Times Imagined and Remembered

Felipe De Brigard, Maria Khoudary and Samuel Murray
Duke University

1. Introduction

People have a remarkable capacity to think about times other than the present. When we remember, for instance, our thoughts pertain to events that happened in the past. When we plan or anticipate things to come, we think about possible events that may or may not occur in the future. And when we imagine how a past event could have gone differently, our thoughts are driven to the past, albeit one that we did not live. In the past few decades, a substantial amount of scientific evidence indicates that these three kinds of mental simulations—*episodic memory*, *episodic future thinking*, and *episodic counterfactual thinking*—are supported by a common set of cognitive and neural mechanisms (Schacter, Benoit, De Brigard, & Szpunar, 2015).[1] Initial evidence came from neuropsychological studies showing that individuals with amnesia due to medial temporal lobe damage have concomitant difficulties with episodic future (Tulving, 1983; Klein, Loftus, & Kihlstrom, 2002; Hassabis, Kumaran, Vann, & Maguire, 2007) and counterfactual thinking (Mullaly & Maguire, 2014). Similar deficits appear in individuals with impairments in episodic memory retrieval, such as older adults, people with severe depression, schizophrenia, amnestic mild cognitive impairment, and Alzheimer's disease—all conditions that have in common volumetric reductions in the hippocampus, a critical region for memory consolidation and retrieval (for recent reviews, see Schacter, Addis, et al, 2012; De Brigard & Gessell, 2016; De Brigard & Parikh, 2019). Developmental psychology also provides consistent evidence, as studies have shown similar patterns of development in children's capacities to remember their actual past and to imagine possible past and future events (Atance & O'Neill, 2001; Redshaw & Suddendorf, 2020). Further support comes from neuroimaging findings showing significant activation overlap in brain regions engaged during episodic past, future and counterfactual thinking (Addis, Wong, & Schacter, 2007; Szpunar, Watson, & McDermott, 2007; De Brigard, Addis, Ford, et al., 2013). Specifically, it has been shown that these three psychological capacities tend to recruit a set of brain regions known as the *default mode network*, which includes the posterior cingulate cortex and adjacent precuneus, the medial prefrontal cortex, the medial temporal lobe and lateral temporal cortices (Buckner & DiNicola, 2019). Finally, a number of studies have reported similarities among the phenomenological characteristics of episodic past, future and counterfactual thinking (e.g., D'Argembeau & Van der Linden, 2006; Szpunar & McDermott, 2008; Winfield & Kamboj, 2010; De Brigard & Giovanello, 2012), lending further credence to the claim that all three abilities may be underwritten by similar mechanisms.

Over a decade ago, the *constructive episodic simulation hypothesis* was put forth to account for the observation that these three kinds of episodic simulations recruit common neural mechanisms (Schacter & Addis, 2007). Based upon the idea that episodic recollection consists, not in the reproduction, but in the reconstruction of prior experienced episodes via computational

---

[1] Episodic thoughts contrast with semantic thoughts. The former are contextually-bounded representations of specific events, while the latter are context-free representations with generic content. Thus, one might have an episodic memory of burning the cake for last year's birthday party, or an episodic future thought about how to avoid burning the cake for this year, in contrast to the semantic memory that birthdays are typically celebrated on the day that one was born or a semantic future thought about the likelihood of an economic depression given current conditions. For more on the episodic/semantic distinction as applied to mental simulations, see De Brigard and Parikh, 2019.

processes that enable us to reassemble them (Schacter, Norman, & Koutstaal, 1998), the constructive episodic simulation hypothesis proposed that these same computational processes allow us to recombine episodic details into novel hypothetical simulations of possible events. Additionally, the constructive episodic simulation hypothesis holds that there is a functional explanation for the common recruitment of these processes. Specifically, it suggests that both the reconstruction of episodic memories and the generation of episodic future and counterfactual thoughts share an adaptive purpose: enabling the mental simulation of possible events in order to hedge future uncertainty (Schacter & Addis, 2007; De Brigard, 2014; Michaelian, 2016). The idea, roughly, is that we can detach our thoughts from the immediate present and imagine hypothetical events so that these mental simulations can—to use the Popperian trope—"die in our stead" (Popper, 1977; *cf.* Dennett, 1996, and his notion of "Popperian creature").

But how exactly do these episodic simulations help us at a later time? Following Szpunar and colleagues (2013; see also Schacter, Addis and Szpunar, 2017), we believe that part of the answer was anticipated by neurologist D. H. Ingvar, who also thought that memory has a future-oriented role because it allows people to imagine "alternative hypothetical behavior patterns in order to be ready for what may happen" (1979, 21). Ingvar further elaborated that these "simulations of behavior"—as he called them—need not be exact reproductions of past events; they can be hypothetical, as is the case with episodic future thoughts (Ingvar, 1985). What mattered for him, though, was that these mental simulations could serve their future-oriented role only insofar as one was able to properly deploy them when the right time comes. More precisely, Ingvar's insight was that in order for an episodic simulation—whether past, future or counterfactual—to be useful in the future, one needed to be able to encode the content of the mentally simulated event into a memory that could be retrieved later on, when the simulated information could be put to use. In the spirit of Ingvar, we call these recollections of previous mental simulations "memories for the future" (Szpunar et al, 2013).[2] To illustrate, suppose that in preparing for a future job interview, you imagine yourself giving a certain answer to a particular question you may be asked—that is to say, you engage in episodic future thinking. And now suppose that later on, during the actual interview, such a question is in fact posed. Your imaginative anticipatory preparation will only pay dividends if, during the actual interview, you remember what you mentally simulated before. Notice that the same is true about both memories and counterfactual thoughts. When planning your future interview, you may have remembered a past situation in which you gave a good answer, or maybe you imagined a possible past event when you wish you would have given a different response than you actually did. Being able to bring back to mind the right memory, the right counterfactual, or the right future thought when the appropriate time comes is, following Ingvar, how mental simulations of episodic past, future and counterfactual events can help us hedge future uncertainty.

In the current paper, we follow Ingvar's insight that the future-directed benefits of simulation depend, in part, on encoding, storing, and retrieving simulated information at the right time. However, it is important to note that there are important asymmetries between different kinds of simulations. After all, the mentally simulated content of an episodic memory involves the re-instatement of a previously experienced content. By contrast, both episodic future and

---

[2] It is important to distinguish between a memory whose content is that of a previous mental simulation versus a memory whose content is the very act of mentally simulating. Our concern here is with memories of a previously simulat*ed* content rather than of a prior simulat*ing*. In fact, one may be able to recall a previously simulated content without remembering anything about the act of simulating per se, and the retrieval of that memory for the future could still be beneficial. For more on the simulat*ed* vs simulat*ing* distinction, see Sellars, 1967.

counterfactual simulations include as contents events that were not experienced. Moreover, episodic memories not only involve contents that have been previously experienced—they often are remembered more than once. Episodic counterfactual and future thoughts, on the other hand, are typically entertained just once and often are not subsequently rehearsed. Given well documented differences in the phenomenology associated with mentally simulated contents of episodic memories versus episodic future and counterfactual thoughts (e.g., De Brigard, Giovanello, Stewart, et al., 2016), it is likely that our recollections of such simulations at a later time are also going to differ between those in which the encoded content was an episodic memory relative to those in which the mentally simulated content represents a possible future or counterfactual event. Unfortunately, the extent to which there may actually be such an asymmetry in our recollections of past mental simulations remains unclear, as research on how we encode and subsequently retrieve "memories for the future" is at its infancy (Szpunar, Addis et al, 2013). One important asymmetry, discussed below, concerns the accurate recall of different kinds of information.

