Word counts

Short abstract: 98, Long abstract: 234, Main text: 16,190, Figure captions: 586,

References: 3,086, Total: 20,271

Thinking in and about time: A dual systems perspective on temporal cognition

Christoph Hoerl, Department of Philosophy, University of Warwick.

Email: c.hoerl@warwick.ac.uk URL:

https://warwick.ac.uk/fac/soc/philosophy/people/hoerl/

Teresa McCormack, School of Psychology, Queen's University Belfast.

Email: t.mccormack@qub.ac.uk URL: https://pure.qub.ac.uk/portal/en/persons/teresa-mccormack(024a3045-accf-42b9-8eba-65bb8b2bf6a7).html

Short abstract

We outline a dual systems approach to temporal cognition, which distinguishes between a *temporal updating system* and a *temporal reasoning system*. We argue that the former is both phylogenetically and ontogenetically more primitive, but also that both of them are still at work alongside each other in adult human cognition. We use this distinction to interpret findings in comparative and developmental psychology – arguing that neither animals nor infants can think and reason about time – and also to explain certain features of adult human cognition discussed in the philosophy of time and in the literature on intertemporal choice, respectively.

Long abstract

We outline a dual systems approach to temporal cognition, which distinguishes between two cognitive systems for dealing with how things unfold over time – a *temporal updating system* and a *temporal reasoning system* – of which the former is both phylogenetically and ontogenetically more primitive than the latter, and which are at work alongside each other in adult human cognition. We describe the main features of each of the two systems, the types of behavior the more primitive temporal updating system can support, and the respects in which it is more limited than the temporal reasoning system. We then use the distinction between the two systems to interpret findings in comparative and developmental psychology, arguing that animals operate only with a temporal updating system and that children start out doing so too, before gradually becoming capable of thinking and reasoning about time. After this, we turn to adult human cognition and suggest that our account can also shed light on a specific feature of our everyday thinking about time that has been the subject of debate in the philosophy of time, which consists in a tendency to think about the nature of time itself in a way that appears ultimately self-contradictory. We conclude by considering the topic of intertemporal choice, and argue that drawing the distinction between temporal updating and temporal reasoning is also useful in the context of characterising two distinct mechanisms for delaying gratification.

Keywords

animal cognition, cognitive development, dual systems, mental time travel, metaphysics of time, planning, intertemporal choice, temporal reasoning, tense

Thinking in and about time: A dual systems perspective on temporal cognition

Is temporal cognition a basic, primitive ingredient in mental life, or is it a complex achievement requiring a great deal of cognitive sophistication? On the one hand, humans and animals are by necessity adept at timing their actions appropriately, at ensuring that sequences of actions unfold in the correct order, and at keeping track of changes in the environment that occur with temporal regularity. In that sense, both people and animals are inherently temporal creatures. On the other hand, the nature of time remains the subject of highly technical debates amongst metaphysicians and theoretical physicists (Bardon, 2013; Carroll, 2010), and different cultures have different systems for marking time, acquiring which requires explicit teaching and occurs over a protracted period of development (Aveni, 1989; Friedman, 1982; McCormack, 2015). In these respects, thinking about time is something that seems very difficult.

It is by no means obvious that being able to think about time is something that is ontogenetically or phylogenetically primitive, despite animals' and infants' prowess at processing a variety of types of temporal information. The idea that animals are cognitively "stuck in time" has a long-standing history (Aristotle, 1930; Bergson, 1911), and has been the subject of a considerable amount of research (for discussion, see Clayton, Bussey, & Dickinson, 2003; Roberts, 2002; Suddendorf & Corballis, 2007; Zentall, 2005). Despite convincing evidence that animals can retain information about things that they have experienced in the past (e.g., Babb & Crystal, 2005; Clayton & Dickinson, 1998; Eacott, Easton, & Zinkivskay, 2005) and act in ways that prepare them for situations that are yet to come (e.g., Mulcahy & Call, 2006; Osvath

& Osvath, 2008; Raby, Alexis, Dickinson, & Clayton, 2007), there is still widespread disagreement about how to interpret this evidence. Similarly, the idea that children have limited ability to think and reason about time has a long-standing history (Fraisse, 1964; Piaget, 1969). Infants can process a variety of types of temporal information even from birth (e.g., de Hevia, Izard, Coubart, Spelke & Streri, 2014), toddlers are adept at learning about event order (Bauer, Hertsgaard, Dropik & Daly, 1998), and the use of tense in language typically appears very early (Weist, 1989). Nevertheless, as we shall discuss, there are good reasons to believe that it takes several years before children can think about time as adults do.

In this paper, we outline a *dual systems* approach to temporal cognition and argue that it is useful not just for framing issues in comparative and developmental psychology but also in considering aspects of adult human cognition. We recognize that there is considerable debate about how the claims of such accounts should be interpreted, in particular whether it might be more appropriate to refer to dual processes rather than dual systems (Evans & Stanovich, 2013). We will not address these debates here, but instead identify four key claims that provide the basis for conceptualizing the distinction we want to make as a distinction between two systems: We claim (i) that one of these systems is less ontogenetically and phylogenetically primitive than the other, (ii) that one depends on experience and learning in a way the other does not, (iii) that one typically involves more cognitively effortful reasoning than the other, and (iv) that they co-exist and can potentially be in conflict, yielding contradictory beliefs or judgments. This set of claims is similar to sets of claims made in the context of other 'dual systems' accounts (Kahneman, 2011; Sloman, 1996; Smith & DeCoster, 2000), and for this reason we believe it is useful to adopt the same terminology. Note, however, that in distinguishing between two systems we are not claiming that they have distinct and discrete neurological bases or operate entirely independently, and we acknowledge that there may be other ways of describing the distinction we make (e.g., as two sets of processes).

Our dual systems approach to temporal cognition distinguishes between a *temporal updating* system and a *temporal reasoning* system. Abilities that have been studied under the heading of temporal cognition include a sensitivity to temporal duration, a sensitivity to repeating temporal periods, ways of keeping track of temporal order, and the ability to judge where in time events are located. What these diverse abilities have in common is that they are all used to solve tasks that involve things unfolding over time in a certain way. Such tasks may differ in ways that mean that they each require a somewhat different explanation of performance, e.g., because they involve different timescales (e.g., seconds versus years), or because they require a sensitivity to duration rather than succession. The distinction between temporal updating and temporal reasoning is intended as a more fundamental distinction that cuts across these differences. There is a basic reason why we want to make this distinction: we believe that there are a number of different tasks that involve things unfolding over time in a certain way that can be solved without the ability to represent and reason about time itself. The temporal updating system is sufficient to solve such tasks, whereas other tasks require the temporal reasoning system. We will first provide an outline characterization of each of the two systems and describe the types of behavior that the more primitive temporal updating system can support, but also describe its limitations. After this, we will use the distinction to interpret findings from the comparative and developmental literatures. To anticipate, we will argue that neither animals nor infants can think and reason about time – they rely entirely on the temporal updating system, although in the case of children there is also an important developmental story to be told about the emergence of the temporal reasoning system. We will then turn to adult temporal cognition and suggest that our account can also shed light on a specific feature of adults' everyday thinking about time that has been the subject of debate in the philosophy of time. We will conclude by considering the phenomenon of intertemporal choice and outlining a way in which the distinction between temporal updating and temporal reasoning bears on existing discussions of that phenomenon too.

1. The two systems

1.1 Temporal updating

A creature capable of what we call temporal updating maintains a representation of how things are in its environment, which can be conceived of as map-like in so far as it contains information about locations, but which can also contain information about the existence of objects and features of those objects whilst leaving the location of those objects underdetermined. As things change over time, the creature will receive new information, and this information may contradict an aspect of its existing representation. But all the creature does in response to receiving such information, using the updating system, is change the relevant aspect of its representation of the environment. That is to say, crucial to the temporal updating system is that it deals with changing input by *changing representations*, rather than by *representing change*. If a change happens, it simply records the new, changed state of affairs, rather than also representing that things were previously different from how they are now. Thus, importantly, the temporal updating system operates with a model of the world that concerns the world only ever as it is at present. Other times, and how things are at other times, are not represented in the model at all.

We should stress that when we say that the model of the world used in temporal updating is a model that concerns the world only as it is at present, this is compatible with it being based on information the creature has gathered over time. It can therefore also include information about features of the world that are currently outside the creature's sensory scope, but which the creature has learnt about – e.g., about features of locations it does not currently occupy, or about the existence of certain objects. There is even a sense in which this provides for a primitive way of representing goal states: The creature can think of certain items represented as existing in its environment as desirable, and of certain locations in that environment as good locations for doing certain things, and it can respond to these opportunities its environment might afford. This would allow it to act in ways that *de facto* prepare it for future encounters with those items or locations, even though other times are not represented in its model of the world as such.

In order for the temporal updating system to work, the creature's model of the world has to maintain information over time and also update as new information is received. This updating will result in changes to the creature's world model; these include changes about, for example, what objects are at which locations, but also about what is or is not desirable or what the right thing to do is at certain places. Thus, some updating will happen as a result of changes that occur in the creature's environment, independent of the creature itself, but some will happen because of changes in the creature's motivational state (we discuss an example of the latter in section 2.3).

Certain parameters will govern these maintenance and updating processes, and we assume that these parameters will vary considerably in a context- or task-dependent way. For example, some types of information may be maintained for only very short periods of time and quickly be lost or updated, consistent with the idea of the contents of a working memory store changing dynamically, whereas others might be maintained for lengthy periods and be resistant to change. The types of mechanisms involved in making updates to the creature's model of the world in response to changes purely extrinsic to the creature are also likely to be quite different from those involved when the changes are changes in the animals' motivational state. What these mechanisms are, though, and what parameters govern their operations are empirical questions, and none of the arguments given below require making any particular assumptions about them, other than the basic assumption that information about objects and locations can be maintained even when they do not remain within sensory range.

1.2 Temporal updating and behavior

What follows is a description of some of the behaviors that we believe the temporal updating system can support. In each instance, we also specify what we take to be the related limitation, i.e., what a creature cannot do if it is operating only with a temporal updating system.

(i) **Single-trial learning.** A creature capable only of temporal updating can acquire new information about the world from even a single learning episode and change its model of the world accordingly. This information can be of a variety of types, combined in various ways (e.g., information about the spatial location and nature of an object), and it can be held in memory without the creature continuously making use of it (e.g., it might only do so at a later occasion after being cued). **Limitation.** Making use of information acquired in the past, even if it stems from just one learning episode, is separate from, and does not require, representing that information as stemming from the past, which is something a creature capable only of temporal updating cannot do.

(ii) **Elapsed-time sensitivity.** A creature capable only of temporal updating can nevertheless be sensitive to how much time has elapsed since a certain event happened. Aspects of its model of the world might have a "shelf life". That is, after incorporating a new piece of information of a certain type into its representation of the world (e.g., where some food is located), it might then only store that information for a certain amount of time and as a result its representation will change yet again at a later point in time. This could be governed systematically by an interval timer. Once a certain amount of time has elapsed on such a timer, the creature would no longer operate with a model of the world that includes the relevant piece of information (see section 2.1 for more detail). **Limitation.** There is no sense in which the creature need be representing how long ago it obtained the relevant piece of information the representation of which is governed by a timer. This piece of information simply is or is not included in the creature's model of the world as a function of how long ago it was obtained.

