personal futures.

Thus, although we initially defined semantic future thinking as the self-*un*related subset of knowledge about the future (following Atance and O'Neill, 2001), we clearly must make conceptual room for representations of the future that are not episodic but are nonetheless self-specific (just as a distinction is drawn between episodic and semantic autobiographical memory; e.g., Levine et al., 2004). The present study suggests that these may be critical to the role of episodic future thinking in intertemporal choice: curbing delay discounting may be less about projecting the self into the future than about calling to mind the (semantic) knowledge that the self has a future.

In future work, we hope to positively test the role of self-relevance *per se* in the cue effect. This may be possible simply by adding existing survey measures of personal importance, goal-relevance (Lehner and D'Argembeau, 2016), and centrality (Özbek et al., 2020) to the present experimental design. In the absence of positive results, ideas about the role of self-relevance must remain speculation. Moreover, negative results cannot, strictly speaking, demonstrate the absence of an effect: for example, the between-participants design of the present experiments may have obscured a true effect of autonoetic consciousness, while a within-participants design (avoided here for concerns about demand characteristics) may offer a more powerful means of detecting mediation. This, among endless other alternative explanations for our null results, would have to be ruled out to definitively conclude that autonoetic consciousness does not mediate the cue effect. Nonetheless, we feel that the null results here are at least suggestive: we did not find even a simple bivariate correlation between autonoetic consciousness and the cue effect (which might be expected simply as an artifact of the experimental manipulation) despite large sample sizes, strong theoretical expectations (as outlined in the introduc-

[1]In a similar vein, McCarroll and Cosentino (2020) argue that third-person perspective in mental imagery of the future can increase connectedness to one's future self by activating semantic self-knowledge and thereby increase intertemporal patience.

tion), and an experimental design that was able to produce other positive results (differences in cue effect and self-reported phenomenology between the EFT and SFT groups). Thus, the present findings may point, in a tentative way, to the need to posit factors beyond autonoetic consciousness, such as self-relevance, to explain the role of episodic future thinking in intertemporal choice.

Figure captions

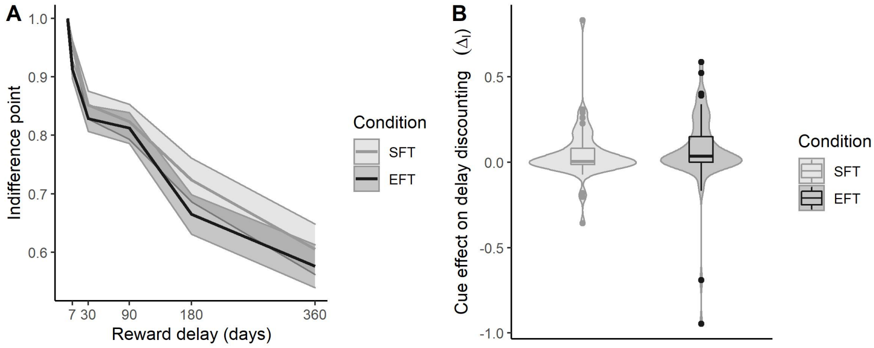

Figure 4.1: The effect of episodic and semantic future thinking on delay discounting. (A) Indifference points as a function of reward delay and condition. Shaded regions reflect 1 standard error of the mean. A deviation from the overall hyperbolic trend occurs for both groups at 90 days, which corresponds to the block of trials in which the cue word appeared. (B) The distribution of the cue effect (Δ_I) for each condition. The EFT group contained many outliers showing a negative cue effect and the SFT group contained many outliers showing a positive cue effect.

Bibliography

Addis, D. R. (2020). Mental time travel? a neurocognitive model of event simulation. *Review of Philosophy and Psychology*, 11(2):233–259.

Amasino, D. R., Sullivan, N. J., Kranton, R. E., and Huettel, S. A. (2019). Amount and time exert independent influences on intertemporal choice. *Nature human behaviour*, 3(4):383–392.

Amlung, M., Marsden, E., Holshausen, K., Morris, V., Patel, H., Vedelago, L., Naish, K. R., Reed, D. D., and McCabe, R. E. (2019). Delay discounting as a transdiagnostic process in psychiatric disorders: a meta-analysis. *JAMA psychiatry*, 76(11):1176–1186.

Andrade, J., May, J., Deeprose, C., Baugh, S.-J., and Ganis, G. (2014). Assessing vividness of mental imagery: the plymouth sensory imagery questionnaire. *British Journal of Psychology*, 105(4):547–563.

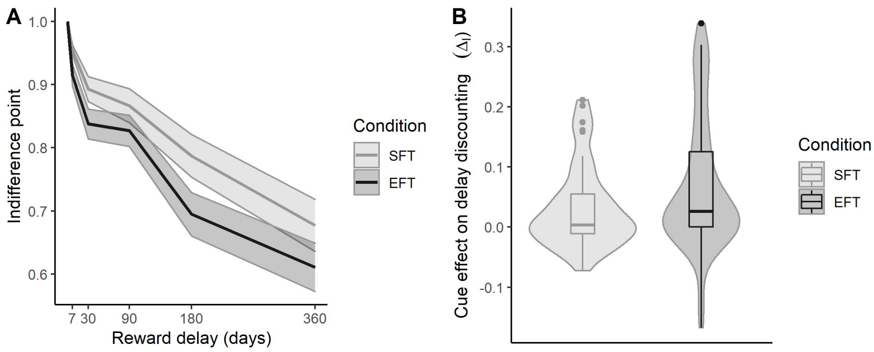

Figure 4.2: The effect of episodic and semantic future thinking on delay discounting after cue effect outlier removal. (A) Indifference points as a function of reward delay and condition. Shaded regions reflect 1 standard error of the mean. In the EFT group, a strong deviation from the overall hyperbolic trend occurs at 90 days, which corresponds to the block of trials in which the cue word appeared. (B) Distribution of cue effects for each group, after removing outliers.