The current paper aims to explain this *recall asymmetry* by examining how we encode and retrieve *temporal* information from episodic simulations. When we remember a previous experience, imagine how it could have occurred instead, or even fancy a possible scenario that may come to be in our future, the contents of our mental simulations are rich: they involve places, people, objects, actions, and so on. And they often involve temporal information as well; i.e., information about *when* an event happened, *when* it could have happened, or *when* it may come to occur. In addition to this kind of temporal information (i.e., when in time the simulated event occurs), there is also temporal information contained in the sequence of events taking place within a particular mental simulation (e.g., when we think that if X were to occur first, then Y and Z would follow). Both varieties of temporal information can be critical for guiding behavior. Consider instances of what is known as *prospective memory*: our psychological capacity to remember to bring about certain actions at a later time. Planning for the future often involves mentally simulating the anticipated episode, a process that involves considering both when the episode will take place and what sequence of actions and events will constitute it. For instance, if one is planning for a job interview, remembering that it is happening next Tuesday at 10:00 am *and* that one needs to print resume copies before getting into the car are both crucial for improving one's chances of getting the job.

Our discussion centers on two interrelated issues: 1) how the time of occurrence—the *when*—of actual and hypothetical events is incorporated when mentally simulating them as "memories for the future", and 2) how such temporal information is retrieved later on. Ultimately, we will defend the thesis that different cognitive processes are required to incorporate temporal information into our mental simulations, and that improving temporal memory requires integrating temporal information with other representational aspects of the simulated content. We argue that our account explains recent results showing that temporal information is poorly remembered for imagined as opposed to remembered simulations. We begin, then, in Section 2, with a critical survey of different views on how we remember temporal information about past events. This review supports a functional account of temporal memory, following Friedman's (1993; 2004) influential proposal. Next, in Section 3, we elaborate on this model and explain how it would extend to memory for imagined, as opposed to remembered, times. Although the proposal here is mostly speculative, we nevertheless discuss some empirical findings that can be explanatorily unified by our proposed framework. Finally, in section 4, we discuss some practical consequences that may follow from our proposed framework regarding how we can improve time-relevant episodic future thinking, with a focus on temporal discounting, prospective memory, and implementation of intention.

2. Memory for past times

We rely on episodic autobiographical memory to recall information about past personal experiences (Tulving, 1985). In recalling a past event—say, a past job interview, to continue with the above example—we may bring to mind who was there, where it happened, which objects were involved, how the event unfolded, or when it occurred. But our capacity to accurately recall these pieces of information isn't uniform; the rate of forgetting and amount of retrieval error varies as a function of the type of cue, the nature of the content, and the length of the retention interval, among other factors. Years of research have revealed that memory provides certain advantages to particular kinds of contents. We know, for instance, that comparatively different or unique items are better remembered later on (von Restorff, 1933), that uncompleted tasks are more likely to be recalled than completed ones (Zeigarnik, 1927), and that items at the beginning and end of lists are better retrieved than those in the middle (Ebbinghaus, 1885). Similarly, we know that people remember places and locations particularly well (O'Keefe & Nadel, 1978), that faces are better recalled than perceptually equivalent images (Homa, Haver & Schwartz, 1976), that pictures are remembered better than words (Paivio, 1971), and that real objects are better recalled than their pictures (Snow, Skiba, Coleman, & Berryhill, 2014).

Similar patterns occur when it comes to remembering times, or *temporal memory*.[3] Thinking back to that one job interview, you may notice that it is easier to bring to mind where it happened, the spatial layout of the office where it took place was, whether they gave you coffee, and even who was there and what were they wearing. But chances are, unless you've created some kind of meaningful conceptual association to the date of that particular event (e.g., if the job interview occurred on your birthday), you won't remember the date it happened. In a foundational paper, Neter and Waksberg (1964) found that participants were systematically biased toward thinking that past events took place more recently than they actually did—a phenomenon known as the "forward telescoping effect"—and that the size and frequency of the error varied as a function of the length of the recall period. A related bias, whereby people judge past events as occurring earlier than they actually did—the "backward telescoping effect"—has also been documented, although it seems to affect only recent memories (Janssen, Chessa & Murre, 2006). Accuracy for recalling temporal information also varies as a function of task and chronological granularity (e.g., years vs days). For instance, Burt (1992) found that while participants are quite good at assessing the relative temporal distance between events, their performance is quite poor when they were asked to retrieve the events' precise dates. Likewise, accuracy in time estimates decreases as retention intervals increase (Thomson et al, 1996), and relative to places or objects, dates are extremely bad mnemonic cues (Wagenaar, 1986; Thomson, 1982). However, the more episodic details one remembers about an event—where it happened, who was there, etc.—the more likely one is to accurately recall when it occurred. Extant evidence also suggests that some levels of chronological granularity, e.g., time of day or season, are better recalled than others, e.g., day or month (Friedman, 1987), and that a particular event is not a good retrieval cue for a temporally

---

[3] We use the term "temporal memory" to cover the encoding and retrieval of temporal information, be it chronological, sequential, and the like. We do not mean to imply that it refers to just one kind of memory. On the contrary, as we argue in what follows, temporal memory likely comprises different kinds of temporal information that depend on distinct cognitive processes.

contiguous one (Friedman & Huttelocher, 1997), unless they are causally, conceptually, or semantically related (Tzeng & Cotton, 1980).

To account for these and related findings, researchers have proposed a number of different views. Following Friedman (1993), we can distinguish three families of theories of temporal memory. First, there are *distance-based* theories, according to which temporal memory is correlated with, and depends upon, properties of the passage of time. One version of this view capitalizes on the linear structure of time and argues that experiences are organized in our long-term memory according to the order of their occurrence, and that the format of this organization is, in fact, spatial, so that we gauge the temporal recency or remoteness of a past event on the basis of how spatially distant the memory appears from the present point at the time of retrieval (Murdock, 1974). By contrast, other distance-based views hold that temporal distance information is to be found, not in the orderly format of our memory representations, but in qualitative properties of the mnemonic contents, such as their intensity (Guyau, 1890), strength (Anisfeld and Knapp (1968; Hinrichs 1970), ease of retrieval (Brown et al., 1985), or "sense of recency" (Glenberg et al., 1983). The underlying rationale was that, since all these features of our mnemonic contents tend to decrease as a function of time, temporal accuracy was thought to worsen as events become more and more distant in the past.