(iii) **Sequential learning.** The type of sensitivity to elapsed time we have just characterized involves a process governing what happens to elements of a creature's model of the world over time, without temporal relations being represented within the model itself. Similar processes might also explain certain basic forms of sequential learning. A creature might become sensitive to the temporal order in which certain kinds of sequences unfold by acquiring a routine for updating its model of the world in that order, rather than that order being represented in the model itself. **Limitation.** A signature limit of the temporal updating system is that the correct functioning of the system depends on the creature receiving information about changes in its environment in the same order in which those changes happen; it will produce errors in situations in which these two orders come apart.

(iv) **Anticipation.** We can distinguish between at least two ways in which such a creature might produce behaviors that serve to prepare for future states of affairs. First, the creature may possess some sort of temporal sensitivity whereby it behaves in a certain way when a phase timing system is in a certain state (e.g., turn up at a certain location at a certain time of the day), thus enabling it to behave in a way that yields future benefits. Second, as we said, such a creature may have a primitive way of representing goal states, by representing items existing in its environment as desirable, or locations in that environment as good locations to do certain things. This may cause it to act in ways that are optimal given certain possibilities its environment may afford (discussed further in section 2.2). **Limitation.** If they are not immediately accessible, the items represented as desirable (or the locations represented as good locations to do certain things) will *de facto*, at best, be encountered by the creature at some point in the future. However, the creature can represent these items (and locations), and respond to their presence in its environment, without having the means to represent its potential encounters with them as events occurring at a separate, future point in time.

1.3 Temporal reasoning

The key difference between a creature capable only of temporal updating and one that is able to engage in what we call temporal reasoning is that the latter operates with a model of the world that includes a temporal dimension. That is to say, the model contains addresses for different points in time, and can therefore contain information not just concerning the world as it is at present, but also information about states of affairs different from those obtaining in the present, which obtained in the past or will obtain in the future. Here we summarize what we take to be the most fundamental kinds of representational resources this involves.

(i) **Representing particular times:** A creature capable of temporal reasoning can represent events as happening at particular times, each of which only comes round once. Creatures not capable of temporal reasoning, by contrast, whilst capable of becoming sensitive to repeated temporal patterns in their environment (e.g., processes that display a circadian rhythm), will not distinguish between individual instantiations of these patterns as distinct unique occurrences (Campbell, 1996).

(ii) **Representing temporal order.** Because a creature capable of temporal reasoning can represent the temporal order relations between events happening at different times, it can use information about this order to arrive at a correct model of the world. It is not restricted to arriving at a correct model of the world only if it receives information about changes in its environment in the same order in which those changes occur.

(iii) **Tense.** A creature engaging in temporal reasoning is also capable of using the system of particular times as a framework for orienting itself in, by using *tense*, i.e. locating events in the past, present or future. By contrast, the representations entertained by a creature capable only of temporal updating are tenseless, or untensed. Its model of the world concerns the world as it is at present not because items in it are represented *as present*, but simply because it is the model entertained at present.

Having briefly outlined some of the key features of both the temporal updating system and of temporal reasoning, we will turn to considering ways in which the distinction between these two systems might be relevant to the interpretation of existing empirical research. Before doing so, we add two brief clarificatory comments.

First, we assume the existence of particular timing mechanisms, which can explain how even a creature capable only of temporal updating might nevertheless display forms of behavior that are sensitive to elapsed time. Timing mechanisms of some sort are widely assumed to be available even to basic creatures (e.g., insects, Bradshaw & Holzapfel, 2010); we remain neutral on their nature (see Grondin, 2010). Clearly there are important further questions about the precise ways in which these mechanisms operate, and the limitations they are subject to. It is important to note, however, that our suggestion is not that there is one form of cognition – temporal updating – that relies on the existence of mechanisms of this type, and another one – temporal reasoning – that does not. Mechanisms keeping track of time can obviously also be involved in contexts in which individuals make explicit judgements, e.g., about how long ago a certain event occurred (Friedman, 2001). The issue at stake concerns the function the relevant mechanisms play as part of the two systems: In one case, they simply govern the updating and maintenance of elements of the creature's model of its present environment; in the other, they ground a representation of a temporal interval extending into the past (see section 2.1).

Second, as we characterize it, the difference between temporal updating and temporal reasoning is fundamentally concerned with what a creature can represent – i.e. whether or not its model of the world contains a temporal dimension. In characterizing the distinction in this way, we have assumed that not all of the ways in which a creature might be sensitive to aspects of its environment involve that creature actually representing those aspects. That is, we are operating with a notion of representation that distinguishes representing an aspect of the world from simply being sensitive to it (see also Peacocke, 2017). This has to be distinguished from a broader notion of representation such as that involved when lower-level brain mechanisms are described as operating on 'representations' (for instance, one might describe the early visual system as 'representing' the differences between the two retinal images, but this is clearly not part of what is visually represented in the viewer's perception). There has been considerable debate on how exactly to characterize the relevant difference between these different notions of representation (for some influential early discussions, see, e.g., Dennett, 1969; Stich, 1978). For present purposes, we want to emphasize that the distinction we are drawing between merely being sensitive to temporal features of the world and representing time is not just a terminological one. If a creature can do the former but not the latter, this has concrete behavioural implications. For instance, as we have explained, the correct functioning of the temporal updating system is dependent on the creature receiving information about changes in its environment in the order in which these changes happen. This is a signature limit of the temporal updating system that the temporal reasoning system is not subject to (see section 3.1 for further discussion).

2. Are animals capable of temporal reasoning?

One type of debate that might be reframed by adopting the approach we are advocating is the debate about the existence of capacities for "mental time travel" (MTT) in non-human animals. Though this issue has received a great deal of attention (Roberts, 2002; Suddendorf & Corballis, 2007; Zentall, 2005), existing debates have reached something of an impasse, arguably because of the way in which some researchers have framed the basic dialectic. One strand of debate has considered whether MTT should be defined in information processing terms, or in terms of a particular kind of conscious awareness (Clayton, Bussey, & Dickinson, 2003; Tulving, 2005). Yet, if the issue is entirely about possessing a type of conscious awareness, the question of whether animals are capable of MTT becomes empirically intractable. Similarly, researchers disputing the existence of MTT capacities in animals often appeal to the operation of low level associative mechanisms that might be sufficient to explain the relevant behavior in each case (e.g., Redshaw, Taylor, & Suddendorf, 2017; Suddendorf & Corballis, 2010). Yet, whilst such alternative explanations might in principle be available in each case, dealing with individual findings in this sort of *ad hoc* way seems unconvincing.

We believe that the distinction that we are drawing, between temporal reasoning and temporal updating, provides a more helpful alternative to the dichotomies in play in these existing debates. Our distinction is one between two different systems of cognition, rather than one between cognition and mere low-level association. In describing animals as capable of temporal updating only, we assume that it nevertheless makes sense to talk about them as operating with a model of the world, and indeed a model that represents objects or places currently outside their sensory scope. We now discuss three empirical paradigms that have been used to make the case for MTT abilities in animals, and explain how animals might show the types of behavior in question even if they are restricted to mere temporal updating. We have chosen these paradigms because they are generally regarded as making the strongest empirical case for MTT capacities in animals. Some of them have been tested on more than one species, but only corvids have so far been shown to be successful in all three of them. We therefore concentrate on the relevant studies with corvids. Since our argument is that these studies are not able to establish whether animals are capable of temporal reasoning rather than just temporal updating, we will also consider later (section 3.3) what sort of alternative empirical paradigms, not yet tested on animals, would be able to establish this.

2.1 "What-When-Where" Memory

The most influential paradigm that has been taken to measure mental time travel to the past in animals is that of Clayton and Dickinson (1998). These researchers aimed to demonstrate that scrub jays can remember three key pieces of information about past events: what happened, where it happened, and when it happened. For present purposes, it is the last of these pieces of information that is crucial. Representing that something happened at a particular point in the past is an instance of temporal reasoning. Thus we need to ask whether these studies provide good evidence that animals can do so.

In the original study, what was taken as evidence that birds can remember “when” information is the fact that which cache site they return to is appropriately delay-dependent. As shown in Figure 1, birds that learned that worms degraded over a period of 124 hours did not return to a cache of worms if given access to it after a 124-hour delay, but instead returned to a cache of non-degrading peanuts. By contrast, if the delay was just 4 hours, birds returned to the cache of worms (their preferred foodstuff) rather than peanuts (see also Clayton, Yu, & Dickinson, 2003). To think that this study measures MTT is to think that the birds can remember the event of caching the worms, and how long ago this caching event occurred (Salwiczek, Watanabe, & Clayton, 2010). Our claim is that there is no need to make such an assumption in order to explain the birds’ behavior. It could be that the birds have some form of interval timing mechanism that governs how long the representation of edible worms remains a part of their model of the world (McCormack, 2001). Such a timer would begin to operate at the time of caching, and if the worms are then found to be rotten upon retrieval after a given interval (as in the learning phase of this study), the timer will subsequently ensure that a caching site is no longer represented as containing edible worms once that interval has elapsed (see Figure 1). In that way, we need not assume that the birds can remember the caching event itself; they simply do or do not continue to represent the hidden worms as a function of the state of their internal timer. Notably, this is quite different from assuming gradual forgetting of the