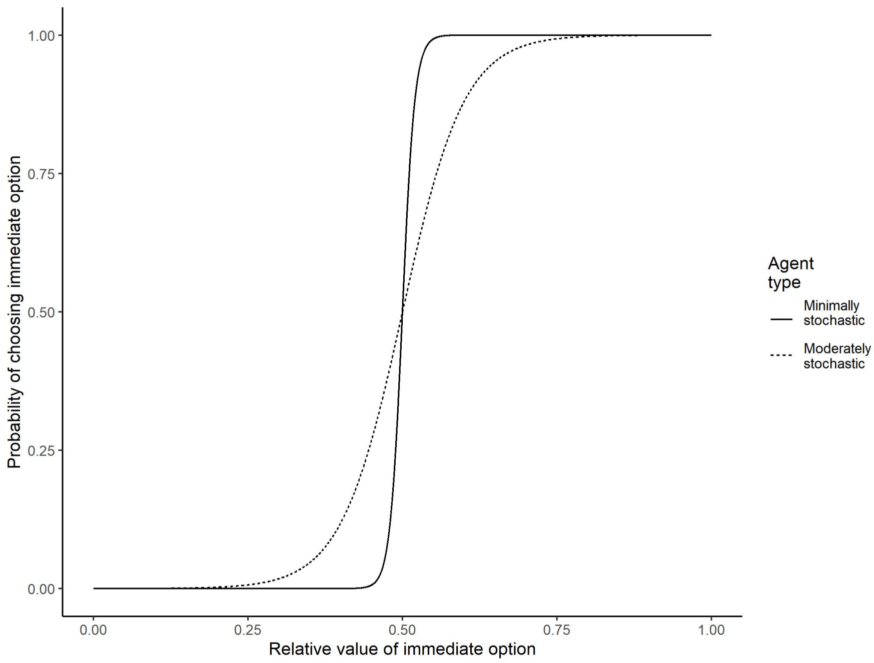

Figure 4.3: Functions generating the decisions of the two decision agents.

Atance, C. M. and O'Neill, D. K. (2001). Episodic future thinking. *Trends in cognitive sciences*, 5(12):533–539.

Benoit, R. G., Gilbert, S. J., and Burgess, P. W. (2011). A neural mechanism mediating the impact of episodic prospection on farsighted decisions. *Journal of Neuroscience*, 31(18):6771–6779.

Burns, P., Atance, C., O'Connor, A. P., and McCormack, T. (2022). The effects

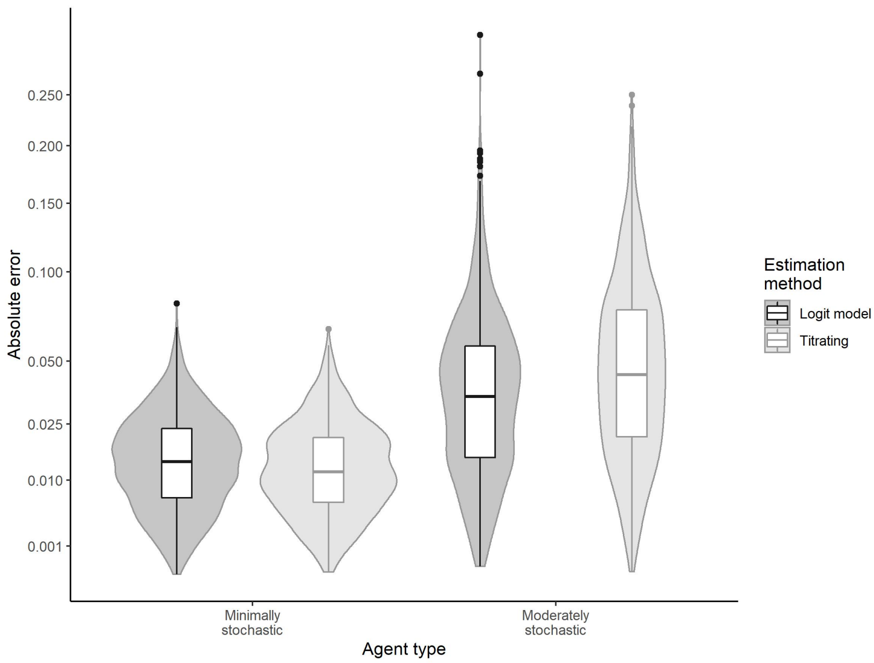

Figure 4.4: Absolute errors in indifference point estimates for each agent type and estimation method. Note that the y-axis is square root-transformed.

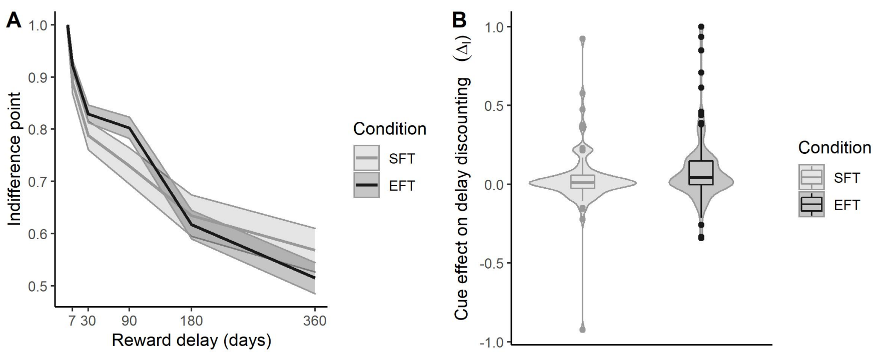

Figure 4.5: The effect of episodic and semantic future thinking on delay discounting. (A) Indifference points as a function of reward delay and condition. Shaded regions reflect 1 standard error of the mean. In the EFT group, a strong deviation from the overall hyperbolic trend occurs at 90 days, which corresponds to the block of trials in which the cue word appeared. (B) Distribution of cue effects for each group.

of cueing episodic future thinking on delay discounting in children, adolescents, and adults. *Cognition*, 218:104934.

Callard, A. (2018). *Aspiration: The agency of becoming*. Oxford University Press.

Chiou, W.-B. and Wu, W.-H. (2017). Episodic future thinking involving the non-

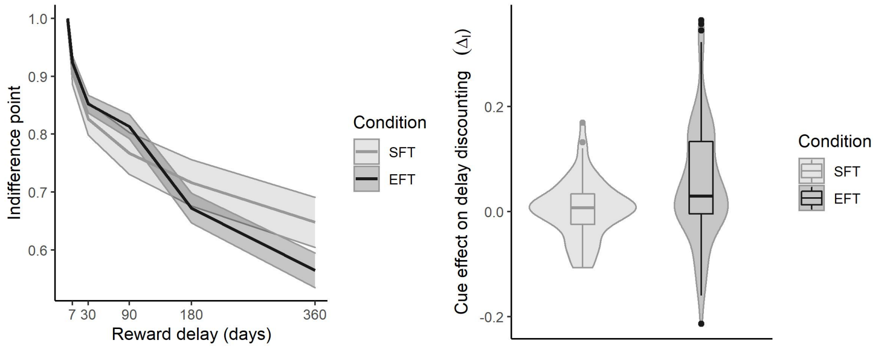

Figure 4.6: The effect of episodic and semantic future thinking on delay discounting after removing cue effect outliers. (A) Indifference points as a function of reward delay and condition. Shaded regions reflect 1 standard error of the mean. (B) Distribution of cue effects for each group.

smoking self can induce lower discounting and cigarette consumption. *Journal of Studies on Alcohol and Drugs*, 78(1):106–112.

Cole, S. N., Morrison, C. M., and Conway, M. A. (2013). Episodic future thinking: Linking neuropsychological performance with episodic detail in young and old adults. *Quarterly Journal of Experimental Psychology*, 66(9):1687–1706.