However, distance-based theories are challenged by evidence demonstrating that temporal memory does not honor the linearity of time. In an elegant study, Friedman and Huttenlocher (1997) asked viewers of the TV show "60 Minutes" to date stories featured in the weekly show between September of 1993 and May of 1994. The stories they chose bore no relation to contemporaneous events, in order to make it more likely for the participants' dating assessment to be done on the basis of remembered time alone. As a control condition, the researchers asked participants to date news stories, for which they would have, in addition to temporal memory, plenty of related knowledge upon which to base their temporal judgments. They found that, for stories from the previous six weeks, accuracy followed a somewhat linear pattern but was relatively poor overall. For stories older than six weeks, however, the linear relation breaks down, such that temporal distance no longer predicts accurate temporal recollection. This is troubling for strength theories in particular, which predict a strongly linear, inverse relation between the event's temporal distance from the time of recollection and accuracy. The authors also showed that there was no correlation between the amount of information participants could recall of a 60-minutes story and accurately assessing its date, casting doubt on distance-based theories based on the amount of propositional knowledge about the event (e.g., Brown et al, 1985). Finally, participants were much better at dating news events relative to stories, suggesting temporal distance alone is insufficient for accurate temporal memory.

The second family of theories are *order-theories*, according to which information about the time of occurrence of past events is stored, not in the memories of those events, but in the *associations* between the remembered events. Friedman (1993) discussed two variants of these theories. On the one hand, there is the *associative chaining theory*, according to which serial order learning does not occur because we encode a distinct position for each particular item or event, but because we encode pairwise associations between successive items/events (Lewandowsky & Murdock, 1989). On the other hand, there are *order code theories*, that expand upon associative chaining theory to argue that, in addition to encoding order, retrieval order can modify and update previously stored temporal information (Hintzman et al, 1985; Tzeng & Cotton, 1980). For instance, you may have learned about Hurricane Katrina before you learned about 9/11, but you were able to properly order these two US national emergencies by updating both your memories upon retrieval.

Yet, order theories also suffer from a number of shortcomings, not least of which is their limited explanatory scope, as they seem to accommodate only a relatively narrow set of findings. Associative chaining theory, for example, has a hard time explaining why our temporal memory is better, rather than worse, when items are not temporally contiguous. If our memory for time depends upon our recollection of the association between sequential items, one would expect memory for temporally contiguous items to be better than for temporally distant ones. The evidence, however, points in the exact opposite direction (e.g., Fozard, 1970; Guenther & Linton, 1975). Moreover, order theories predict that recalling temporal information should be better when cued by a temporally associated event relative to an event that is not temporally associated. However, extant findings suggest that recall is not improved when cued by an immediate predecessor relative to non-temporally related cues (Tzeng and Cotton, 1980). As mentioned before, time is a pretty lousy cue, but order theories presume the opposite.

Finally, the third family of theories are *location-based theories*. Unlike distance-based theories, in which temporal information of an autobiographical episode depends upon some changing property of the memory for the event, location-based theories assume no such change. Instead, they postulate that memory for time depends on information that is encoded and stored when the event is experienced. As such, they also differ from order-theories, as they accept that temporal information is to be found in the content of the memories per se, rather than in their associations. Friedman (1993) distinguishes three main kinds of location-based theories. First, there are *time-tagging theories*, which suggest that temporal information is one of the constituents of the encoded memory. More precisely, time-tagging theorists suggest that when an event is experienced and encoded, some sort of automatic time-stamping occurs so that the memory is tagged with its particular time of occurrence (Flexser & Bower, 1974; Hasher & Zacks, 1979). A second view, known as the *encoding perturbation theory*, suggests that in the course of encoding sequences of events, our memory marks certain pieces of information as control elements—junctures, if you will—that help to organize experience temporally by way of serving as reference points (Estes, 1985).[4] Finally*, reconstructive theories* hold that temporal information is *not* a constituent of the encoded content, but it is rather reconstructed—i.e., inferentially generated—at the time of retrieval out of stored episodic information from the encoded event as well as general knowledge (Hintzman et al, 1973; Friedman & Wilkins, 1985). Accordingly, when you remember that the old job interview took place on such-and-such day, you do so not because the memory comes to you wearing its date on its sleeve, as it were, but rather because your memory system enables you to infer, however approximately, the date of the interview on the basis of episodic information from the event you successfully encoded, stored and retrieved, in addition to related semantic information, both from general knowledge about the world as well as your own autobiographical past.

Location-based theories, however, do not enjoy uniform empirical support. For instance, time-tagging theories imply that the temporal assessment of episodic memories should not be affected by practice (Zacks et al, 1984), but evidence shows that performance can improve as a function of trial (e.g., Underwood & Malmi, 1978; Naveh-Benjamin, 1987). The automaticity

---

[4] Shum (1998) proposed a similar view for autobiographical memory, according to which autobiographical episodes are structured around temporal landmarks, the cuing of which helps to access specific autobiographical memories and, thus, improve accuracy in temporal memory. Shum's model is very congenial with the proposals of both Huttenlocher, Hedges & Prohaska (1998) as well as Conway and Pleydell-Pearce (2000), which we discuss below.

assumption of time-tagging theories is also at odds with the fact that people are more accurate at remembering the time of past events when they are told, during encoding, to memorize them, relative to when they are not told to remember them or to actively forget them. If the time-tags were automatic and depended upon involuntary processes, then we should not expect to see this difference (Michon & Jackson, 1984). Another concern with time-tagging theories is the lack of clarity regarding the temporal scale with which, allegedly, our memory system temporally tags experiences. Is it by the hour? The day? The week? As it turns out, the evidence suggests that there is considerable variation in accuracy across different time scales (Friedman, 1987; Thomson, 1982) and, interestingly, time scales that are best remembered—such as season of the year or time of day—are usually those for which participants report using inferential strategies, rather than the direct and immediate access at retrieval assumed by time-tagging theories (Friedman & Wilkins, 1985). In fact, as documented by Thomson and colleagues (1987), participants report remembering directly the dates and times of events only for a small proportion of remembered events—about 10%. Moreover, for this small subset of remembered times, memory accuracy was perfect, whereas for all other remembered times, accuracy was only approximate. In fact, differences in accuracy across time-scales puts pressure on the encoding perturbation-theory, as reference points do not seem to be fixed along a particular time scale. To be fair, the theory could potentially handle this difference in granularity by enabling different markers across temporal hierarchies. Doing so, however, turns the view into a hierarchical re-constructivist one, such as Huttenlocher et al's (1998) model, which incorporates reference period boundaries across different levels of the temporal hierarchy (i.e., years, moths, days) in order to predict systematic dating bias toward temporal boundaries at different levels of the hierarchy.