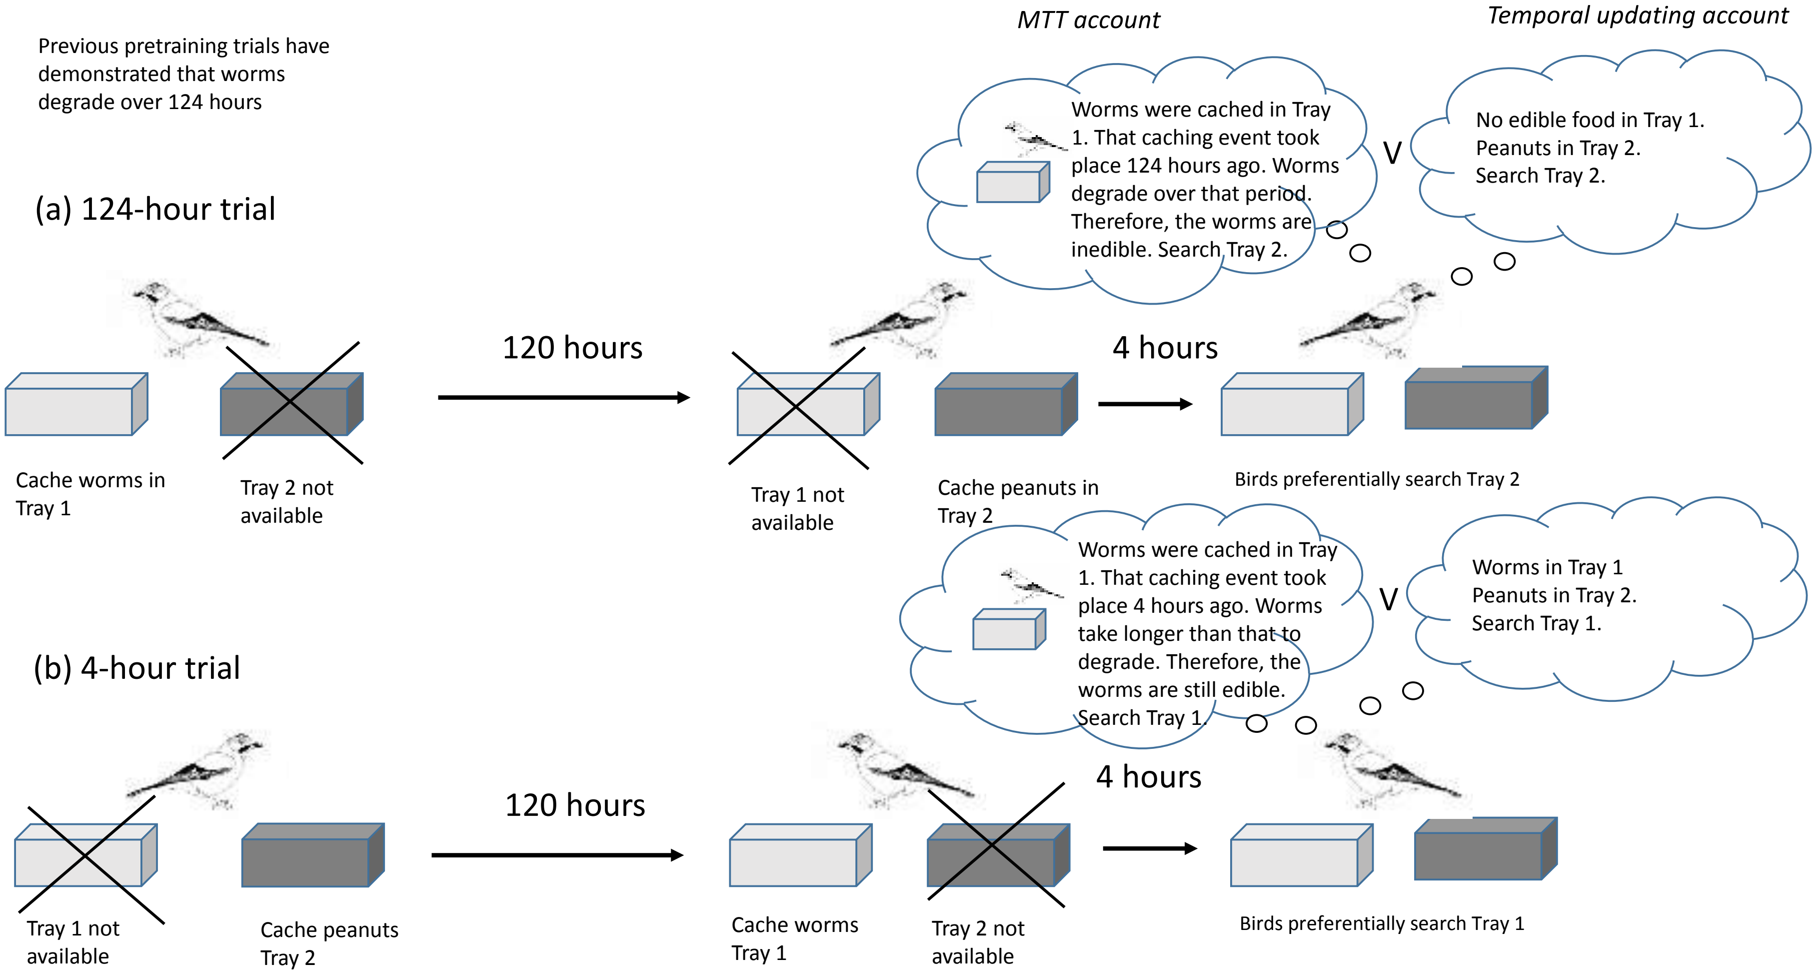

MTT account
Temporal updating account
Previous pretraining trials have demonstrated that worms degrade over 124 hours
(a) 124-hour trial
Worms were cached in Tray 1. That caching event took place 124 hours ago. Worms degrade over that period. Therefore, the worms are inedible. Search Tray 2.
V
No edible food in Tray 1. Peanuts in Tray 2. Search Tray 2.
120 hours
4 hours
Cache worms in Tray 1
Tray 2 not available
Tray 1 not available
Cache peanuts in Tray 2
Birds preferentially search Tray 2
(b) 4-hour trial
Worms were cached in Tray 1. That caching event took place 4 hours ago. Worms take longer than that to degrade. Therefore, the worms are still edible. Search Tray 1.
V
Worms in Tray 1 Peanuts in Tray 2. Search Tray 1.
120 hours
4 hours
Tray 1 not available
Cache peanuts Tray 2
Cache worms Tray 1
Tray 2 not available
Birds preferentially search Tray 1

Figure 1.

Illustration of test trials in Clayton and Dickinson's (1998) study. In (a) 124-hour trials, scrub jays cached worms in Tray 1 and peanuts 120 hours later in Tray 2. After 4 hours, they were allowed search in both trays. In (b) 4-hour trials, the birds initially cached peanuts in one tray and then cached worms in another tray after a 120 hour delay; 4 hours later they were allowed to search in both trays. Worms are the birds' preferred food, but birds in the Degrade condition in the study received a series of pretraining trials in which they learnt that worms had degraded after a 124 hour period. These birds preferentially searched for peanuts in 124-hour trials and worms in 4-hour trials. The figure contrasts the type of representation assumed to underpin the birds' preferences according to an MTT account with that which is assumed by the temporal updating account. On the former account, representations have tensed content that leads to birds to infer that the worms are not edible. On the latter account, it is assumed that whether the birds' model of the world continues to include a representation of edible worms in Tray 1 is governed by an interval timer.

locations of cached foodstuffs, although one could potentially describe it as a form of interval-timer controlled forgetting. Characterized in this way, what is distinctive about this type of forgetting is that it is appropriately flexible to the interval in question.

There are obviously further questions to be asked as to how exactly such an interval timing mechanism, which is triggered by the initial caching event, might function. Studies of interval timing in animals typically use considerably shorter intervals; indeed, Buhusi and Meck (2005) define interval timing as covering the range from under a second to 24 hours. However, there is no reason in principle to assume that animals do not have timing mechanisms that would allow them to be sensitive to the length of a 124 hour period (or even longer). Certainly, we are not committed to the idea that such a mechanism need be a dedicated internal clock; for example, a mechanism that keeps track of how many (fractions of) circadian cycles have passed would also be sufficient. Moreover, note that it is not just our account that needs to postulate the existence of such a mechanism that keeps track of time. Any account that holds that the birds actually remember the caching event and how long ago it occurred must also assume the existence of a mechanism that allows them to make an accurate judgment about the distance in the past at which that event occurred.

Clayton et al. (2003) explicitly consider, and do not reject, the possibility that the scrub jays' search behavior in their task could be governed by some interval timing mechanism, even in the case of somewhat more complex experimental designs. However, they do not seem to believe that this has a bearing on whether the task can be taken to measure MTT to the past. Our argument is that such a mechanism would allow animals to show the appropriate level of temporal sensitivity in the absence of any capacity to represent past events, by facilitating temporal updating in a way that is delay-sensitive. The crucial point is that in order for the timing mechanism to fulfil its purpose, it is enough for it to govern what happens to elements of the animal's model of its current environment; we need not postulate any ability to represent other times within this model (see also Hoerl, 2008).

2.2 "Mental Time Travel to the Future"

We now turn to considering whether a temporal updating account can also explain animal behavior in studies purporting to measure mental time travel to the future. Two types of tasks have been used in this area: tasks involving tool saving (inspired by Tulving's, 2005, original "Spoon test"; Dufour & Sterck, 2008; Mulcahy & Call, 2006; Osvath & Osvath, 2008), and tasks involving animals caching food that they do not currently desire eating (Correia, Dickinson, & Clayton, 2007; Cheke & Clayton, 2012). We will consider each in turn.

2.2.1 Tool saving tasks.

A study by Kabadayi and Osvath (2017) can illustrate the general structure of the tool saving tasks: In this study, ravens first learnt that a certain tool – a stone – could open an apparatus containing a food reward (see Figure 2). The following day, they were re-introduced to the baited apparatus, but now the tool to open it was not available, and the apparatus was removed again after a while. One hour later, and at a different location, the ravens were offered a forced-choice selection between the functional tool and three non-functional distractors. Fifteen minutes after that, they were re-introduced to the apparatus. Kabadayi and Osvath found that ravens selected the functional tool on the first test trial of this kind, and that they did so also on the majority of further trials on which they had to make their selection 15 minutes before being given access to the apparatus.

Does success in this tool saving task require MTT, and thus temporal reasoning, or can it be achieved using only the more basic temporal updating system? As we said before, an animal capable only of temporal updating maintains a model of the world only as it is at present. However, we allowed that this model could include items that the animal currently has no perceptual access to. Even if the ravens are capable only of temporal updating, the apparatus is still likely to figure in their model of the world when they are presented with the tool and the distractors, since they have learnt about its existence. Upon being re-introduced to the tool, the birds might thus realize that this gives them the opportunity potentially to open the apparatus, and this might be enough to motivate them to select the tool. That is to say, they select the tool because they want to open the apparatus, which they think of as part of their current environment.

Kabadayi and Osvath say that their study “suggests that ravens make decisions for futures outside their current sensory contexts” (p. 203). It is important, however, to distinguish between at least two ways in which this claim might be interpreted. As indicated, we agree that the birds’ behavior takes account of an *object* not currently within their sensory scope. As it is not within their current sensory scope, it is also true they will therefore have access to that object only after a delay (if at all). Nevertheless, it does not follow from this that, in acting on the basis of their representation of this object and its properties, the ravens must also be representing