Conway, M. A., Justice, L. V., and D'Argembeau, A. (2019). The self-memory system revisited. *J. Mace, The organization and structure of autobiographical memory*, pages 28–51.

Craver, C. F., Kwan, D., Steindam, C., and Rosenbaum, R. S. (2014). Individuals with episodic amnesia are not stuck in time. *Neuropsychologia*, 57:191–195.

Dalley, J. W. and Robbins, T. W. (2017). Fractionating impulsivity: neuropsychiatric implications. *Nature Reviews Neuroscience*, 18(3):158.

Daniel, T. O., Sawyer, A., Dong, Y., Bickel, W. K., and Epstein, L. H. (2016). Remembering versus imagining: When does episodic retrospection and episodic prospection aid decision making? *Journal of Applied Research in Memory and Cognition*, 5(3):352–358.

De Leeuw, J. R. (2015). jspsych: A javascript library for creating behavioral experiments in a web browser. *Behavior research methods*, 47(1):1–12.

De Luca, C. R., Wood, S. J., Anderson, V., Buchanan, J.-A., Proffitt, T. M., Mahony, K., and Pantelis, C. (2003). Normative data from the cantab. i: development of executive function over the lifespan. *Journal of clinical and experimental neuropsychology*, 25(2):242–254.

Duval, C., Desgranges, B., de La Sayette, V., Belliard, S., Eustache, F., and Piolino, P. (2012). What happens to personal identity when semantic knowledge degrades? a study of the self and autobiographical memory in semantic dementia. *Neuropsychologia*, 50(2):254–265.

D'Argembeau, A. and Van der Linden, M. (2012). Predicting the phenomenology of episodic future thoughts. *Consciousness and cognition*, 21(3):1198–1206.

Ebert, J. E. (2001). The role of cognitive resources in the valuation of near and far future events. *Acta psychologica*, 108(2):155–171.

Ebert, J. E. and Prelec, D. (2007). The fragility of time: Time-insensitivity and valuation of the near and far future. *Management science*, 53(9):1423–1438.

Engelhard, I. M., van den Hout, M. A., Dek, E. C., Giele, C. L., van der Wielen, J.-W., Reijnen, M. J., and van Roij, B. (2011). Reducing vividness and emotional intensity of recurrent "flashforwards" by taxing working memory: An analogue study. *Journal of anxiety disorders*, 25(4):599–603.

Franck, C. T., Koffarnus, M. N., House, L. L., and Bickel, W. K. (2015). Accurate characterization of delay discounting: A multiple model approach using approximate bayesian model selection and a unified discounting measure. *Journal of the experimental analysis of behavior*, 103(1):218–233.

Hardisty, D. J. and Weber, E. U. (2009). Discounting future green: money versus the environment. *Journal of Experimental Psychology: General*, 138(3):329.

Hershfield, H. E. (2011). Future self-continuity: How conceptions of the future self transform intertemporal choice. *Annals of the New York Academy of Sciences*, 1235:30.

Hoerl, C. and McCormack, T. (2016). Making decisions about the future. *Seeing the Future*, pages 241–66.

Hollis-Hansen, K., O'Donnell, S. E., Seidman, J. S., Brande, S. J., and Epstein, L. H. (2019). Improvements in episodic future thinking methodology: Establishing a standardized episodic thinking control. *PloS one*, 14(3).

Holmes, E. A., Crane, C., Fennell, M. J., and Williams, J. M. G. (2007). Imagery about suicide in depression—"flash-forwards"? *Journal of behavior therapy and experimental psychiatry*, 38(4):423–434.

Johnson, M. K., Foley, M. A., Suengas, A. G., and Raye, C. L. (1988). Phenomenal characteristics of memories for perceived and imagined autobiographical events. *Journal of Experimental Psychology: General*, 117(4):371.

Johnson, M. W. and Bickel, W. K. (2008). An algorithm for identifying nonsystematic delay-discounting data. *Experimental and clinical psychopharmacology*, 16(3):264.

Kim, H., Schnall, S., and White, M. P. (2013). Similar psychological distance reduces temporal discounting. *Personality and Social Psychology Bulletin*, 39(8):1005–1016.

Kinley, I., Amlung, M., and Becker, S. (2022). Pathologies of precision: A bayesian account of goals, habits, and episodic foresight in addiction. *Brain and Cognition*, 158:105843.

Kinley, I., Porteous, M., Levy, Y., and Becker, S. (2021). Visual perspective as a two-dimensional construct in episodic future thought. *Consciousness and Cognition*, 93:103148.

Kirby, K. N., Winston, G. C., and Santiesteban, M. (2005). Impatience and grades: Delay-discount rates correlate negatively with college gpa. *Learning and individual Differences*, 15(3):213–222.

Klein, S. (2013). Making the case that episodic recollection is attributable to operations occurring at retrieval rather than to content stored in a dedicated subsystem of long-term memory. *Frontiers in behavioral neuroscience*, 7:3.

Klein, S. B. (2014). Sameness and the self: philosophical and psychological considerations. *Frontiers in psychology*, 5:29.

Koffarnus, M. N. and Bickel, W. K. (2014). A 5-trial adjusting delay discounting task: Accurate discount rates in less than one minute. *Experimental and clinical psychopharmacology*, 22(3):222.

Kwan, D., Craver, C. F., Green, L., Myerson, J., Boyer, P., and Rosenbaum, R. S. (2012). Future decision-making without episodic mental time travel. *Hippocampus*, 22(6):1215–1219.

Kwan, D., Craver, C. F., Green, L., Myerson, J., Gao, F., Black, S. E., and Rosenbaum, R. S. (2015). Cueing the personal future to reduce discounting in intertemporal choice: Is episodic prospection necessary? *Hippocampus*, 25(4):432–443.

Kwan, D., Craver, C. F., Green, L., Myerson, J., and Rosenbaum, R. S. (2013). Dissociations in future thinking following hippocampal damage: evidence from discounting and time perspective in episodic amnesia. *Journal of Experimental Psychology: General*, 142(4):1355.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478.

Lehner, E. and D'Argembeau, A. (2016). The role of personal goals in autonoetic experience when imagining future events. *Consciousness and Cognition*, 42:267–276.

Lempert, K. M. and Phelps, E. A. (2016). The malleability of intertemporal choice. *Trends in Cognitive Sciences*, 20(1):64–74.

Levine, B., Turner, G. R., Tisserand, D., Hevenor, S. J., Graham, S. J., and McIntosh, A. R. (2004). The functional neuroanatomy of episodic and semantic autobiographical remembering: a prospective functional mri study. *Journal of cognitive neuroscience*, 16(9):1633–1646.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. *Commons, ML.; Mazur, JE.; Nevin, JA*, pages 55–73.