Of all the location-based theories, the reconstructive view enjoys the most empirical support (Friedman, 1993; 2004). According to it, if our goal is to remember when, in calendar terms, a past event occurred, we typically make use of inferential processes that latch onto stored episodic details about the encoded event as well as semantic information from public and autobiographical knowledge. For example, in trying to remember when a job interview took place, you may recall wearing a sweater and that there was a winter jacket hanging on a coat rack. These bits of episodic information help to constrain the search space by reducing the number of possible seasons to one: winter. Confirmatory evidence may also be provided by your semantic knowledge that interviews in your job market tend to occur in January and February, so the temporal window gets even narrower. Likewise, you may also remember light coming through the window, and perhaps the smell of coffee and freshly baked croissants, further suggesting that the interview might have taken place in the morning. Notice, though, that you may not be able to retrieve any further information that could narrow down the precise week or day, in which case your recollection will be approximately accurate at certain levels of granularity—e.g., the season, month, and time of day—but rather inaccurate at others—e.g., week, day—which is exactly what the evidence suggests (Huttenlocher et al., 1998).

In addition to inferential processes, we sometimes employ direct retrieval processes when voluntary semantic tagging has occurred at encoding, likely as a result of elaborative or deep encoding (Craik & Lockhart, 1972). This is what occurs with important events in our lives, such as the birth of one's child or winning the lottery. In both cases, in addition to all the episodic information that gets encoded while experiencing the events, there is also semantic information that becomes inextricably associated with the episode, enabling you to directly retrieve temporal information whenever you remember the autobiographical experience. Suppose, for instance, that right as you leave the job interview, you receive news of your parent's death. Presumably, the date becomes inseparably linked to your memory of hearing this news, and as long as you remember that the job interview occurred just before you heard the news, you will be able to directly access the precise date of the job interview. Direct retrieval of temporal memories due to this kind of semantic tagging is comparatively less common than reconstructive processes, and unlike them, it yields not approximate but exact dates (Thomson et al, 1987).

The two processes postulated by the reconstructive view—i.e., inference and semantic tagging—help to explain how we remember when events happened in our calendar past. However, more recently, Friedman (2004) acknowledges that not all kinds of temporal information from our memories can be explained by these two processes alone, and he suggests that two additional processes may be necessary. One such process is what he calls a "rudimentary sense of temporal distance" (Friedman, 2004: 599), in virtue of which we can recognize information as pertaining to a very recent past. This process, Friedman speculates, may have been an earlier adaptation that evolved for foraging, enabling our ancestors to discriminate areas that have been recently tapped from those that have not been explored—at least in a while. He further elaborates that this sense of temporal distance—which is not that different from the intensity or strength-based distance theories discussed above—explains why participants in the Friedman and Huttenlocher (1997) experiment were able to gauge with approximate accuracy the time of recent 60-minutes stories, but as the stories dated further back in time, their accuracy plummeted. This sense of temporal distance, however, ceases to be informative for events that are older than just a few weeks.

Finally, Friedman (2004) also suggests that information about the precise order in which specific events succeed one another within a remembered episode may also require a distinct process. Suppose that, during the job interview, one of the interviewers pulls out a folder with your résumé and realizes that it is missing a page. Thankfully, having anticipated the need to bring some extra copies, you immediately took one out from your bag and gave it to her—a gesture the hiring committee clearly saw with admiration. Now, when you remember the event, you know which action brought about your taking a copy of your résumé out of your bag, and you also know the judgment that your action brought about among the members of the hiring committee. You can, as it were, causally partition these actions into sequences and, as a result, recall which happened first and which followed. It is likely that the capacity for remembering the temporal order of actions within episodic memories depends, in part, on the fact that remembering an autobiographical episodic memory involves the mental simulation of a dynamic content that is causally structured (De Brigard & Gessell, 2017), as opposed to simple unconnected scenes. Incidentally, Friedman suggests that the process that allows us to extract temporal order from autobiographical memories helps to explain why we are better at remembering events that are causally connected relative to disconnected, even when you keep temporal distance constant (Zeng & Cotton, 1980).

3. Memory for imagined times

In the previous section, we summarily reviewed a number of theories that have been proposed to explain experimental findings from research on temporal memory. We sided with Friedman's assessment that, although typically only approximately accurate, people's memory for past times is nevertheless systematically biased in ways that can be accounted for by one of four temporal memory processes: reconstructive inferences, semantic-tagging, sense of recency, and causal order. In this section we suggest that these processes can also underwrite temporal memory for imagined events. Specifically, we conjecture that the processes implicated in remembering time for episodic memories also operate in remembering times in episodic future and counterfactual thinking (Ingvar, 1979; 1985; Szpunar et al., 2013). Initial support for this conjecture stems from the commonalities between the neural and cognitive mechanisms underlying episodic memory, future and counterfactual thinking mentioned in the Introduction.

To see how this may be the case, we start with recent developments on the well-known Self-Memory System (SMS) model of autobiographical memory (Conway & Pleydell-Pearce, 2000). According to this model, autobiographical memories are reconstructed out of semantic and episodic autobiographical information, both of which are hierarchically organized according to generality. The upper levels of the hierarchy comprise autobiographical knowledge—some of which comes from general semantic knowledge, and some of which comes from semantizised autobiographical information about ourselves. The top of the hierarchy includes *life themes*, such as "work" or "relationships". The next level down corresponds to *lifetime periods*, such as "Working at University X" or "Job at company Y". At the next level down there are *general episodes*, such as "Department talk" or "Graduation". Finally, at the lowest level, we find episodic information, particular pieces of mnemonic information of spatiotemporally located events within a general episode, such as the memory of the precise moment in which your graduate advisor hooded you, or the moment in which you received the phone call for a job promotion. Importantly, the SMS model distinguishes between *stored* and *transitory* representations (Conway, 2005; 2009). Stored representations are the pieces of mnemonic information from which the autobiographical memories are reconstructed, whereas transitory representations are the individual instances of reconstructed autobiographical recollections we consciously entertain as a result of interacting with a particular cue.

Thus, a second characteristic of the SMS model is that each and every instance of autobiographical recollection—i.e., each transitory memory representation reconstructed out of stored information—is caused by a *cue*. Pretty much anything can become a cue: a smell, an object, a word. Cues can interact with stored information at any level of the hierarchy, and then spread the activation of associated information to re-construct a transitory representation of an autobiographical memory typically in the form of an episodic mental simulation entertained in working memory. Although the spread of this activation need not be voluntary (Berntsen, 2012), most autobiographical recollections occur under our conscious control and in response to specific goals (more on this soon). Hence, the final component of the SMS model is a *central control process* that can access, supervise and direct the spread of activation of stored information elicited by the cue during the process of memory re-construction (Williams and Hollan, 1981). How exactly this process controls the spread of the activation remains unclear, but a likely candidate is the voluntary deployment of internal attention, which enables us to highlight and integrate pieces of stored information from long-term memory into a dynamically structured mental simulation consciously entertained in working memory (De Brigard, 2011; De Brigard and Gessell, 2017).