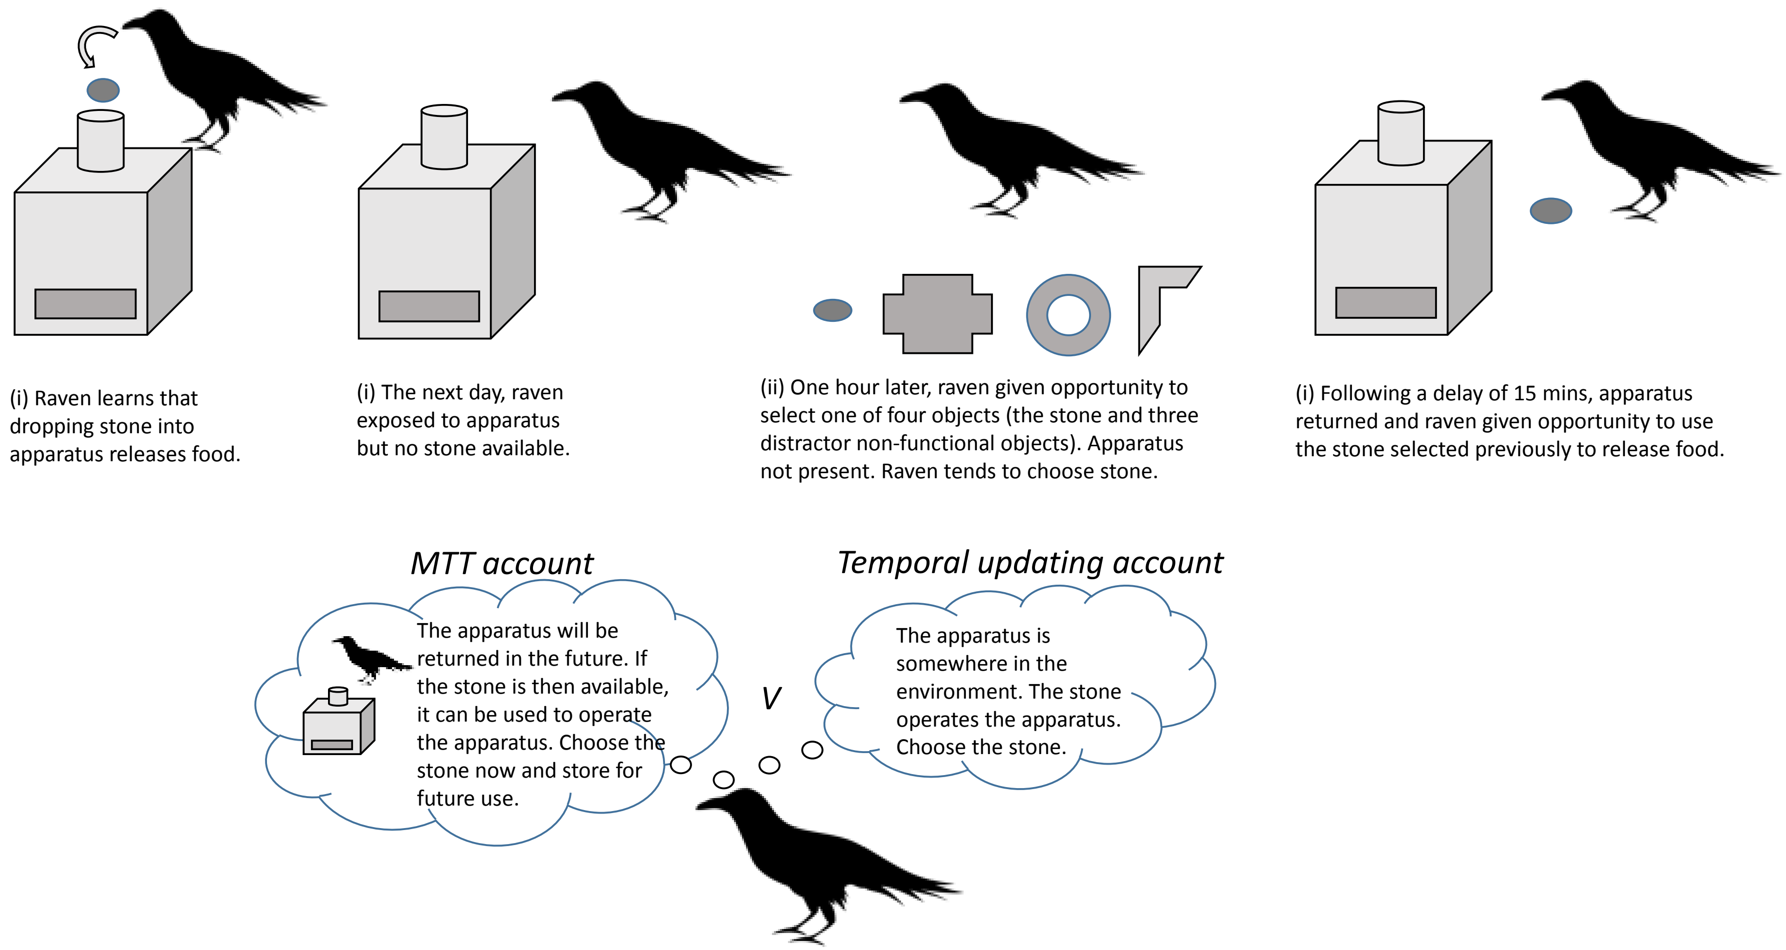

(i) Raven learns that dropping stone into apparatus releases food.
(i) The next day, raven exposed to apparatus but no stone available.
(ii) One hour later, raven given opportunity to select one of four objects (the stone and three distractor non-functional objects). Apparatus not present. Raven tends to choose stone.
(i) Following a delay of 15 mins, apparatus returned and raven given opportunity to use the stone selected previously to release food.
MTT account
The apparatus will be returned in the future. If the stone is then available, it can be used to operate the apparatus. Choose the stone now and store for future use.
V
Temporal updating account
The apparatus is somewhere in the environment. The stone operates the apparatus. Choose the stone.

Figure 2.

Illustration of the tool-saving procedure in Kabadyi and Osvath's (2017) study, in which ravens preferentially select a stone that can be used to operate an apparatus to release food (only part of the study carried out by Kabadyi and Osvath is described here). The figure contrasts the type of representation assumed to underpin the ravens' performance according to an MTT account with that which is assumed by the temporal updating account. On an MTT account, birds infer the need for the tool by representing a future event in which they will re-encounter the apparatus. On a temporal updating account, the birds' behavior is governed by a representation that includes information about objects outside of their current sensory scope (the apparatus) but does not include tensed content.

their future encounter with that object as such, as happening at a separate point in time distinct from the present. It is true that the experiments in question *de facto* involve a delay between the time when the tool is selected and the time when it can be used. But it is far from obvious that this delay also plays any role in the reasoning that leads the birds to choose the tool. (It is interesting to note, in this context, that the birds chose the functional tool already on the first test trial, before they had any opportunity to learn about getting re-introduced to the apparatus a set delay after tool choice.)

Part of what might motivate the idea that tool saving behavior demonstrates temporal reasoning capacities is that it seems to require some form of grasp of potential opportunities the environment might afford. But there is a more demanding and a less demanding way of understanding what the latter involves. In a different context Osvath and Osvath (2008, p. 662) describe planning in humans as involving "a capacity to construct mental experiences of potential events, something that could be expressed as a projection of the self into possible future events". To think of the ravens as engaging in planning in this sense would be to ascribe to them a capacity to represent modalities themselves – the ability to represent the future time of being confronted with the apparatus as a point in time distinct from the present at which two possible states of affairs could obtain – the animal having the functional tool or its not having it – depending on what the animal does now. What we are arguing, in effect, is that success in tool saving tasks does not necessarily require a grasp of possibility in this sense (for a study suggesting that non-human primates are in fact unable to represent such dual possibilities, see Redshaw & Suddendorf, 2016). It requires a grasp of possibility only in the more basic sense of requiring a grasp of the apparatus as an object that is potentially accessible, even though it is not within the animal's sensory scope. Such a grasp can be grounded in a representation simply of how the world actually is now (e.g., 'there is a baited apparatus that can be opened with this tool'), albeit one that might leave some aspects of states of affairs in that world underdetermined. More specifically, it can be grounded in a representation of the form we take to be involved in the temporal updating system – a representation on which the object is represented as existing somewhere in the birds' current spatial environment, and in this sense as potentially accessible. Together with the motivational state of desiring the contents of the apparatus, this seems sufficient to explain why the animals choose the tool.

2.2.2 Caching tasks

To illustrate the second type of task that has been taken as evidence of future mental time travel in animals, we will describe the study by Cheke and Clayton (2012). This study, depicted in Figure 3, relies on the fact that the birds show specific satiety effects, i.e., if they are sated on one food type, they prefer to eat a different food (note the specific set of foods used varied between birds; we illustrate one example set, and also only part of the overall study). This study started with a baseline caching trial in which birds were pre-fed with their maintenance diet that did not include peanuts or raisins, and then given the opportunity to cache each of the latter two food types freely in two trays. This was followed, after a delay, by a first retrieval stage in which they were pre-fed with raisins and subsequently given access to Tray 1 only. After a further delay there was a second retrieval stage in which they were pre-fed with peanuts and subsequently given access to Tray 2 only. In the test trials, which occurred some time after these retrieval stages, birds were pre-fed peanuts and then allowed to again cache both raisins and peanuts freely in the two trays. Of interest was whether the birds would cache the two foodstuffs in a manner that corresponded to which food they would prefer to retrieve from which tray, given their specific satiety at retrieval from the relevant tray. Would they cache peanuts in Tray 1, because they will be sated on raisins when Tray 1 becomes available, and would they cache raisins in Tray 2, because they will be sated on peanuts when Tray 2 becomes available? In fact, 3 out of 4 birds showed this pattern.

Cheke and Clayton's (2012) discussion frames this finding in terms of the idea that animals can "act for a future need that is different from their current one" (p. 171), or "overcome their current desire to anticipate […] future needs" (ibid.), implying that it demonstrates MTT capacities. As we will argue, though, this is a somewhat misleading characterization, and an explanation of the birds' behaviour purely in terms of temporal updating is readily available. According to this alternative explanation, birds will, under normal circumstances, represent both caching trays as equally good locations for caching food. This is what happens in the baseline caching trials, in which they cache foods in equal amounts in both trays. The birds are then given the opportunity, in the two retrieval stages, to retrieve food from Tray 1 only when pre-fed raisins and from Tray 2 only when pre-fed peanuts. In response to this, we suggest, they update their model of the world, such that now Tray 1 is represented as a good location for peanuts but not raisins, whereas Tray 2 is represented as a good location for raisins but not peanuts. This change in their model of the world explains why subsequently they differentially cache each foodstuff in a different location.