McCarroll, C. J. and Cosentino, E. (2020). Rewarding one's future self: Psychological connectedness, episodic prospection, and a puzzle about perspective. *Review of Philosophy and Psychology*, 11(2):449–467.

McCue, R., McCormack, T., McElnay, J., Alto, A., and Feeney, A. (2019). The future and me: Imagining the future and the future self in adolescent decision making. *Cognitive Development*, 50:142–156.

Miglin, R., Kable, J. W., Bowers, M. E., and Ashare, R. L. (2017). Withdrawal-related changes in delay discounting predict short-term smoking abstinence. *Nicotine & Tobacco Research*, 19(6):694–702.

Myerson, J. and Green, L. (1995). Discounting of delayed rewards: Models of individual choice. *Journal of the experimental analysis of behavior*, 64(3):263–276.

Odum, A. L. (2011a). Delay discounting: I'm a k, you're a k. *Journal of the experimental analysis of behavior*, 96(3):427–439.

Odum, A. L. (2011b). Delay discounting: trait variable? *Behavioural processes*, 87(1):1–9.

Özbek, M., Bohn, A., and Berntsen, D. (2020). Characteristics of personally important episodic memories, counterfactual thoughts, and future projections across age and culture. *Applied Cognitive Psychology*, 34(5):1020–1033.

Palombo, D. J., Keane, M. M., and Verfaellie, M. (2015). The medial temporal lobes are critical for reward-based decision making under conditions that promote episodic future thinking. *Hippocampus*, 25(3):345–353.

Parfit, D. (1982). Personal identity and rationality. *Synthese*, pages 227–241.

Parfit, D. (1984). *Reasons and persons*. OUP Oxford.

Perrin, D. and Rousset, S. (2014). The episodicity of memory. *Review of Philosophy and Psychology*, 5(3):291–312.

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

Prebble, S. C., Addis, D. R., and Tippett, L. J. (2013). Autobiographical memory and sense of self. *Psychological bulletin*, 139(4):815.

R Core Team (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Rachlin, H. (2006). Notes on discounting. *Journal of the experimental analysis of behavior*, 85(3):425–435.

Radu, P. T., Yi, R., Bickel, W. K., Gross, J. J., and McClure, S. M. (2011). A mechanism for reducing delay discounting by altering temporal attention. *Journal of the experimental analysis of behavior*, 96(3):363–385.

Rasmussen, K. W. and Berntsen, D. (2014). Autobiographical memory and episodic future thinking after moderate to severe traumatic brain injury. *Journal of Neuropsychology*, 8(1):34–52.

Reimers, S., Maylor, E. A., Stewart, N., and Chater, N. (2009). Associations between a one-shot delay discounting measure and age, income, education and real-world impulsive behavior. *Personality and Individual Differences*, 47(8):973–978.

Rendell, P. G., Bailey, P. E., Henry, J. D., Phillips, L. H., Gaskin, S., and Kliegel, M. (2012). Older adults have greater difficulty imagining future rather than atemporal experiences. *Psychology and Aging*, 27(4):1089.

Rosseel, Y. (2012). lavaan: An R package for structural equation modeling. *Journal of Statistical Software*, 48(2):1–36.

Rung, J. M. and Madden, G. J. (2018). Experimental reductions of delay discounting and impulsive choice: A systematic review and meta-analysis. *Journal of experimental psychology: general*, 147(9):1349.

Rung, J. M. and Madden, G. J. (2019). Demand characteristics in episodic future thinking ii: The role of cues and cue content in changing delay discounting. *Experimental and clinical psychopharmacology*, 27(5):482.

Samuelson, P. A. (1937). A note on measurement of utility. *The review of economic studies*, 4(2):155–161.

Sasse, L. K., Peters, J., and Brassen, S. (2017). Cognitive control modulates effects of episodic simulation on delay discounting in aging. *Frontiers in Aging Neuroscience*, 9:58.

Sasse, L. K., Peters, J., Büchel, C., and Brassen, S. (2015). Effects of prospective thinking on intertemporal choice: The role of familiarity. *Human brain mapping*, 36(10):4210–4221.

Schacter, D. L. and Addis, D. R. (2007). On the constructive episodic simulation of past and future events. *Behavioral and Brain Sciences*, 30(3):331–332.

Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, pages 461–464.

Sokol, Y. and Serper, M. (2019). Development and validation of a future self-continuity questionnaire: A preliminary report. *Journal of Personality Assessment*.

Stein, J. S., Tegge, A. N., Turner, J. K., and Bickel, W. K. (2018). Episodic future thinking reduces delay discounting and cigarette demand: an investigation of the good-subject effect. *Journal of Behavioral Medicine*, 41(2):269–276.

Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction.* MIT press.

Szpunar, K. K. (2010). Episodic future thought: An emerging concept. *Perspectives on Psychological Science*, 5(2):142–162.

Szpunar, K. K. and Tulving, E. (2011). Varieties of future experience. *Predictions in the brain: Using our past to generate a future*, pages 3–12.

Tausen, B. M., Carpenter, S., and Macrae, C. N. (2020). Just another variant of psychological distance? the role of visual perspective in mental simulation. *Psychology of Consciousness: Theory, Research, and Practice*, 7(4):351.

Tooming, U. and Miyazono, K. (2021). Vividness as a natural kind. *Synthese*, 199(1):3023–3043.

Trope, Y. and Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological review*, 117(2):440.

Tulving, E. (1972). Episodic and semantic memory. *Organization of memory*, 1:381–403.

Tulving, E. (1985). Memory and consciousness. *Canadian Psychology/Psychologie canadienne*, 26(1):1.

Umanath, S. and Coane, J. H. (2020). Face validity of remembering and knowing: Empirical consensus and disagreement between participants and researchers. *Perspectives on Psychological Science*, 15(6):1400–1422.

van Schie, C. C., Chiu, C.-D., Rombouts, S. A., Heiser, W. J., and Elzinga, B. M. (2019). When i relive a positive me: Vivid autobiographical memories facilitate autonoetic brain activation and enhance mood. *Human brain mapping*, 40(16):4859–4871.

Wileyto, E. P., Audrain-McGovern, J., Epstein, L. H., and Lerman, C. (2004). Using logistic regression to estimate delay-discounting functions. *Behavior Research Methods, Instruments, & Computers*, 36(1):41–51.

Zaman, A. and Russell, C. (2022). Does autonoetic consciousness in episodic memory rely on recall from a first-person perspective? *Journal of Cognitive Psychology*, 34(1):9–23.

Chapter 5

Discussion

5.1 Summary and main contributions of chapters

The chapters presented in this thesis are connected by a motivation to identify moderating and mediating factors in the effect of episodic future thinking on delay discounting, and to suggest mechanisms by which this effect occurs. In this final discussion, I will highlight the preceding chapters' individual contributions and explain their connections to each other. I will then consider their limitations and suggest future directions for research on the relationship between episodic future thinking and delay discounting.