To illustrate the SMS model, consider again the example of remembering a job interview. The first thing to notice is that such recollections can occur in response to different cues and for different purposes, each of which will influence the pattern of activation that results in a transient representation of an autobiographical memory. For instance, you may smell a particular kind of coffee and then involuntarily reactivate stored information that results in the simulation of your memory of sitting at the interview table while being served a particularly strong cup of coffee. Here the cue interacts with the lowest level of the hierarchy, episodic information, and propagates activation upwards, so that the event is situated within a particular instance of job interview, which in turn is associated to a particular life period (e.g., post-doc), under a particular life theme (e.g., work). A cue can also interact with stored information at higher levels of the hierarchy to spread activation downwards. For example, someone may ask you if you remember a job interview in which coffee was served. Upon hearing this particular question, central control processes are recruited, enabling you to voluntarily search within the life theme "work" and general event "job interview" a specific episode in which coffee was served. The result is a reconstructed episodic autobiographical memory in which the content is constituted by stored semantic and episodic information, transiently represented in working memory as a conscious mental simulation.

Recently, the SMS model has been extended to cover not only episodic autobiographical memory but also episodic future thinking (Conway et al., 2019; D'Argembeau, 2020). According to these developments, the hierarchical structure of autobiographical information also extends to our capacity to project ourselves mentally to the future (Figure 1). As D'Argembeau (2020) insists, most research on mental time travel focuses on shared mechanisms at the level of episodic simulation, and the different theories that account for these communalities—e.g., the constructive episodic simulation hypothesis (Schacter & Addis, 2007), the scene construction theory (Hassabis & Maguire, 2007), or the event memory approach (Rubin & Umanath, 2015)—often say little about the general autobiographical context in which these episodic mental simulations are embedded. In the past few years, however, it's become clear that semantic memory is needed to "scaffold" the episodic components that constitute the contents of our episodic simulations (Irish & Piguet, 2013). But the precise nature of this semantic scaffolding remains unclear. Extending the SMS model toward the autobiographical future promises to clarify how semantic knowledge—both general and autobiographical—can interface with episodic simulations to generate not only recollections of past experiences but also thoughts about possible future personal events. More precisely, we can extend the hierarchical structure of SMS so that lifetime themes, periods, and general and specific events can be extended both backward and forward looking, depending on whether one is trying to remember an episodic autobiographical memory or to imagine an episodic future thought.

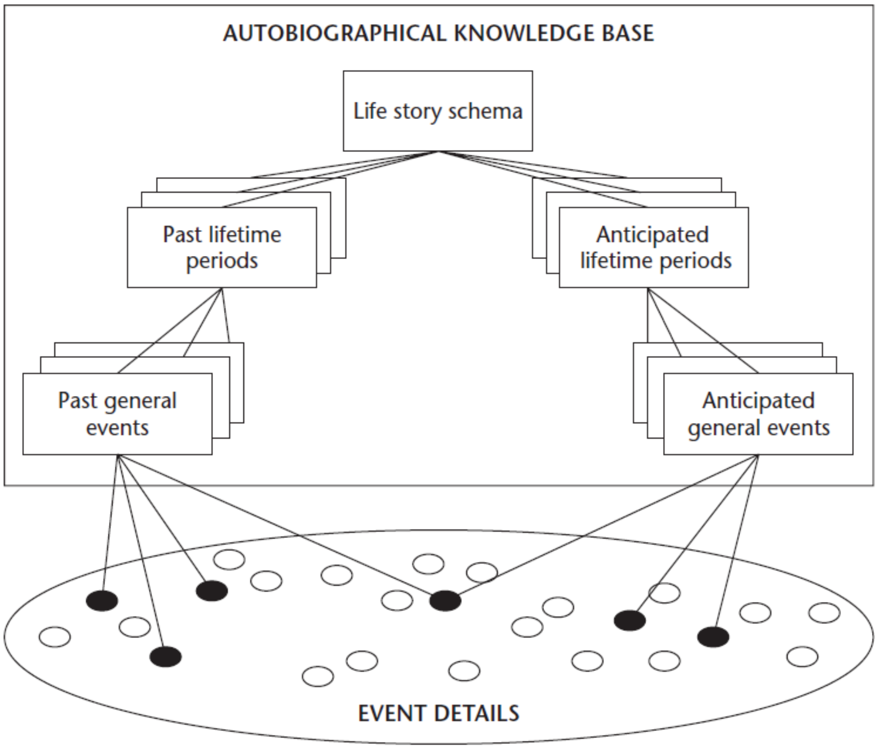

Figure 1. Extension of the SMS model for episodic future thought. Reprinted from Conway, Justice and D'Argembeau, 2019.

The degree to which a mental simulation integrates information across the hierarchy of autobiographical knowledge represented in the SMS framework depends on a number of factors. To illustrate, consider now an episodic simulation of a possible future job interview. Such simulation may occur, for instance, while you are involuntarily daydreaming about an unspecified future job market. In this case you'd be entertaining an episodic simulation of a possible future event that likely won't have much associations with higher levels of autobiographical knowledge, akin to a sudden recollection of an autobiographical memory you cannot precisely pinpoint in your personal past. Often these transient representations are fleeting, like snapshots that capture slices of possible experiences; albeit more frequently they are dynamic, depicting episodes that are causally structured and unfold over time (Anderson et al, 2015). Alternatively, you might voluntarily generate an episodic future thought as a result of a particular request—say, a friend asking you to imagine what it would be like to interview for your dream job. Here, as in the case of voluntary autobiographical recollection, the request (the cue) establishes a search goal for the central control process, which strategically constructs a transient representation—a mental simulation—out of episodic details from memory as well as semantic information from stored knowledge pertaining to general events, periods and lifetime themes. Importantly, the degree of complexity and detail of these episodic mental simulations is going to depend not only on the amount of time devoted to the imagistic elaboration but also on the amount of episodic details available for constructing the simulation and the degree to which it incorporates semantic information from other levels of autobiographical knowledge.

The consequences of the differences in degree of complexity and detail in episodic future simulations are varied, but for present purposes we will focus on those pertaining to memorability. Recent evidence indicates that, similar to episodic autobiographical memory, certain components of episodic future simulations are better remembered than others. Szpunar, Addis, and Schacter (2012) showed, for instance, that participants were better at recalling imagined persons relative to objects and locations, and that they were able to recall these details better when simulated as positive future events relative to negative ones. Similarly, McClelland and collaborators (2015) had participants construct episodic future simulations involving people, places, and objects randomly culled from their autobiographical memories, and asked them to rate them along a number of phenomenological dimensions. A surprise cued-recall test, administered a few minutes later, revealed that episodic future simulations constructed out of more familiar components and experienced as more detailed and more plausible were better remembered than episodic future simulations that involved less familiar places and people, and that were experienced as less detailed and plausible. More recently, Jeunehomme and D'Argembeau (2017) used free- and cued-recall tests to examine the accessibility and characteristics of recalled episodic future simulations (i.e., memories of the future). They found that, when tested a week later, participants freely recalled only about half of the simulations previously created, suggesting that not all mental simulations are equally accessible in memory. However, when properly cued, participants only forgot about 7% of their previous mental simulations. Persons and places were better remembered than objects and actions, and they were also better cues—consistent with evidence reviewed in Section 2 regarding their mnemonic advantage in autobiographical memory. Additionally, clarity and familiarity during mental simulation—particularly in relation to imagined people—predicted subsequent recollection. Finally, their results also show that when episodic future simulations are better incorporated into one's lifetime goals, and thus draw more on autobiographical knowledge of one's desires and intentions, they are overall better remembered than when they are not so easily incorporated (see also, Ben Malek, Berna and D'Argembeau, 2017; 2018).