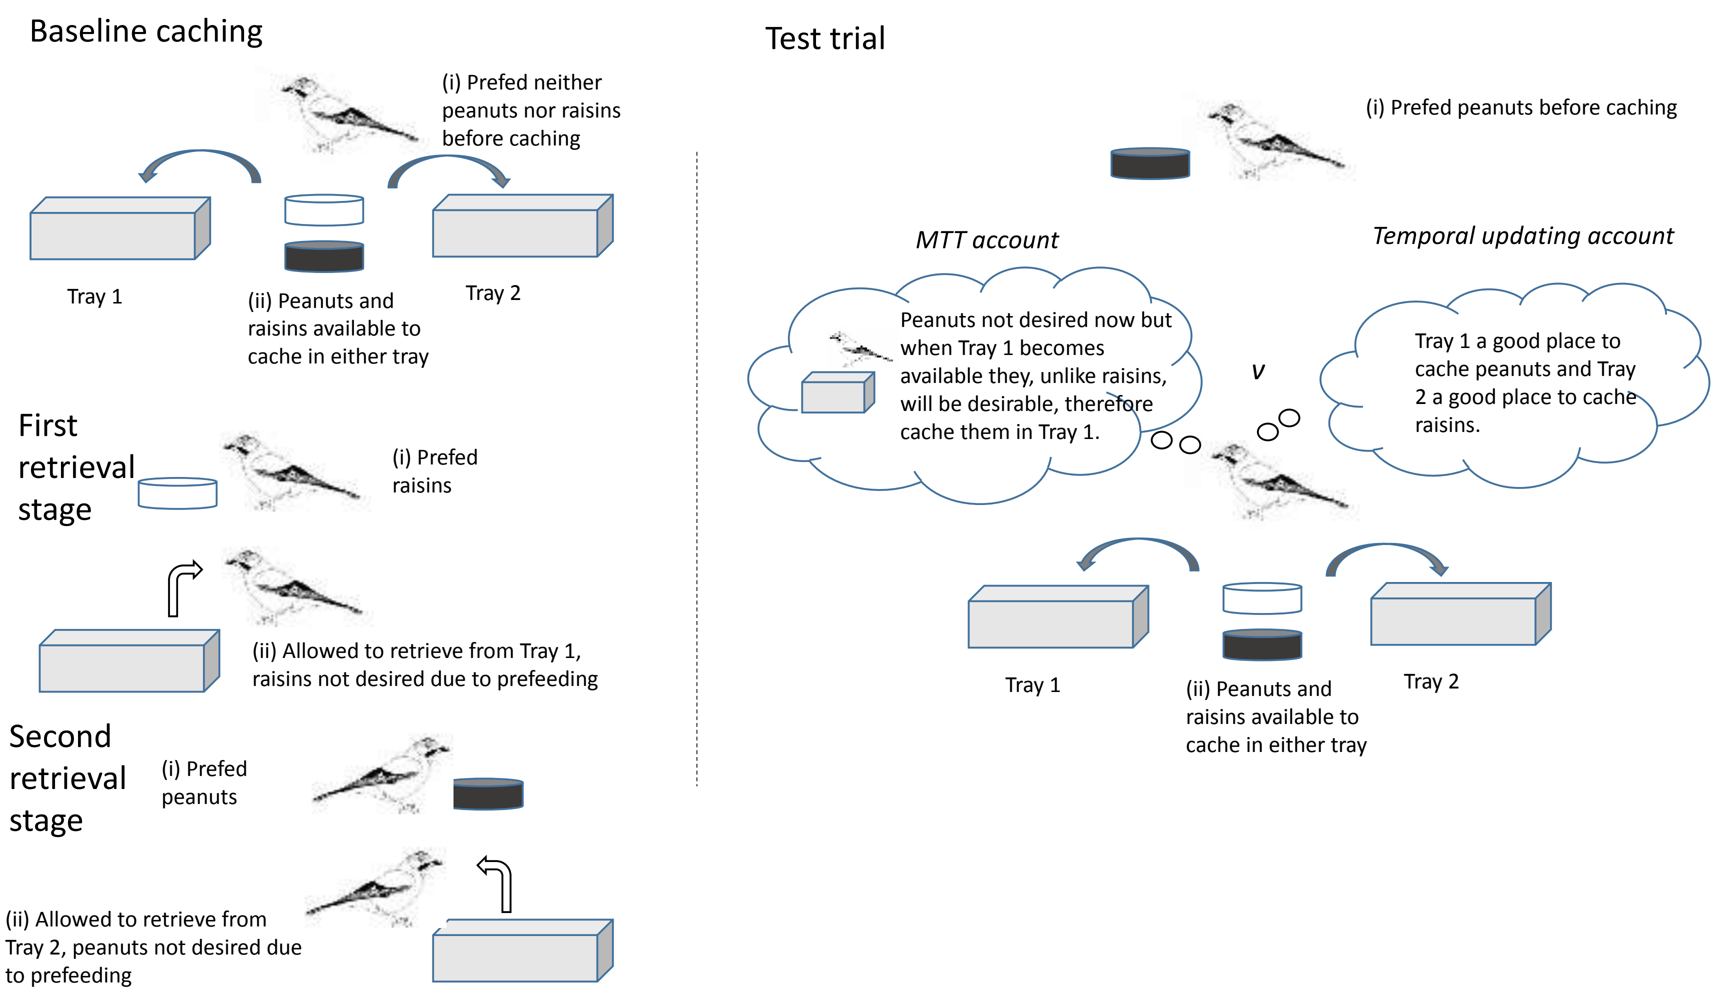
Baseline caching
(i) Prefed neither peanuts nor raisins before caching
Tray 1
(ii) Peanuts and raisins available to cache in either tray
Tray 2
First retrieval stage
(i) Prefed raisins
(ii) Allowed to retrieve from Tray 1, raisins not desired due to prefeeding
Second retrieval stage
(i) Prefed peanuts
(ii) Allowed to retrieve from Tray 2, peanuts not desired due to prefeeding
Test trial
(i) Prefed peanuts before caching
MTT account
Temporal updating account
Peanuts not desired now but when Tray 1 becomes available they, unlike raisins, will be desirable, therefore cache them in Tray 1.
v
Tray 1 a good place to cache peanuts and Tray 2 a good place to cache raisins.
Tray 1
(ii) Peanuts and raisins available to cache in either tray
Tray 2

Figure 3.

Illustration of the procedure in Cheke and Clayton's (2012) study for one pair of food types (peanuts and raisins). Boxes represent trays made available as caching locations; discs represent containers with food – black for peanuts, white for raisins. In the initial phase of the study (left-hand side of the figure), the birds were allowed to freely cache peanuts and raisins in two trays. In this initial phase, birds had not been pre-fed either food type. After a delay, they were pre-fed raisins and then allowed to retrieve from Tray 1. The pre-feeding of raisins meant that they no longer desired the raisins in the tray, preferring the peanuts. After a second delay, they were pre-fed peanuts and allowed to retrieve from Tray 2. The pre-feeding of peanuts meant that they no longer desired the peanuts in the tray, preferring the raisins. In the subsequent test trial (right-hand side of the figure), the birds were first pre-fed peanuts and were then again allowed to freely cache raisins and peanuts in both trays. The birds now preferred to cache peanuts in Tray 1 and raisins in Tray 2. On a mental time travel account, the birds infer that they should cache peanuts in Tray 1 on the basis of a representation of a future event in which they will find peanuts but not raisins desirable. On a temporal updating account, it is assumed that the birds update their model of the world, learning during the two retrieval stages that Tray 1 is a good place to cache peanuts and Tray 2 is a good place to cache raisins.

One might object to this explanation that it leaves out the significance of the fact that the birds are pre-fed peanuts at the start of the test-trial, and are thus, when subsequently allowed to cache, already sated on this food. Despite this satiation, they nevertheless cache peanuts in Tray 1. This seems to be what is behind Cheke and Clayton’s claims about the birds overcoming their current desires in favour of future ones, which suggest something like the following argument: The birds are currently sated on peanuts. Yet, they cache peanuts in the tray that will become available after they have been pre-fed raisins. Thus, they must be able to realize that, at that future time, they will desire to eat peanuts, and this is what motivates them to cache peanuts in that tray, even though they do not desire eating peanuts at present (see Figure 3). As characterized, this argument rests on the assumption that an animal restricted to thinking only about how things are at present is thereby also restricted to acting only based on its present appetitive desires. Note that in the case of food-caching birds in particular, there is no reason for making such an assumption. As Cheke and Clayton (2012) themselves note, there is separate evidence that the motivational systems for eating and caching operate semi-independently from one another. Furthermore, they also describe a separate experiment, consisting only of the equivalent of the test trial in the study described above, which shows that, whilst specific satiety reduces the motivation to cache the pre-fed food type, it does not eliminate it.

Thus, in so far as Cheke and Clayton’s study can be described as one in which the birds ‘overcome’ a current desire in favour of one that will serve their future needs, the issue at stake cannot be that the motivation to cache the pre-fed food clashes with a current lack of motivation to eat that food brought about by specific satiety. Rather, the only sense in which the birds can be said to ‘overcome’ an existing motivation is that, by default, if pre-fed one food type, they have a motivation to cache less of that food-type, and this changes into a motivation to selectively cache the pre-fed food-type in a particular tray once they can learn that that food-type is desirable at retrieval from the relevant tray (see also Cheke & Clayton, 2012, p. 174). This is entirely compatible with an account on which the relevant retrieval trials simply cause the bird to update its representation of which tray is a good location for caching which food.

2.3 Past and future thinking and animals: concluding remarks

Comparison across Figures 1-3 should make one aspect of our account clear. In each instance, the representation that we are assuming underpins the birds' behavior must necessarily *also* be part of what is represented according to an MTT account (i.e., that Tray 1 does not contain edible worms, that the stone should be chosen because it operates the apparatus, that Tray 1 is a good place to cache peanuts). Because of this, the representations we are positing cannot be considered to be implausibly *ad hoc*. Other questions one might raise with respect to our account, too, arise in the same way for accounts that postulate MTT abilities in animals. Thus, we have already noted that an MTT based interpretation of Clayton and Dickinson's (1998) study also has to assume the existence of some sort of mechanism that keeps track of time, which underpins animals' putative ability to remember "when" information. Similarly, in connection with our interpretation of Kabadayi and Osvath's (2017) study, questions might be asked about the conditions under which animals continue to represent an object as part of their current environment, even if it is outside their sensory range, and the conditions under which they stop doing so. Identifying these conditions is an empirical matter, and it is highly likely that the conditions vary by context and across species. But note that an explanation of the birds' behaviour that ascribes to them an ability to imagine future events raises exactly parallel questions as to the conditions under which the animal does or does not assume the future event will occur.

The key point that the plausibility of our account hinges on is whether we are correct about the general means by which the birds arrive at the representations that underpin their behaviour, i.e., purely as a result of the operation of the temporal updating system. And, as should be clear, there are real constraints on what is available to such a system: only a model concerning the world as it is at present is maintained, tensed content is absent from that model, and the correct functioning of the system depends on information about changes in the environment being received in the same order in which those changes happen. These constraints mean that, in explaining animal behavior, our account only allows for a distinctively narrow set of options. The alternative MTT account assumes that the birds arrive at these (very same) representations necessary to guide behavior as a result of first remembering a past or imagining a future event. Note, though, (and this is not always made clear) that MTT alone does not deliver these representations: e.g., in order to guide behavior, remembering the event of caching worms has to be combined with other information about how long it takes worms to degrade and how much time has actually elapsed, and then an inference has to be reached about the contents of Tray 1. Thus, the plausibility of that account hinges not only on whether animals can be thought to engage in MTT, but on whether they can be thought to use the information it delivers alongside any other information that is required to yield a conclusion as to what needs to be done right now.

Before leaving the issue of animal cognition, we want to briefly consider two recent theoretical accounts that have also attempted to provide an alternative explanation of animals' performance on these sorts of tasks. Redshaw (2014) has argued that animals may have "uncontextualized" representations of events, by which he means representations that fail to locate these events in any specific temporal context (in his view, such contextualizing would require metarepresentational abilities). And in articulating his own dual systems theory, Keven (2016) has argued that animals (and young children) do not possess episodic memory but may possess event memories that are "perceptually-based [and] snapshot-like". Unlike episodic memories, he believes these event memories are not organized into narratives with temporal-causal structure. Both accounts thus share the idea that animals have some type of free-floating representations of past (or future) events. Furthermore, although the two authors do not make this explicit in their discussions, their accounts are both compatible with one idea we are pressing, namely that there is no reason to believe that animals are capable of thinking about particular, unique, times. Nevertheless, the account we have put forward differs fundamentally from their accounts. The temporal updating system simply maintains a model of the world that records information about the environment and is updated in response to new information. Situations that obtained or will obtain at other times do not feature in such a model, even in an uncontextualized way. One way to put this is to say that on Redshaw and Keven's accounts, animals are not truly cognitively stuck in time: they can mentally meander through time even if they have no idea where in time their meandering takes them. On our account, by contrast, animals really are cognitively stuck in time: they cannot think about other times at all.