Chapter 2 lays the groundwork for the subsequent chapters with a formal modelling approach. It contextualizes the topic of episodic future thinking and delay discounting within the setting of reinforcement learning, a field that examines the essential elements of decision making problems and of the agents that attempt to solve them. The chapter proposes a novel formulation of Bayesian Q-learning that, in contrast to existing formulations (Daw et al., 2005; Dearden et al., 1998), generalizes to real-valued (vs binary) rewards and offers an exact solution for the reward expectancies of model-free agents rather than relying on approximations. Within this mathematical framework, episodic future thinking can be understood as a form of model-based decision making. The chapter demonstrates that, given an unbiased choice of parameters, model-based agents exhibit shallower discounting than their model-free counterparts. Such a difference between agent types is assumed by dual-systems theories of self-control in which steep discounting arises from the dominance of a habitual system over a deliberative system (McClure and Bickel, 2014), but had not, to my knowledge, been demonstrated computationally.

Although model-based agents exhibit shallower discounting for an unbiased choice of parameters, some authors argue that the brain's model-based system should be sensitive to current goals, such as the goal of forgoing immediate rewards in favour of long-term gains (Story et al., 2014). This is simulated in chapter 2 by adding bias to the model-based agent's parameters to more heavily favour delayed rewards. Yet by the same token, such flexibility means there is in principle no reason why a model-based agent could not also preferentially select immediate rewards and exhibit even steeper discounting than a model-free agent. This, along with the fact that model-based decision making is not immune to overvaluing addictive rewards, is marshalled in a critique of the habit theory of addiction.

More importantly for the core theme of this thesis, chapter 2 offers a framework for understanding the role of episodic future thinking interventions in curbing delay discounting: if discounting is a function of the expected value and uncertainty of future rewards, then episodic future thinking interventions can be expected to have their effect by influencing one or both of these. The chapter suggests that, in structured episodic future thinking interventions, positive imagery might correspond to high expected reward value, while vivid and detailed imagery might correspond to low uncertainty about the future. Thus, chapter 2 lays the groundwork for an empirical search for moderators and mediators of the effect of episodic future thinking on delay discounting.

Chapter 3 takes up this task, moving from the pristine realm of mathematical abstractions to the real world of empirical measurement. It uses the well-established cued discounting paradigm (Peters and Büchel, 2010) to measure changes in participants' discount functions in response to episodic future thinking cues (henceforth the "cue effect"). To search for factors moderating this effect, the study uses a broad questionnaire querying the phenomenology of participants' mental imagery, measuring self-reported variables such as emotional valence/intensity and vividness. Of particular interest was the visual perspective from which participants imagined future events, as this influences both the overall phenomenology and the behavioural impact of imagery (Libby and Eibach, 2011).

To impose minimal assumptions on the data, participants were able to separately specify their degrees of first- and third-person perspective. In contrast to existing research that assumes these are mutually exclusive, we found that participants report *both* first- and third-person perspectives on the same imagined event, perhaps switching between the two. Also, while much research finds that third-person perspective dampens affect, we found third-person imagery where the future self is visible to be associated with higher ratings of emotional intensity. We interpreted this finding according to a theoretical model proposed by Sutin and Robins (2008) in which third-person perspective can heighten affect by increasing the salience of the visualized self. Finally, the degree to which participants reported *both* first- and third-person imagery was correlated with their degree of trait dissociation, as measured by the Dissociative Experiences Scale [cite]. This surprising finding has since been replicated in a larger study, demonstrating that "dual perspective" imagery is not a curious inconsistency in self-report data, but a genuine phenomenon related to established psychological constructs.

Thus, chapter 3 contributes to a basic understanding of episodic future thinking. Interestingly, however, where its effect on delay discounting is concerned, we did not identify any moderating factors: no variable measured on the phenomenological questionnaire was related to the strength of the cue effect. One explanation may have been a relative lack of variance: perhaps with a sample of healthy undergraduates who were all given the same instructions, responses to the phenomenological questionnaire did not vary enough to reveal any moderation. One solution might be to use an experimental manipulation to introduce greater variance into participants' self-reported phenomenology.

Chapter 4 attempts to do just this, assigning participants either to an experimental condition designed to elicit vivid, positive, personally relevant imagery or to a condition designed to elicit relatively vague, affectively flat imagery. In the course of explaining the motivation for this design, the chapter makes several theoretical contributions. First, it offers a critique of the prevailing experimental method of comparing cue effects elicited by episodic future thinking versus recent episodic memories (Hollis-Hansen et al., 2019). This method systematically precludes any investigation of whether the cue effect is actually due to *episodic* future thinking: are the distinctive characteristics of episodic future thinking necessary for the cue effect, or is any conceptual awareness of the future (versus the recent past) sufficient? We propose that episodic future thinking must be compared to *semantic future thinking* to determine whether episodicity is essential to the cue effect, and another of the chapter's contributions is to clarify and operationalize this latter notion.

Episodic future thinking was originally defined in contrast to semantic future thinking, following the distinction between episodic and semantic memory (Atance and O'Neill, 2001). However, as the chapter points out, the episodic–semantic distinction in future thinking is not exactly analogous to the episodic–semantic distinction in memory: we can imagine essentially any future event but cannot recollect everything in our past; moreover, semantic future thinking must be future-oriented whereas semantic memory can be atemporal. Following Atance and O'Neill (2001), we suggest semantic future thinking is distinguished by non-self-specificity and a lack of autonoetic consciousness. Thus, in the experimental manipulation, the episodic and semantic future thinking groups were operationally distinguished according to whether participants were asked to describe self-specific events. This manipulation successfully produced differences in self-reported autonoetic consciousness. Chapter 4 thus offers an experimental method by which semantic future thinking can be compared to episodic future thinking.

The chapter proposes that autonoetic consciousness, as a factor differentiating episodic and semantic future thinking, may explain any unique effect of episodic future thinking on delay discounting. This proposal connects back to chapter 2, which argued for an analogy between model-based decision making, which is based on "simulated experience", and episodic future thinking, which similarly entails mentally simulating the experience of a future event (i.e., autonoetic consciousness). The chapter consists of two studies and a simulation demonstrating how to overcome a subtle disadvantage of the widely-used "titrating" method for measuring discount functions. However, despite rigorous quantitative methods, we did not find that autonoetic consciousness (or any other construct measured by our phenomenological questionnaire) mediated the EFT cue effect. We suggest that, given that the self-relevance of imagined future events was what operationally distinguished the two experimental groups, perhaps self-relevance *per se* is what mediated the EFT cue effect. We relate this to the concept of "semantic" self-continuity (Prebble et al., 2013), in which conceptual representations of an individual's life narrative underlie their temporally extended sense of self and, perhaps, their temporally extended self-interest. Later in this discussion, I will pick up this thread to suggest the need to look beyond "experiential" factors explaining the effect of EFT on decision making. First, however, it is necessary to consider various limitations of the current studies.