But how is the temporal information of our memories for the future remembered? In a recent study, De Brigard, Gessell, Yang, Stewart and Marsh (2020) sought to explore precisely this question. They asked participants to recall specific episodic autobiographical memories, episodic future thoughts and episodic counterfactual thoughts featuring a unique person, place, object and time (month and year). In addition to recording each of these components for each of the three kinds of episodic simulations, they also provided ratings of vividness and valence. A day later, participants received a surprise memory test, in which three of the four simulation components were provided as cues, and their task was to recall the remaining one. Consistent with the aforementioned studies, location and person information was equally well remembered across the three kinds of mental simulation (Figure 2), whereas objects were better remembered in the memory relative to the future condition—a result consistent with the findings from Jeunehomme and D'Argembeau (2017).

Moreover, consistent with the evidence reviewed in Section 2, memory for dates was comparatively poorer than that for location, object, and person in the episodic autobiographical memory condition. Importantly, however, it was *disproportionately worse* for episodic counterfactual and future simulations than it was for episodic memories. Furthermore, a binomial logistic regression indicated that differences in vividness did not explain any variance associated with the poor performance of time information above and beyond condition. Taken together, these results suggest that temporal information (at least at the level of month/year dates) is more poorly recalled in imagined episodic simulations of possible past and future events than it is in recalled episodic autobiographical memories.

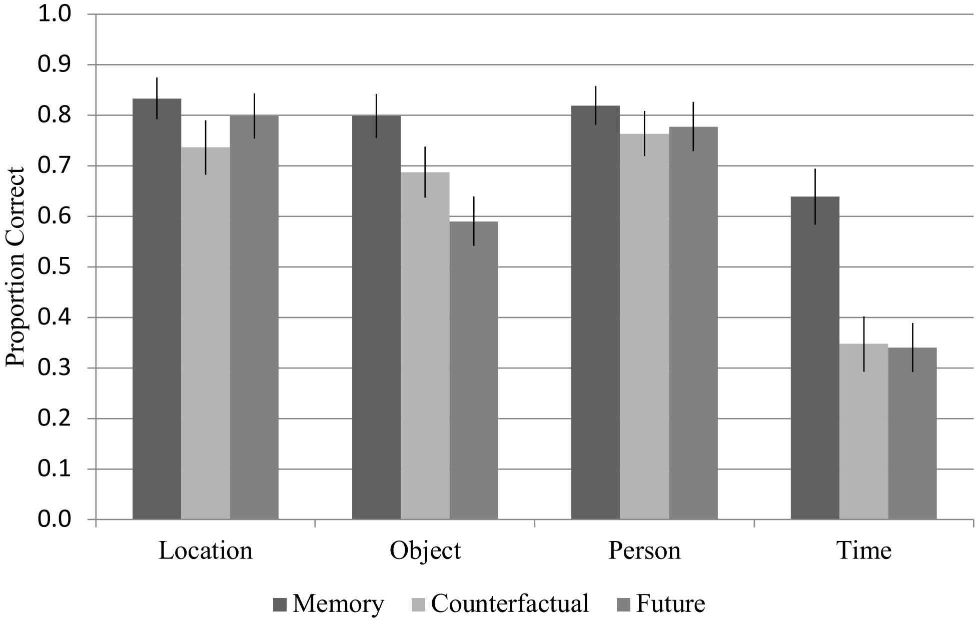

Figure 2. Proportion of correct responses for each episodic component as a function of condition. Error bars indicate SEM. (Reprinted from De Brigard et al., 2020)

How can we explain these results? According to the framework we are proposing here, accuracy for remembered times is higher in the memory condition relative to both the episodic future and counterfactual conditions for at least three reasons. First, the time component during the simulation of episodic autobiographical memories is probably better integrated with information across all levels of the SMS' hierarchy and, thus, it is much more likely to be encoded as part of a "memory for the future". Second, because it is better integrated during encoding, the inferential reconstruction of the temporal content at retrieval is more likely to be accurate, just as it is more likely that other episodic details are more accurately remembered in previously remembered as opposed to previously imagined events; this latter observation is evidenced by the fact that, on average, accuracy for memory details was higher relative to episodic future and counterfactual conditions (Figure 2). Third, since semantic temporal information is built into the hierarchical structure of the autobiographical knowledge base in the upper levels of the SMS, it can help to better constrain the inferential reconstructive processes participants employed—both consciously and unconsciously—at retrieval, which in turn may have helped them to improve their estimates of remembered times.

The same is not the case when it comes to the episodic future and counterfactual conditions. First, because in these two conditions there is less information in the autobiographical knowledge base relative to episodic autobiographical memories (Conway et al, 2019; D'Argembeau, 2020), the time component in the episodic future and counterfactual conditions isn't as well incorporated into the hierarchical structure of the encoded simulation and, as a result, it is less likely to be retrieved. Additionally, perhaps because participants weren't directly instructed to, it is unlikely that they attended to how the imagined events fit into their SMS model, which in turn decreased the likelihood that the time component would be directly retrieved by way of semantic associations with elements of the autobiographical knowledge base. Finally, because the information along all levels of the hierarchy—both in the autobiographical knowledge base as well as the level of episodic detail—for both episodic future and counterfactual simulations is sparser than for episodic memories, there are fewer inferential constraints operating on the reconstructive retrieval processes in those two conditions, which in turn increases the likelihood of inaccurately retrieving the time component.

Therein lies the nature of the asymmetry we alluded to in the Introduction. Recall that according to the constructive episodic simulation hypotheses, the same constructive processes underlying episodic memory retrieval are recruited for the construction of episodic future and counterfactual simulations (Schacter & Addis, 2007). In turn, this common recruitment is thought to be the result of a shared adaptive purpose: to generate episodic simulations that can be encoded as "memories for the future" in order to be retrieved when the right time comes (Szpunar et al., 2013). But we have now seen that not all previously simulated episodic contents are equally well remembered, and that while places and people tend to be better recalled regardless of the nature of the encoded simulation, the retrieval of temporal information (at least dates) is disproportionally worse for previously imagined relative to previously remembered episodic simulations. And while we offered a possible explanation about the psychological basis of this asymmetry, it is worth asking why there is such an asymmetry in temporal memory to begin with. One possibility is that temporal information—specifically, information about when the event occurs—is typically not as useful for improving future behavior as other kinds of episodic details, such as specific places, people, and objects. As such, when people engage in episodic hypothetical thinking, the temporal element—the *when*—of the simulated event is not as prioritized as other elements of the simulation, e.g., the *where*, *what*, *who* and *how* components. Another possibility, compatible with the previous one, is that our attention needs to be strategically allocated while encoding the simulated episode, so unless directly targeted, the temporal components of the simulated content, even if generated, won't be encoded. Needless to say, further research is needed to fully understand the psychological reasons behind this asymmetry.