Yet, whilst our account is thus more radical than those offered by Redshaw and Keven, it can actually be seen to construe animals' behavior as more purposive than their accounts do. Note that, precisely because of their supposedly uncontextualized nature, it is ultimately not clear how exactly event memories of the type envisaged by Redshaw and Keven are supposed to explain the types of animal behaviors we have described – a point Redshaw himself seems to acknowledge in the context of discussing tool saving behaviors in animals, which he characterizes simply as cases in which an uncontextualized representation of using the tool in question induces a bias to select it again. That is, it is not clear how uncontextualized representations of how things were or will be at unspecified other times can systematically and appropriately guide present action (e.g., how can they generate the types of representations described at the beginning of this section, such as "Tray 1 does not contain edible worms"?). By contrast, even though on our account animals are not capable of representing situations obtaining at other times at all, the model of the current world they operate with can clearly give them good reasons to act in certain ways. For instance, we assumed that the ravens in Kabadayi and Osvath's study choose the functional tool because it allows them to open the apparatus containing the food reward, which they represent as an item existing in their environment.

3. When do children acquire temporal reasoning abilities?

We have argued that there are good reasons to doubt whether animals can think about the past or the future, and therefore that they have anything more than the temporal updating system. We believe that the same is true of infants, although unsurprisingly due to their limited motor skills the paradigms used to look at memory (e.g., Barr, Dowden & Hayne, 1996; Rovee-Collier, 1999) and future planning (e.g., McCarty, Clifton, & Collard, 1999) are typically quite different from those described above that have been used with animals (though see Atance, Louw, & Clayton, 2015; Martin Ordas, Atance, & Caza, 2017; Russell, Alexis, & Clayton, 2010; Russell, Cheke, Clayton, & Meltzoff, 2011). The challenge that our distinction provides for developmental psychologists is a more substantial one, though, than simply making the case that infants are capable only of temporal updating: it is to characterize the subsequent developmental emergence of the temporal reasoning system. In this section, we highlight some of the limitations in young children's temporal cognition and also some of the key developmental achievements that are required for mature temporal thought, focusing on two related areas: (i) the shift from relying on temporal updating to being able to reason about temporal order and (ii) the emergence of thought about other times. To anticipate, although we believe that infants operate only with the temporal updating system, and that temporal reasoning does not emerge until around 4-5 years, we think 2 to 3-year-olds may be at an intermediate developmental stage in which they are beginning to represent non-present situations and discriminate between them in a way that correlates with the difference between past and future situations. However, this is not genuine temporal reasoning, and children of this age may fall back on the temporal updating system.

3.1 Temporal updating versus reasoning about temporal order

Even infants can be sensitive to and learn about event order. Numerous studies of deferred imitation in infants have conclusively established that, at least by the second year of life, they can observe a short sequence of actions and reproduce those actions in the correct order even after a delay (e.g., Bauer & Mandler, 1989; Bauer, Wenner, Dropik, & Wewerka, 2000). Preschool children also rapidly acquire a repertoire of what have been described as ‘scripts’ for routines (Nelson, 1996).

These studies of infants and pre-schoolers suggest that children are very good at remembering and reproducing actions in the correct order. This basic ability in itself, however, is something that can fall within the scope of the temporal updating system rather than the temporal reasoning system. A key limitation of the temporal updating system that might be used to show whether children do indeed rely on it in learning about sequences is that the correct functioning of the system depends on it receiving information about events in the same order in which those events occur. We have described this as a signature limit of temporal updating, which contrasts sharply with a much more sophisticated way in which adult humans can deal with change over time, by engaging in genuine temporal reasoning about what happens when.

Although relatively few studies have examined temporal reasoning skills in young children, the evidence suggests that they struggle in situations in which the order in which they find out about successive changes may not reflect the order in which they happen – that is, in situations in which they cannot rely purely on temporal updating. In one study carried out by McCormack and Hoerl (2005), children learned that pressing (e.g.) a red button caused a marble to be released into a window of a box, and that pressing a blue button caused a toy car to be released. There was only ever one object in the window of the box at a time, so if the red button was pressed to yield a marble, on pressing the blue button the marble dropped away and was replaced by a car. Children also learned about two dolls who always acted in a particular order. The window of the box was then covered over, and two types of tasks were carried out. In one version of the task, children watched as one doll pressed the red button and then the other doll pressed the blue button. Under this condition, when asked which object was in the window of the box, even three-year-olds were able to answer correctly. This task can be solved by temporal updating: children can sequentially update their model of what is in the box window: initially representing it empty, then as containing a marble, and then as containing a car. However, in another version of the task the dolls pressed their buttons behind a screen, out of sight of children, and the dolls were then left beside the buttons they had pressed. Temporal updating could not be used to pass this version of the task, which required that children infer the window's contents by reasoning about the order in which the dolls had pressed their buttons, and even 4-year-olds struggled to do this (see also McCormack & Hoerl, 2007). Similar results were found in an earlier study by Povinelli, Landry, Theall, Clark, and Castille (1999), which used video clips to decouple the order in which children found out about two events from the order in which they actually happened.

Other studies have indicated that children of this age also have difficulties appropriately reasoning about temporal order relations in planning tasks, such as in circumstances in which they have to think ahead about the order in which events are going to unfold in the future (Lohse, Kalitschke, Ruthmann, & Rakoczy, 2015; McColgan & McCormack, 2008; relatedly see also Kaller, Rahm, Spreer, Mader, & Unterrainer, 2008; Tecwyn, Thorpe, & Chappell, 2014). Martin-Ordas (2017) asked children to select one of three objects to bring back to two rooms they had visited earlier; the correct answer was to choose the key needed to open a marble box containing marbles to use on a marble run. Three and four-year-olds correctly selected the key, but they were unable to judge which room they should then visit first – the room with the marble box or the marble run room. It was not until children were five that they could reason appropriately about the order in which these future visits needed to happen.

The claim that children below around 5 years find it difficult to reason about before-and-after relations in time might sound surprising given that children actually acquire the verbal terms “before” and “after” at an earlier age (Busby-Grant & Suddendorf, 2011). There is, though, also evidence regarding children’s competence in using and interpreting those terms that is in line with what our account would predict. Specifically, they have difficulty correctly interpreting these terms when the order in which events are mentioned in a sentence does not match their order in the world (e.g., “Anna took off her coat after she took off her hat.”; Blything & Cain, 2016; Blything, Davies, & Cain, 2015), suggesting they use an order-of-mention strategy to interpret them. While there are a variety of interpretations of this finding (Blything et al., 2015; Pyykkönen & Järvikivi, 2012), one possibility is that this reflects pre-schoolers’ difficulties with temporal reasoning and their tendency to use the temporal updating system.

Taken together, the findings we have discussed here suggest that while even infants can learn about and be sensitive to event order, difficulties in reasoning about temporal order persist into the preschool years, with the findings from some studies suggesting that the errors children make may be due to falling back on the temporal updating system. By the time children are five, existing evidence suggests that they have consolidated some important new temporal reasoning skills.

3.2 Thinking about other times

As with animals, there is considerable debate over how infants' and preschoolers' memory abilities should be characterized (e.g., Bauer, 2007; Fivush, 2011; Howe & Courage, 1997). There have been attempts with various degrees of success to use a supposed "what-when-where" paradigm analogous to that used with animals (Burns Russell, & Russell, 2015; Martin-Ordas et al., 2017; Russell et al., 2011). Notably, even relatively old children struggle with tasks analogous to that of Clayton and Dickinson (1998) that require sensitivity to how long ago an item was hidden, which on our account of animal performance on such tasks is unsurprising because children would have no need in everyday life to have their search behavior governed by sensitive interval timing mechanisms. Performance on the tasks more typically used to measure infant memory (Bauer et al., 2000; Rovee-Collier, 1999) can be straightforwardly accounted for in terms of the temporal updating system. However, it is less straightforward to explain pre-schoolers' verbal descriptions of non-current events merely in terms of the idea of temporal updating. Existing studies suggest that children who are 2-3 years can talk about both past and future events, albeit often providing limited and fragmented information (Hayne, Gross, McNamee, Fitzgibbon, & Tustin, 2011; Peterson, 2002; Weist & Zevenbergen, 2008). It is widely accepted that children of this age can refer to things that are outside their current sensory scope (Sachs, 1983), which is something that the temporal updating system could handle if these things are still part of their model of the world as it is now. However, children do not just refer to things outside their current sensory scope, they often use tense to describe non-current events, and we have assumed that the temporal updating system does not operate with tensed representations. Given that 2-3-year-olds can refer to events in the distant past and in the future, and 3-year-olds also use temporal adverbs

(Weist & Buczowska 1987; Weist & Zevenbergen, 2008), it might seem paradoxical to argue that they cannot think about the past or future.

We accept that children of this age do not rely only on the basic temporal updating system. However, we want to argue that they are at an intermediate stage at which they nevertheless do not yet possess genuine temporal reasoning abilities, and, in that sense, do not have proper concepts of the past or future. To make this argument it is necessary to consider more carefully the nature of the domain over which the temporal reasoning system operates. Temporal reasoning operates over the domain of times, with times arranged in a linear array such that each time occupies a unique unrepeated location in the array. Reasoning about such a domain involves a grasp of two distinctive types of systematicity that obtain within it. First, systematic before-and-after relations obtain between points in that array due to its linearity – e.g., for a sequence of times A, B, and C, if A happens before B and B happens before C, A must also happen before C. Second, which times are in the past, present, and future changes systematically with the progression of time: for the sequence of times just mentioned, if A is now present, then both B and C are in the future. But when B will be present, A will be past and C future, and when C will be present, both A and B will be past.

There is no reason to believe that young pre-schoolers can reason about the domain of time in this way. While they may talk about events that are in the past or the future, there is no reason to believe that they have a sense of where in the past or future those events are located or a grasp of the systematic temporal relations that obtain between these events. In Tillman, Marghetis, Barner, and Srinivasan's (2017) recent study, 3-

year-olds were unable to make judgments about the relative order of a set of past and future events (their previous and next birthdays, breakfast this morning, and dinner this evening). Indeed, 3-year-olds in this study were unable to reliably judge the deictic status of these events (nor the deictic meaning of time words such as "yesterday"). Similarly, Busby-Grant and Suddendorf (2009) found that children of this age could not discriminate the relative distances in the future of even very widely-separated events (for instance, going home from day-care versus next Christmas).