5.2 Methodological limitations

Our search for moderating and mediating factors of the cue effect may have been hampered by the way candidate variables were measured. In the studies described in this thesis, participants provided ratings of vividness, valence, etc. during an initial future thinking writing task used to generate cues for a subsequent intertemporal choice task. However, it is possible that an event vividly imagined during the writing phase might not be vividly evoked by cues during the decision phase and vice versa. In contrast to this approach, other studies have collected ratings from participants *after* decisions, pertaining to their imagery *during* decisions (Bulley and Gullo, 2017; Hollis-Hansen et al., 2019; Peters and Büchel, 2010). Thus, by measuring imagery elicited by cues, this retrospective rating method appears to be more suited to identifying moderators of the cue effect than our prospective rating method. Indeed, based on this consideration, our more recent work has collected trial-by-trial ratings of vividness during the decision task to maximize the "resolution" of the data.

Moreover, in the studies reported here, participants were cued to imagine future events that did not necessarily involve the delayed rewards. For example, while deciding whether to choose $100 in 3 months, a participant might be cued to imagine being on a camping trip (where presumably that money would not be useful). In contrast, Rösch et al. (2022) note that in some studies, the imagined future event was specific in some way to the delayed reward (for example, imagining spending the larger-later amount of money; these studies tended to measure larger cue effects). The "content specificity" of the imagined future events in these studies aligns with the theoretical account of the cue effect described in the introduction to this thesis: episodic future thinking is proposed to provide, through mental simulation, a preview of future reward. However, it can only do this if the future reward is actually imagined. Perhaps it is not surprising, then, if the vividness with which someone imagines a camping trip in 3 months is unrelated to their willingness to wait for an unrelated monetary reward in 3 months.

Nonetheless, despite these methodological shortcomings, a strength of the study in chapter 4 is that it used an experimental design and attempted to identify mediators of the cue effect rather than merely moderators. This type of work will be important to evaluate causal claims about the role of episodic future thinking in decision making. After all, there is a range of aspects of episodic future thinking that are plausibly related to decision making but are correlated with one another. For example, as discussed in chapter 2, vividness and autonoetic consciousness are correlated—if vividness is a moderator of the cue effect, is this an artefact of its correlation with autonoetic consciousness? Moreover, events related to one's personal goals tend to elicit autonoetic consciousness (Lehner and D'Argembeau, 2016)—if autonoetic consciousness is a moderator of the cue effect, is this an artefact of its relationship to goal-relevance? Questions such as these must be addressed by experimental means.

5.3 Theoretical limitations and future directions

A deeper limitation of the studies in this thesis is their reliance on the cued discounting paradigm, in which participants are cued to imagine a future event while deliberating between a smaller monetary reward sooner and a larger one later. As part of this deliberation, participants engage in episodic future thinking directed, in some sense, toward the delayed reward. This suggests a template for intertemporal choice in which the role of episodic future thinking is straightforward: in simulating the experience of the future reward, episodic future thinking provides a mental "preview" of it and allows its undiscounted value to be appreciated in the present (Boyer, 2008; Peters and Büchel, 2010). As outlined in chapter 1, we can derive the prediction that, the higher the fidelity of this simulation (i.e., the greater the vividness) and the more value ascribed to the future reward (i.e., the more positive the valence), the greater the effect on decision making.

However, there are cases of intertemporal choice that do not seem to fit this template. First, the relationship between imagined emotion and temporal discounting is not always so straightforward: a future-oriented choice could be framed not as the obtainment of future reward but as the avoidance of future consequence (e.g., the consequence of future regret; Hoerl and McCormack, 2016). In this case, imagining a future consequence to be more severe would motivate more farsighted decision making—contrary to the expectation that negatively valenced episodic future thinking will increase impatience or have no effect. Cosentino et al. (2022) argue that metacognition is required to "bracket" imagined future emotions to keep them from unduly influencing our current affect (thereby maintaining the distinction between *anticipated* and *anticipatory* emotion; Baumgartner et al., 2008) and to understand their relationship to present actions. Thus, although there is no simple linear relationship between the emotional valence of episodic future thinking and its influence on intertemporal patience, the model of the cue effect outlined in the introduction to this thesis is still viable if it is supplemented by metacognition.

A deeper difficulty arises from the fact that episodic future thinking can motivate us toward goals that are *conceptual* rather than experiential, i.e., toward rewards that cannot be simulated by imagery alone. The distinction between conceptual and experiential rewards can be understood as follows: in a famous thought experiment, Nozick (1974) invites us to imagine an "experience machine" that we could plug our brains into to simulate any conceivable experience, such as reading or even writing a great novel. Few of us would choose this, Nozick argues, because we aspire not just to *have* certain experiences, but to *be* certain types of people (e.g., to be a great novelist beyond merely having a great novelist's experiences). Indeed, our goals can be conceived as lying on an abstraction gradient from high-level "be" goals (e.g., "be healthy") to mid-level "do" goals (e.g., "do some exercise") to low-level "action" goals (e.g., "put on running shoes"; Carver and Scheier, 2001). Monetary delay discounting tasks largely overlook the abstract end of this gradient: while monetary rewards are abstract compared to primary reinforcers such as food (O'Doherty et al., 2001), they are fairly concrete in absolute terms—even the future-oriented monetary goal of "opt for \$100 in 3 months over \$50 now" is concrete compared to "save for the future", which is concrete compared to "be financially shrewd" and so on.

Abstract or conceptual "be" goals are difficult to represent via imagery (as are concepts in general—e.g., we can picture a triangle, but not triangularity *per se*; Fodor and Pylyshyn, 2014). After all, if a reward is fundamentally not experiential, how could we mentally simulate the experience of it? Yet episodic future thinking *does* motivate us toward abstract goals, and in fact seems to do so *specifically by emphasizing their abstractness*: for example, imagining success at an academic task leads to more motivation when it causes the necessary effort to be construed in terms of abstract rather than concrete goals (e.g., "being the best I can be" rather than "trying hard at a task"; Vasquez and Buehler, 2007). Interestingly, visual perspective appears to be an important determinant of whether imagined experiences are construed abstractly (encouraged by a third-person perspective) or concretely (encouraged by a first-person perspective), and thus whether motivation toward abstract goals is enhanced (Libby and Eibach, 2011; Niese et al., 2022). Cases where episodic future thinking enhances motivation toward abstract goals are difficult to explain through the simulated experience of future rewards, metacognitively "bracketed" (Cosentino et al., 2022) or not, if these rewards are about *being* rather than *experiencing*[1].