Nevertheless, it is important to notice that not all temporal information is equally inaccurate. As it turns out, when De Brigard and colleagues used a much coarser accuracy measure on their data, by way of re-coding answers as correct as long as their temporal orientation was right (i.e., if past and counterfactual simulations were placed in the past, and future simulations were placed in the future), participants' dating accuracy was much higher, but not equivalent, across conditions—being much better for episodic memories (98%) relative to episodic future thoughts (89%, $t(29) = 2.60$, $p = .038$, corrected). This suggests that even though people may not accurately recall the specific month and year in which they placed an episodic mental simulation, they can still discriminate whether their imaginations were oriented toward a possible future or a possible past. Why is memory accuracy better for the coarser measure of temporal direction relative to the more fine grained measure of date? While this is still an open question that none of the existing data can speak to, one intriguing possibility is that there may be a "sense of direction" process, tantamount to Friedman's (2005) proposed "sense of recency", that imbues our episodic simulations with a very general feeling of temporal orientation. Clearly this is early a conjecture, but it is one we think may merit further research.

What about other effects that may follow from combining the SMS models and the temporal memory mechanisms suggested by Friedman? If the framework we are proposing here is on the right track, it follows that temporal information more strongly integrated with a particular episode across the hierarchical structure of our autobiographical knowledge would be better remembered than temporal information that is not so well integrated. Since we expect temporal information to be better integrated in episodic memories than episodic future thoughts, we anticipate that the size of dating errors would be much smaller in the former than the latter. Moreover, we should expect to see an increase in error size with more temporally distant episodic future thoughts relative to closer ones, as the latter involve more familiar, concrete and detailed simulations than the former (Arnold et al, 2011). Episodic counterfactual thoughts, by hypothesis, should probably fall in the middle, as they share quite a bit of structure with episodic memories but, unlike them, they still involve marked deviations from the actual past.

To test this hypothesis, we re-analyzed data from De Brigard et al. (2020), and evaluated the size of the error as a function of temporal distance. Specifically, we conducted a linear mixed effect regression of temporal distance and type of episodic thought on temporal error, with random intercepts for participants (more complicated random effect structures resulted in non-convergence or singular fits). As predicted, this analysis revealed a significant interaction between temporal distance and episodic future thoughts, $\beta = -0.21$, 95% CI = (-0.32, -0.10), SE = 0.06, $t(411) = -3.75$, $p < 0.001$, such that error size was larger for more distant imagined future simulations relative to temporally closer ones (Figure 3). By contrast, there was less error dispersion in recalled dates for remembered times. In the case of episodic counterfactual thoughts, while they were numerically more dispersed than date in the episodic memory condition, the results were not statistically different from it ($\beta = 0.07$, 95% CI = (-0.02, 0.17), SD = 0.05, $t(409) = 1.45$, $p = 0.147$). While this is of course not conclusive evidence of a forward-looking telescope effect for temporal memory in episodic future thoughts, we believe it invites further research on how memory may distort temporal judgments in recalled episodic simulations.

The evidence (and the theory) reviewed so far, only begins to scratch the surface of what clearly constitutes a much more complex area of research: memory for temporal information in episodic simulations. Nevertheless, we believe that by combining the resources of Friedman's mechanisms for temporal memory with D'Argembeau's recent developments of SMS on mental time travel, we may be able to make progress in our understanding of how people retrieve temporal information from their "memories of the future".

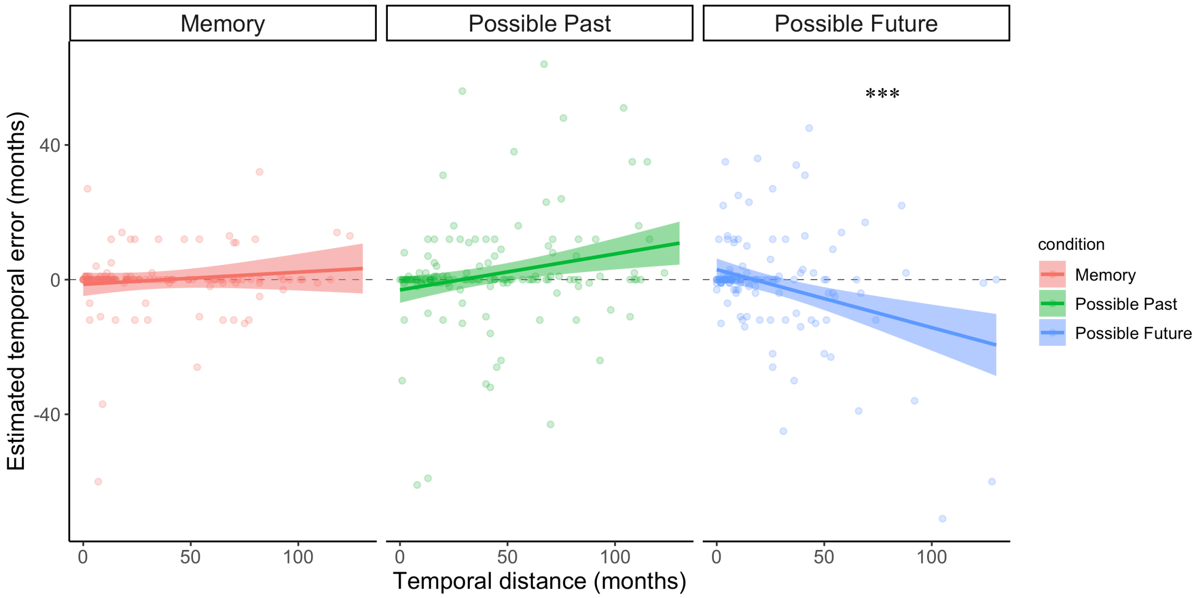

Figure 3. Plots of temporal error size as a function of temporal distance per episodic thought condition. Dashed line at zero indicates the present (month when participants were tested). Shaded regions represent 95% confidence intervals of model estimates. *** indicates a *p* value <.001. Data from De Brigard et al, 2020. Available at https://github.com/IMC-Lab/MemTime.

4. Improving our memory for imagined times

In the previous section we reviewed recent developments aimed at expanding the SMS model from episodic autobiographical memory to episodic future thinking. We also suggested that Friedman's approach to understanding temporal memory could be expanded to cover 'memories for the future' within the context of the extended SMS. In brief, the suggestion is that the hierarchical structure that constrains temporal recollection in autobiographical memory is tantamount to that postulated by the SMS, and thus can be extended to episodic future (and likely counterfactual) simulations. As a result, individuals should be better at recalling temporal information from episodic mental simulations that are more integrated with the hierarchical representations provided by our autobiographical knowledge as well as more strongly associated with other episodic details. Likewise, we hypothesize that mental simulations in which careful attention is paid to the causal structure of the imagined events should yield better order information than fleeting simulations of possible future snapshots in which no dynamical unfolding of events is simulated. Relatedly, we conjecture that attending to different aspects of the episodic simulation at encoding should yield different effects on temporal memory at retrieval. For instance, attending to particular details of the episodic simulation such as the clarity of the day or the weather, would help to encode certain temporal information, such as the time of day or the season. Likewise, enforcing semantic tagging during encoding may help to improve one's recollection of temporal information that isn't easily inferred from the simulated contents alone, such as the day of the week or month of the year.