Important improvements in children's ability to think about the temporal locations of events occur between 3 and 5 years. There is evidence that by the time children are four they can begin to make some discriminations about the relative recency of unrelated events in the past (Friedman, 1991; Friedman, Gardner, & Zubin, 1995; Friedman & Kemp, 1998; McCormack & Hanley, 2011; though see Pathman, Larkina, Burch, & Bauer, 2013). However, even four-year-olds struggle to order the times of events in the future (Friedman, 2000; McCormack & Hanley, 2011), or to judge the remoteness of past and future events (Busby-Grant & Suddendorf, 2009; Tillman et al., 2017). Moreover, a number of studies have shown that they tend to confuse near past and near future times (Friedman, 2003; Friedman & Kemp, 1998). Indeed, the ability to order events in time continues to improve substantially over the next few years (Hudson & Mayhew, 2011; Pathman et al., 2013; Tillman et al., 2017).

We have suggested that reasoning over the domain of times involves not just understanding the relations that obtain between points in time but also understanding that which points in time are in the past, present, or future changes systematically with the progression of time. Extremely few studies have addressed children's ability to engage in the kind of temporal perspective-taking that grasping this requires, but those that have indicate that children below 4 do not have this ability (Cromer, 1971; Harner, 1980, 1982). In Cromer's task, children had to consider the deictic status of an event from a point in time that was not the present. For example, children were told a story about a girl who visited a farm where a number of things happened, including her seeing birds, and who then returned from the farm. Children had to judge (e.g.) at what stage in the story the girl could refer to seeing the birds in the past tense. It was not until children were 4-5 years that they were able to answer these types of questions correctly.

Put together with the findings discussed in the last sub-section, the evidence weighs heavily in favor of the idea of an important transition in the period from 3-5 years in children's ability to engage in temporal reasoning. During this period, children acquire mature concepts of time and start to be able to reason about the domain of linearly-ordered times. This leaves the question, though, of how we want to characterize temporal cognition in the early preschool years, if we want to argue that children of this age cannot properly think about the past and future. We suggest that at this age children are able to make some sort of discrimination between situations that have obtained and situations that are yet to come. One way to put this is to say that they may retain models of the world that have been superseded (i.e., of past states of affairs), or models of the world as it has yet to be (i.e., of future states of affairs). As demonstrated by generally accurate use of past and future tense, they can usually appropriately discriminate between these models, according to which of these two types they belong to. Nevertheless, we do not believe that children of this age are treating some models as descriptions of situations located at specific past times and others as descriptions of situations located at specific future times. Rather, we believe that children of this age may simply make a categorical distinction that marks a difference between these two types of situations; specifically, we believe that children discriminate between situations that are no longer alterable and situations that are still potentially alterable (see McCormack, 2015; McCormack & Hoerl, 2017, for considerably more detail on this proposal and our developmental model). However, this is different from having one unified model of the world within which time itself is represented. Having such a unified model goes beyond just representing certain sets of states of affairs and being able to discriminate between them. It involves representing time as one of the dimensions along which reality is extended, and as a linear dimension along which these sets of states of affairs can therefore be organized, so that they can all be captured in one set of systematic temporal relations in which they stand to each other. Temporal reasoning, in other words, operates with a four-dimensional picture of reality, on which everything that happens can be described by giving its location and the time at which it happens. It is the discovery of time as this fourth dimension that is the crucial step in the transition to a temporal reasoning system.

Space precludes us from commenting more than very briefly here on further developmental questions regarding cognitive prerequisites for the development of temporal reasoning – in particular regarding claims that have been made to the effect that temporal reasoning requires language (Bennett, 1964), or that it requires a capacity for metarepresentation (Redshaw, 2014). With regard to the role of language, we believe it is plausible (although we will not develop the argument here) that acquisition of the basic temporal concepts discussed here requires language, perhaps because it is only through discussing non-current events with others that children begin to grasp how such events are temporally organized (Hoerl & McCormack, 2005; Welch-Ross, 2001). This does not mean, though, that children learning different languages or growing up in different cultures acquire different concepts of time. We assume that our description of the temporal updating system captures basic and universal features of human thinking about time, and this includes a notion of time as linear. Although cultures may differ in the extent to which they emphasize cyclical or repeating patterns in time, we follow Gell (1992) in assuming that linearity is a universally basic feature (McCormack, 2015, McCormack & Hoerl, 2017). The features of the temporal reasoning system that we have highlighted will of course be overlaid by further culturally specific constructs, such as different ways of metaphorically mapping time onto space or of measuring time using a calendar system. These further culturally-specific aspects of development have a protracted developmental time course (Friedman, 2003; McCormack, 2015).

With regard to the issue of metarepresentation, we note an interesting structural parallel between the account we have proposed of the development of temporal reasoning capacities and Perner's (1991) influential account of the development of metarepresentation. On Perner's account, children move from having only one model of the world through an intermediate phase of being able to switch between different such models, before finally being able to conceive of them as different representations of the same reality. On our account of the emergence of temporal reasoning, children are initially capable only of temporal updating, and thus operate with a tenseless model of the world. They then begin to be able to maintain models describing non-current states of the world but they do not represent these as states organised along a single temporal dimension, and thus do not grasp the systematic temporal relations between them. In acquiring genuine temporal reasoning, a unified model emerges that allows children to represent how these states are temporally organized and interrelated.

Although the structural parallel is interesting, we note, however, that there is an important difference between Perner's description of the emergence of metarepresentation and our description of the emergence of temporal reasoning, in that the accounts differ in terms of the type of systematic relations between models that children need to learn to grasp. His claim is that children need to grasp how different models of the world are related to the actual world – specifically, that they are representations of reality – and it is through this that children understand how the models relate to each other, as different possible ways of representing the very same world. That is, grasping the type of systematic relations that Perner is interested in is a consequence of understanding the representational nature of mental states. In our account, children must grasp the systematic temporal relations between different models of non-current states of the world by realizing they are located at points on the same time line. It cannot be straightforwardly assumed that grasping this type of relation necessarily requires grasping the representational nature of mental states, although we recognize that there is a more detailed developmental story to be told about how this new ability comes about, further discussion of which we provide elsewhere (Hoerl & McCormack, 2005; McCormack, 2015; McCormack & Hoerl, 1999, 2001, 2017).

3.3 Animal cognition revisited: lessons from developmental studies?

Before turning to adult human cognition, we want to briefly consider whether the findings of developmental studies might help illuminate what types of animal research would be capable of providing a test of the hypothesis that animals rely only on the temporal updating system. Developmental studies have looked in detail at children's ability to represent and reason about event order information, and two types of findings are of note. First, developmentalists have devised studies that provide children with information about changes that have happened in a sequence, but do not provide this information in the order in which the changes themselves unfolded (both by showing mis-ordered videoclips that children need to mentally re-order, Povinelli et al., 1999, and by requiring children to infer the sequence from two pieces of information that are presented simultaneously, McCormack & Hoerl, 2005, 2007). The temporal updating system as we have described it can only provide an accurate model of how the world is now if changes are encountered in the sequence in which they actually occur, so reliance on that system would result in task failure. Second, we also note that some developmental planning tasks require that children do more than simply select a tool that is functional for obtaining a reward that exists somewhere in their current environment. In McColgan and McCormack's (2008) study, children must bear in mind the order in which events are going to unfold in the future and appropriately place an object so that it will be encountered later at the right point in a sequence of future events; in the study of Martin-Ordas (2017), children must think ahead about the order in which they need to visit two rooms so that they are appropriately prepared when they later encounter the reward apparatus. We anticipate that it might be possible to devise animal versions of both these types of studies, and such studies would prove particularly illuminating in testing our hypothesis about limitations in animal cognition.

4. Dual systems in adult human temporal cognition

Once the capacity for temporal reasoning has developed, does the temporal reasoning system simply replace the more primitive temporal updating system, or does the temporal updating system remain in operation even in adults? In this section, we consider some evidence – albeit from a somewhat unusual source – suggesting that even adults still operate with the temporal updating system alongside engaging in temporal reasoning.

A claim familiar from some of the existing literature on dual systems perspectives on cognition is that one of the hallmarks of two systems being at work alongside each other in people's thinking about a particular domain is that they can give rise to cases of what Sloman (1996, p. 12) calls "simultaneous contradictory beliefs" about aspects of that domain, where 'belief' is to be understood as "a propensity, a feeling or conviction that a response is appropriate even if it is not strong enough to be acted on" (ibid.). The idea is that in a case in which the two systems yield diverging outputs, the more primitive, automatic, system still delivers its verdict, even if it is not endorsed by the more deliberate reasoning system, giving rise to a felt pull towards making one judgment, despite this judgment being rejected as incorrect.

We believe that a careful consideration of aspects of what might be called adults' 'naïve theory of time' provides evidence for the existence of such simultaneous contradictory beliefs also in the domain of temporal cognition. Unlike naïve theories of some other domains (Gelman, 2006), the nature of this naïve theory of time still awaits systematic empirical attention, although there is growing interest in this topic amongst philosophers (see, e.g., Braddon-Mitchell & Miller, 2017; Callender, 2017). Their interest stems from their belief that there are particular elements of people's naïve theory of time that cannot simply be explained in terms of the physics of time. i.e. as reflecting features of time as it figures in physical theory. This raises the question as to how exactly these elements of people's naïve theory of time should be characterized and what their actual psychological sources are. Using the term 'manifest time' to refer to humans' naïve theory of time, Callender (2017, pp. 23f.) puts the point as follows: "[Once] one removes the project of explaining manifest time [in terms of physical reality, one] places it on the desks of psychologists. The psychologists, however, don't know it's on their desk. The end result is that manifest time remains unexplained."

One core ingredient of humans' naïve theory of time that has been argued to require such a psychological explanation is the belief that there is an objective flow or passage of time for which there is no spatial equivalent. This belief appears to be universal (Gell, 1992), despite cultural variations in a number of other aspects in which time is conceptualized (Boroditsky, 2011). It has also been documented in the context of research in psycholinguistics on limitations on the extent to which aspects of time can be captured by spatial metaphors. There it has been argued that, whilst spatial metaphors for time are ubiquitous, there is also a set of metaphors for time that take change or motion as their source, pointing to a unique attribute of 'transience' involved in the concept of time, which cannot be captured other than in terms of notions that themselves invoke time (Galton, 2011).

What exactly is involved in thinking of time as flowing or passing? Philosophers have offered an analysis of this ingredient of people’s naïve theory of time on which it involves the combination of two components: the belief that there is just one objectively present moment in time, and the belief that which moment in time is objectively present changes over time (Leininger, 2015). The claim that people’s naïve theory of time involves the belief that there is one *objectively* present moment in time is best illustrated by contrasting what people take to determine “now” with what they take to determine “here”. What counts as “here” is clearly just a matter of where the speaker using that word is located, and thus a place is “here” only for the speaker who is using the word. A particular place’s being “here” is not a property of space itself; it is purely a matter of perspective. By contrast, people’s naïve theory of time does not conceive of a particular time being “now” as similarly being purely a matter of perspective: rather, which moment in time is present is taken to be a property of time itself (and hence the same for everyone), and moreover a property of time that itself changes over time. It is in this sense that the naïve theory of time involves the belief that there is one objectively present moment in time. Moreover, there is a crucial connection between the belief that there is one objectively present moment of time – one objective now – and the belief that there is such a thing as the passage or flow of time. The idea of the flow or passage of time requires that there is an objective now because it assumes that there is something about *time itself* that changes over time – i.e., there is a property that time itself has that is different from one moment to another, namely which moment of time is present. Time is assumed to be fundamentally unlike space in this respect, because space does not possess such a property – where “here” is does not change because space itself changes, but only because the location of the speaker using the word changes. By contrast, the idea that time flows or passes assumes that which moment in time is "now" is an objective matter, a property possessed by time itself. By the same token, though, it also requires that which moment in time is thus objectively present is something that does change over time. As time goes on, the objective now occupies successively later points in time. It is this change that time itself is assumed to undergo over time that the passage or flow of time consists in.