I would like to propose the following solution: perhaps episodic imagery does not *simulate* abstract rewards so much as it *symbolizes* them. In semiotics, a distinction is drawn between icons, which signify by resemblance (e.g., a picture of a bicycle), and symbols, which signify by convention (e.g., a written word). Applying this distinction to episodic future thinking, we might take our imagery to merely simulate some experience or instead to *symbolize a concept*—perhaps a valenced concept, i.e., an abstract goal, which our imagery makes salient and which then influences subsequent decision making. For example, when I imagine my graduation, I do not take my imagery to merely simulate of the feeling of a gown on my shoulders, the stage floor under my feet, etc.—instead, my focus is on the *symbolic meaning* of the imagined event and its relationship to my abstract

[1]A possible objection here is that one could mentally simulate the abstract reward of being a wise, kind, etc. person by simulating the "epistemic feeling" (Arango-Muñoz, 2014) of knowing that one is a wise, kind, etc. person. Then episodic future thinking could provide a mental "preview" of abstract rewards through this simulated "feeling of knowing" just as it provides a mental preview of concrete rewards through simulated sensory experience. However, knowledge is only knowledge (rather than suspicion, belief, etc.) if it reflects something true about the world. Therefore simulating the "feeling of knowing" presupposes a representation of some state of the world that one knows about, and in fact this state (rather than the experience *per se* of knowing about it) *is* the reward: it is pleasant to imagine being wiser, kinder, etc. one realizes one is; it is not pleasant to imagine believing one is wiser, kinder, etc. than one really is. That is, the "feeling of knowing" is rewarding not because of the feeling, but because of the knowing (vs suspecting, believing, etc.). Thus, the question still remains: if episodic future thinking simulates experience, but abstract rewards are conceptual rather than experiential, how does episodic future thinking motivate us toward abstract goals?

goal of being an educated person[2].

This proposal is related to one recently made by Mahr (2020), who argues that the "imagistic" content of episodic simulations (which comprise future thoughts, memories, counterfactual imaginings, etc.) must be supplemented by propositional content in order for these simulations to be specified *as* future thoughts vs memories vs counterfactual imaginings, etc. For example, if I picture myself sitting in a cafe, nothing in this image itself (in "what [the] image looks like"; Mahr, 2020) specifies whether I am sitting there in the future or the past—this information is propositional and cannot be represented imagistically (Fodor and Pylyshyn, 2014). Mahr (2020) suggests four major "dimensions" of episodic simulations that depend on non-imagistic content: temporal orientation (whether an event is in the future, past, neither), subjectivity (whose perspective, if anyone's, is being simulated), factuality (whether an event really did/will/could occur), and specificity (whether an event is unique to one point in time). This last dimension is most relevant to the present discussion of symbolization, because it implies that a specific episodic simulation can be taken to represent a general class of events. For example, I can imagine my commute, including specific things I experience as part of my commute, without necessarily imagining a specific time I did/will commute (Mahr, 2020). However, the present proposal goes further: episodic simulations can represent not only general classes of events, but also concepts at large—for example, my commute imagery might represent not only my commute in general, but also the valenced concept of my industriousness.

Abstract "be" goals are closely tied to our personal identities (Berkman et al., 2017). If episodic future thinking motivates us toward these goals, this raises interesting considerations about the relationship between future thinking and identity. As in the case of anticipated/anticipatory emotion described earlier, the sequence of connections from imagery to abstract goals to identity to motivation are straightforward as long as each element of the sequence is aligned positively: imagining an event that symbolizes an abstract goal central to one's self-concept should increase motivation toward that goal. However, things are not always so straightforward: for goals peripheral to one's self-concept, third-person perspective imagery (which encourages abstract construal) enhances *negative* self-conscious emotion (e.g., shame) and *inhibits* goal pursuit (Stornelli et al., 2020). To make sense of this result, we can note that goals peripheral to one's self-concept are not necessarily *irrelevant* to it: we are most likely to feel shame and other negative self-conscious emotions in relation to goals we believe we will not achieve but wish we could (notably, the Stornelli et al. study used weight loss goals, which notoriously fit this template for many of us). That is, abstract goals are not simply more

[2] This account does not fully specify the relationship between imagery and concepts except to say that it is symbolic. For example, does imagery *per se* symbolize a concept, or does one simply imagine an event of symbolic importance? Likewise, should we understand abstract goals to be "activated", "made salient", or "primed" (Papies, 2016)? Ultimately, it is not clear that distinctions like these could produce contrasting experimental predictions. The core idea here is simply to observe that imagery motivates us by somehow increasing the influence of valenced concepts on our decisions, and to recognize that to get from image to concept, we must at some point invoke a symbolic relationship

or less central or peripheral to the self, but are involved propositional attitudes (Fodor, 1978) that might form part of the self-concept (e.g., "I believe I will never lose weight", "I wish I could lose weight") and that determine affective responses to these goals being made salient.

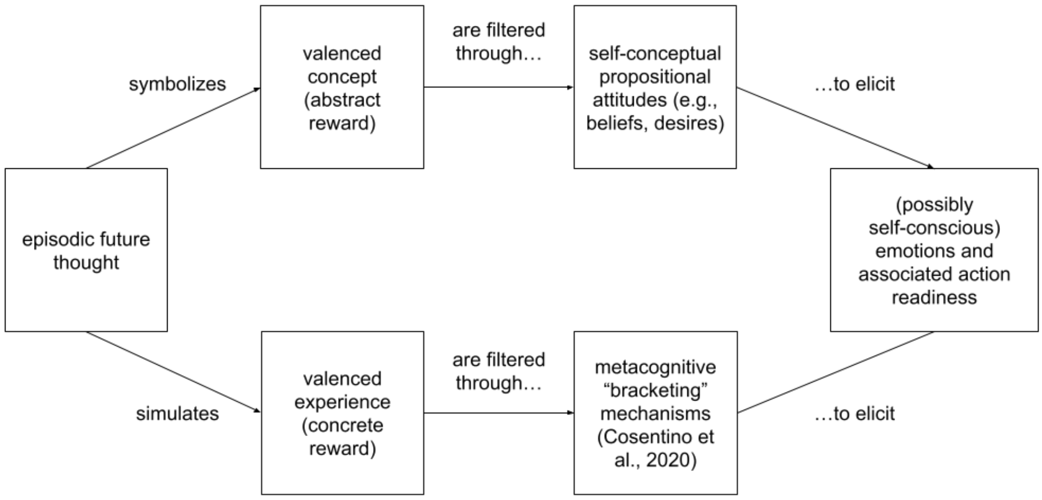

Figure 5.1: A proposed expansion on the model of episodic future thinking and decision making that was described in the introduction to this thesis. In the expanded model, episodic future thinking simulates the experience of concrete rewards, and these simulated experiences are metacognitively "bracketed" (Cosentino et al., 2022) to motivate the pursuit of said rewards. In contrast, episodic future thinking symbolizes (and thereby makes salient) abstract or conceptual rewards, and an individual's beliefs about and attitudes toward these rewards (which may form part of their self-concept) determine motivation toward them.