To conclude, we suggest that these insights can be leveraged to use temporal memory to potentially improve planning and decision-making. Many aspects of our agency require memory for time: we are expected to remember birthdays and anniversaries, we deliberate about whether to spend money now or save for later, and we resolve to fit more time for exercise into our daily schedule (among other things). Doing these things requires deploying temporal memory correctly. Notably, failure to properly deploy temporal memory can lead to mistakes or missed opportunities: failing to call a friend on her birthday, being unable to buy a new computer, or leaving the pie in the oven too long all reflect failures to remember time. However, the reconstructivist framework outlined above indicates several empirically tractable strategies that can be used to bolster temporal memory and, hopefully, avoid some of these characteristic mistakes. In particular, we will focus on applications to three phenomena: temporal discounting, prospective memory, and implementation of intention.

Most people exhibit a bias toward preferring rewards in the near future compared to the distant future, even when the rewards in the distant future are larger than those in the near future (Ainslie, 2001; Frederick, Loewenstein, & O'Donoghue, 2002). This is known as *temporal discounting*. Oftentimes, temporal discounting results in suboptimal choices, for example, spending too much money on a fancy dinner when one needs to save money to buy a new computer for work (Ainslie & Haslam, 1992; though see Brink, 2010 for a modest defense of some discounting). Temporal discounting often rears its head during deliberation: it's when deciding whether to go out or stay in that you weigh the relative merits of either option. Because deliberation sometimes utilizes mental simulation, it is unsurprising that improved episodic future thinking attenuates temporal discounting (Daniel et al., 2013; Kwan et al., 2015; Sasse, Peters, Büchel, & Brassen, 2015). This might be because when one has more detailed representations of rewards in the distant future, it is less likely that one will discount them relative to rewards in the near future; Benoit et al., 2011; Palombo et al., 2015). The framework proposed here posits a potential explanation for this: developing representations of future rewards that are more integrated with hierarchical representations of biographical knowledge might facilitate better intertemporal choice. Moreover, temporal memory might be used to thwart temptation that induces temporal discounting. An important aspect of persistence is calling to mind one's commitment to pursuing the long-term reward rather than the short-term reward. Stronger temporal information (i.e., highly integrated, fine-grained content about the timing of experiencing the long-term reward) might facilitate better recall of one's intention to pursue the long-term reward, thereby undercutting the force of the tempting stimulus. When forming commitments to pursue long-term goals, especially those that are likely to waver in the face of temptation, it could be beneficial to rehearse temporal information encoded in the future simulation so as to strengthen recall at later times.

When we form plans, we need not continue to think about them right up until the moment of action. Plans are useful because we can form and store them without consciously attending to them, thereby freeing up resources to pursue more locally salient activities. Thus, after forming and settling on a plan, we often store those plans and recall them at the right time (Cohen et al., 1996). For instance, if you plan to pick up milk on your way home from work, you don't need to think about getting milk for the entire workday. Instead, you store the plan in long-term memory and use associative cues to trigger recall of the plan at the appropriate moment (Braver, 2012). The capacity for prospective memory consists in the network of processes implicated in storing and recalling plans (McDaniel & Einstein, 2007). While prospective memory enables us to form many complicated plans without taxing our limited cognitive architecture, it is also susceptible to interference and failure. If you are too distracted with a pressing problem at work, you'll likely miss the store and arrive home empty-handed. The aforementioned evidence related to the SMS model suggests that some prospective memory failures might result from how people structure their plans. Memory for persons and places tends to be recalled better than objects and actions. Hence, formulating a plan as "Get milk" might be more susceptible to interference and goal neglect than a materially equivalent plan of "Go to the store". In general, then, when relying on prospective memory, the SMS model predicts that formulating plans in terms of persons and places are less likely to be neglected than those formulated in terms of objects and actions (holding fixed, of course, other relevant factors such as semantic tagging). Additionally, some temporally extended actions have causal structures that make them difficult to maintain in working memory. For example, Nigella Lawson's (2017) recipe for Turkish eggs requires whipping yogurt, poaching eggs, and browning butter all within a short amount of time. Because the components should be combined as quickly as possible, Lawson recommends executing these steps in tandem. However, each component is complicated, and the time-sensitive nature of the activities means that attention cannot be devoted to any particular component for very long. In this kind of situation (whether in the kitchen or elsewhere), our framework suggests that simulating the dynamic causal structure of the components and their relations to each other should facilitate better performance relative to someone who memorizes instructions in list-form.

Finally, some of the plans that we form are open-ended and lack well-defined terminal states. For instance, someone might resolve to exercise more or learn a new instrument. For these open-textured plans, it is useful to form *implementation intentions*, or precise intentions that automate action control in goal-relevant ways. An implementation intention related to exercise might take the form of an 'if-then' statement, such as "If I have free time in the morning, then I will spend 30 minutes on the treadmill" (see Wieber & Gollwitzer, 2010). Forming such implementation intentions requires mental simulation to understand what kinds of antecedent conditions should function as signals to perform the goal-relevant action and what such conditions might look like. Notably, evidence about the strength of particular cues from De Brigard et al. (2020) is directly relevant to thinking about how to form such implementation intentions, as those with temporal cues are unlikely to be as effective as those with location, object, or person-based cues. Consider, for example, the implementation intention mentioned above. Successful use of the intention requires noticing that one has free time, which triggers the intention to use the treadmill. However, such temporal cues are likely less effective. resulting in higher rates of goal neglect when implementation intentions are formulated with them. While not every implementation intention takes temporal information as part of its antecedent, there are times when it is tempting to do so (e.g., "I'll log off in 5 minutes", "I'll make the reservations next week", etc.). One consequence of this framework is that these implementation intentions are likely less effective than materially equivalent counterparts that are formulated in terms of locations, objects, and places (e.g., "I'll scroll until I see a dog picture"). Moreover, as with prospective memory, the more integrated the content of some implementation intention is with one's autobiographical hierarchy, the stronger recall of the intention is likely to be.

To be clear, although empirically testable, our thoughts in this section are mere speculations. However, if on the right track, they suggest that the reconstructivist framework of temporal memory, combined with the extension of the SMS model to episodic future thinking, can help us to understand *how* we structure our mental simulations bears on their effectiveness in guiding future conduct. Future research exploring how temporal information is encoded and retrieved from episodic past, future and counterfactual simulations may help us to better understand, not only the nature of our “memories for the future”, but also how we can better use them to improve future performance.

References

TBD