The model sketched in this discussion is summarized in figure 5.1. It proposes a second "route" from episodic future thought to behaviour, through symbolized concepts rather than simulated experiences. The non-imagistic components of this second route (concepts, propositions) suggest that careful qualitative analysis of individuals' written descriptions of imagined future events, beyond simple queries of their imagery, is necessary to understand how episodic future thinking produces motivation (or not) toward abstract goals. Moreover, a fascinating question for future research might be whether there is a *bidirectional* relationship between episodic future thinking and identity: our goals and self-concepts are important determinants of which future events we find plausible (Ernst and D'argembeau, 2017; Ernst et al., 2019), but future events can also come to feel more plausible as they are repeatedly imagined (Szpunar and Schacter, 2013). Is it possible, then, that in imagining and coming to believe in a different future for oneself, one might also come to see oneself as a different kind of person?

In sum, the empirical and theoretical work contained in this thesis highlights the need to move beyond the straightforward conception of intertemporal choice implied by the cueing paradigm that dominates the literature, and to expand our focus beyond the content concretely depicted in our imagery. Instead, we will have to face up to the complexities of our real decisions about the future, which can be framed in a variety of ways and can involve rewards of varying degrees of abstraction. Moreover, we will have to recognize how our capacity to simulate future experience integrates with the rest of our psyche, necessarily invoking notions of meaning and identity. As scientifically elusive as these are, their enormous role in our mental lives means that a science of the mind—including its capacity to project itself forward in time—cannot be complete without them.

Bibliography

Arango-Muñoz, S. (2014). The nature of epistemic feelings. *Philosophical Psychology*, 27(2):193–211.

Atance, C. M. and O'Neill, D. K. (2001). Episodic future thinking. *Trends in cognitive sciences*, 5(12):533–539.

Baumgartner, H., Pieters, R., and Bagozzi, R. P. (2008). Future-oriented emotions: Conceptualization and behavioral effects. *European Journal of Social Psychology*, 38(4):685–696.

Berkman, E. T., Livingston, J. L., and Kahn, L. E. (2017). Finding the "self" in self-regulation: The identity-value model. *Psychological Inquiry*, 28(2-3):77–98.

Boyer, P. (2008). Evolutionary economics of mental time travel? *Trends in cognitive sciences*, 12(6):219–224.

Bulley, A. and Gullo, M. J. (2017). The influence of episodic foresight on delay discounting and demand for alcohol. *Addictive Behaviors*, 66:1–6.

Carver, C. S. and Scheier, M. F. (2001). *On the self-regulation of behavior*. cambridge university press.

Cosentino, E., McCarroll, C. J., and Michaelian, K. (2022). Resisting temptation and overcoming procrastination: The roles of mental time travel and metacognition. *Phenomenology and the Cognitive Sciences*, 21(4):791–811.

Daw, N. D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*, 8(12):1704–1711.

Dearden, R., Friedman, N., and Russell, S. (1998). Bayesian q-learning. *Aaai/iaai*, 1998:761–768.

Ernst, A. and D'argembeau, A. (2017). Make it real: Belief in occurrence within episodic future thought. *Memory & Cognition*, 45:1045–1061.

Ernst, A., Scoboria, A., and D'argembeau, A. (2019). On the role of autobiographical knowledge in shaping belief in the future occurrence of imagined events. *Quarterly Journal of Experimental Psychology*, 72(11):2658–2671.

Fodor, J. A. (1978). Propositional attitudes. *The Monist*, 61(October).

Fodor, J. A. and Pylyshyn, Z. W. (2014). *Minds without Meanings: An Essay on the Content of Concepts*. MIT Press.

Hoerl, C. and McCormack, T. (2016). Making decisions about the future. *Seeing the future: Theoretical perspectives on future-oriented mental time travel*, page 241.

Hollis-Hansen, K., O'Donnell, S. E., Seidman, J. S., Brande, S. J., and Epstein, L. H. (2019). Improvements in episodic future thinking methodology: Establishing a standardized episodic thinking control. *PLoS One*, 14(3):e0214397.

Lehner, E. and D'Argembeau, A. (2016). The role of personal goals in autonoetic experience when imagining future events. *Consciousness and Cognition*, 42:267–276.

Libby, L. K. and Eibach, R. P. (2011). Visual perspective in mental imagery: A representational tool that functions in judgment, emotion, and self-insight. In *Advances in experimental social psychology*, volume 44, pages 185–245. Elsevier.

Mahr, J. B. (2020). The dimensions of episodic simulation. *Cognition*, 196:104085.

McClure, S. M. and Bickel, W. K. (2014). A dual-systems perspective on addiction: contributions from neuroimaging and cognitive training. *Annals of the New York Academy of Sciences*, 1327(1):62–78.

Niese, Z. A., Eibach, R. P., and Libby, L. K. (2022). Picturing yourself: A social-cognitive process model to integrate third-person imagery effects. *Journal of Cognitive Psychology*, 34(1):24–44.

Nozick, R. (1974). *Anarchy, state, and utopia*, volume 5038. new york: Basic Books.

O'Doherty, J., Kringelbach, M. L., Rolls, E. T., Hornak, J., and Andrews, C. (2001). Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature neuroscience*, 4(1):95–102.

Papies, E. K. (2016). Goal priming as a situated intervention tool. *Current Opinion in Psychology*, 12:12–16.

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

Prebble, S. C., Addis, D. R., and Tippett, L. J. (2013). Autobiographical memory and sense of self. *Psychological bulletin*, 139(4):815.

Rösch, S. A., Stramaccia, D. F., and Benoit, R. G. (2022). Promoting farsighted decisions via episodic future thinking: A meta-analysis. *Journal of Experimental Psychology: General*, 151(7):1606.

Stornelli, J., Pereira, B., and Vann, R. J. (2020). Big picture, bad outcomes: When visual perspectives harm health goal pursuit. *Journal of Consumer Psychology*, 30(2):368–378.

Story, G. W., Vlaev, I., Seymour, B., Darzi, A., and Dolan, R. J. (2014). Does temporal discounting explain unhealthy behavior? a systematic review and reinforcement learning perspective. *Frontiers in behavioral neuroscience*, 8:76.

Sutin, A. R. and Robins, R. W. (2008). When the "i" looks at the "me": Autobiographical memory, visual perspective, and the self. *Consciousness and cognition*, 17(4):1386–1397.

Szpunar, K. K. and Schacter, D. L. (2013). Get real: effects of repeated simulation and emotion on the perceived plausibility of future experiences. *Journal of Experimental Psychology: General*, 142(2):323.

Vasquez, N. A. and Buehler, R. (2007). Seeing future success: Does imagery perspective influence achievement motivation? *Personality and social psychology bulletin*, 33(10):1392–1405.