As we mentioned, much recent debate in the philosophy of time has focused on whether components of the naïve theory of time such as the idea that time passes or flows reflect how time really is, particularly in the light of contemporary physics. We are not directly concerned with this debate here (see Rovelli, 2018, for an accessible discussion). Rather, we want to consider the idea of the passage or flow of time in the light of our dual systems theory of temporal cognition. In what follows, we wish to defend two claims: (i) There is an inherent contradiction in people's naïve theory of time, in so far as it contains within it both the belief that there is an objective present and also the belief that which moment in time is objectively present changes. (ii) This contradiction in people's naïve theory of time can be explained in terms of the co-existence of the dual systems we have identified.

Arguments to the effect that the naïve view of time is inherently contradictory in virtue of containing within it the idea of the passage or flow of time are in fact long-standing within philosophy (McTaggart, 1908; Bardon, 2013). Price (2011) has argued that the basic problem with any picture of time involving that idea is that it wants to be *exclusive* and *inclusive* at the same time. The picture is exclusive insofar as one moment in time is supposed to enjoy some form of objective privilege in virtue of being the one moment in time that is present. Yet it is also supposed to be part of the picture that which moment in time is present changes, meaning that more than one moment in time gets to be present. This implies inclusivity rather than exclusivity: because each moment in time gets to be present when its time comes, no one moment can be objectively privileged. That is, each moment in time is just on a par with all others in being the present moment in time when it is that moment in time. One might try to respond to this by holding on to the claim that the present moment is special, but also claim that which moment is present depends on what time it is. The difficulty with responding in this way is that it makes which moment in time is present dependent on what time it is considered from, rather than it being an objective property of time which moment is present. This difficulty becomes clear if one returns to the case of space. One could say in a similar way that the place indicated by ‘here’ is special in some sense, e.g., because it is the place that the person referring to ‘here’ is located in. Clearly, different places then get to be special in that sense as the speaker moves around (and other places get to be special for other people). But this means that their specialness does not stem from something objective about the spatial locations themselves, their specialness is simply a matter of perspective. Thus, accepting that which moment in time is present is similarly just a matter of perspective requires giving up on the idea ‘objective now’ and with it the idea of the passage or flow of time.

In as far as the naïve theory of time, and more specifically the idea that time passes or flows, does indeed involve a set of simultaneous contradictory beliefs along these lines, how might the dual systems account of temporal cognition we have outlined account for how they arise? The picture of time that Price describes as ‘inclusive’ is one on which all times are indeed on a par with one another and events and states of affairs are thought of in a way that is temporally qualified: they are thought of as taking place or obtaining at certain times, with other events or states of affairs taking place or obtaining at other times in exactly the same manner. This is broadly speaking the kind of way of thinking about time people employ when using the temporal reasoning system, in which time is thought of as a linear dimension along which reality is extended. Some of the key manifestations of this type of thinking are the ability to use a clock-and-calendar system and the ability to engage in MTT. Using the temporal reasoning system can enable people to recognize that the present does not really have an objectively special status, and that thinking of certain events as present is actually simply a matter of locating them with respect to one out of many possible perspectives on time, just as thinking of a place as ‘here’ is a matter of locating it with respect to one out of many possible perspectives on space.

Insofar as the idea of there being one objectively privileged present moment in time nevertheless also figures in everyday thinking about time, we want to suggest, its source is a residual tendency people still have to think of the world using the other, less sophisticated temporal updating system (see also Falk, 2003, and Prosser, 2006, for related ideas). How exactly might the existence of the temporal updating system explain the sense in which it seems to people that the present somehow has a privileged status? When the latter idea is discussed in the literature on the metaphysics of time, it is sometimes further specified as the idea that people have an impression that present things exist *simpliciter*, without temporal qualification (Skow, 2011; Zimmerman 2005). That is to say, it is not that people’s impression is that past and future things do exist, just not right now in the present. Rather, the impression is that past and future things simply do not exist at all – only present things do. As we have seen, representing present things in a way that is not temporally qualified is also a feature of the model of the world maintained by temporal updating system. It is a model of the world that concerns the world only ever as it is at present. When that model is updated in response to a change in input it is simply replaced by a new model, with nothing in the new model representing that things were previously different from how they are now. But this implies that each model fails to identify the then-current situation as only one amongst many, temporally speaking. Nothing within the system signals that its representations represent just how things are at one time, with there also being other times at which things are different. In this way, the operation of the temporal updating system might explain a bias people have towards thinking that there is only one objectively present moment in time and that only what is present exists or is real. People seem to have this bias even though they can also recognize, using the temporal reasoning system, that what they think of as the present is only one perspective onto time amongst many others. In this way, the operation of the dual systems gives rise to contradictory elements in people's naïve theory of time.

Note that the issue here is not just one of diagnosing a contradiction implicit in people's everyday thinking about time and providing an explanation of it. Unlike most other contradictions, in this case the contradiction does not simply get eliminated when people notice it. Philosophers who have come to the firm conviction that time does not really pass or flow, or even that the idea of time as passing or flowing is incoherent, frequently admit that they nevertheless still have this impression of time (for one example, see Ismael, 2017, p. 35). Similarly, even

Einstein, who saw clearly that his physics ruled out any ‘objective now’, continued to be troubled by what he called the “problem of the Now” (Carnap 1963, p. 37). It is the phenomenon of ‘simultaneous contradictory beliefs’ in this more specific sense that dual system accounts have been claimed to be able to be particularly well suited to explain (Sloman, 1996).

5. The two systems and intertemporal choice

We now turn to one final area of research that we think our dual systems account of temporal cognition can bear on in interesting ways, which is research on intertemporal choice (i.e., choices that involve assessing the relative value of rewards available at different time points). The idea that some form of dual systems view might usefully be brought to bear in explaining some of the phenomena surrounding intertemporal choice can already be found in the existing literature. Perhaps the most influential two-systems approach to intertemporal choice is Metcalfe and Mischel’s (1999) hot/cool systems analysis of delayed gratification. However, their distinction is one between an emotional system and a cognitive system, whereas we have drawn a distinction between two different cognitive systems: one for temporal updating and one for temporal reasoning. Our aim is therefore not to replace their distinction or similar existing dual process approaches to delayed gratification. Rather, the idea of two different cognitive systems that deal with how things unfold over time might provide us with a more fleshed-out picture of how exactly the processes already appealed to in the existing literature might be involved in delayed gratification. In this section, we want to argue that the two different systems provide for two quite different types of mechanisms that might facilitate delay of gratification. Note that our aim is simply to describe these two possible mechanisms for delaying gratification, and how they relate to the distinction between the temporal updating system and the temporal reasoning system; we do not attempt to explain the broad set of phenomena that have been extensively studied under the heading of intertemporal choice (e.g., the shape of the temporal discounting curve).

The first, more basic, form of delayed gratification might be illustrated by considering a proposal Boyer (2008) has put forward concerning what he sees as the function of future episodic thinking. Boyer thinks that simulation of future outcomes “can act as a calibration device by triggering emotional rewards that accurately reflect the emotional impact of [these outcomes] and are immediate and, therefore, bypass the usual discounting of future consequences of actions” (Boyer, 2008, p. 221). That is, mentally ‘pre-experiencing’ the reward allows people to give it due weight in their decision making and avoid discounting. Boyer advertises this as a way in which future episodic thinking can aid delaying gratification. Interestingly, though, its net effect could actually also be described as that of diminishing the relevance of time by allowing the chooser to ‘pre-experience’ the reward. Whilst it is true that the mechanism is triggered through a representation of the future reward, and in this sense involves temporal reasoning, we want to suggest that the subsequent effect of simulating that reward is actually to bring it within the purview of the model of the world maintained by the temporal updating system. The reward can then figure within that model as a more valued alternative to the immediately obtainable award, motivating the chooser to take steps to obtain it instead. To flesh out this idea further, we can draw on some features of the temporal updating system that we have already noted.

As we explained, the model of the environment maintained by the temporal updating system is best understood as a model in which time simply does not figure. Thus, there is no way for a future reward to figure in that model as such, i.e. as one that exists, but at a time different from the present. However, things can figure in the model of the world maintained by the temporal updating system even if they are not immediately accessible. For instance, the model may represent a reward as existing in the environment, even though it may not be obtainable immediately, e.g., because it is not within current sensory scope. We hypothesize that simulating a future reward, in the way envisaged by Boyer, allows the chooser to represent the (de facto) delayed reward in its model of the world in this way: both the delayed and the immediate reward are represented as potential objects for present choice, thus overcoming the discounting of the future that normally comes with using the temporal updating system. Note, though, that we are not arguing that this way of overcoming discounting requires only the temporal updating system. The temporal reasoning system is required because it triggers the project of simulating the future reward; the updating system itself does not represent future states of affairs. However, once the future reward is simulated, it can then feature in the model of the world used by the temporal updating system as an available choice.

We want to distinguish this first way of mitigating temporal discounting, which primarily exploits cognitive resources available within the temporal updating system, from a second, more sophisticated one, in which temporal reasoning plays a much more central role. Put briefly, the key distinction might be put as follows: The kind of mechanism just described involves the chooser representing two choice objects – an immediately available smaller reward and a not immediately available larger reward –

and putting themselves into a position to be able to weigh them up against each other by bringing both of them within the purview of the model of the world maintained by the temporal updating system. In delayed gratification involving temporal reasoning, by contrast, it is not just such individual choice objects that are weighed up against each other; rather, what becomes a matter of choice is how the chooser wants events in their future to unfold over time. Specifically, they need to decide how they want rewards to be distributed over the time line that they represent as stretching out into the future, and in this sense they are reasoning about time itself. As we might also put it, whereas the mechanism we have previously described allows the agent to give the future reward due weight through setting aside its futurity, temporal reasoning can mitigate discounting by allowing the agent to give the future itself due weight in their deliberations.

In the preceding section, we described the picture of time that the temporal reasoning system operates with as an inclusive picture of time, following Price, on which different times are seen as being on a par with one another, rather than one time, the present, having a special status. This feature of temporal reasoning, we want to suggest, is what allows agents to switch from simply trading individual rewards off against each other to considering how they want rewards to be distributed over the time line stretching out into the future. There are a number of different factors that have already been shown to impact on intertemporal choice which might be seen to affect it through facilitating reasoning about time in this way.

For instance, using a number of different real-life measures such as saving for retirement, adopting a healthy lifestyle and practising safe sex, Chen (2013) has shown that speakers of languages that allow for future events to be spoken of in the same grammatical forms as present events show more future-oriented behavior than speakers of languages that require a grammatical future marker. It is unlikely that this is due to the former simply ignoring the difference between present and future events, or imagining future events as present and using a mechanism of the type described by Boyer to delay gratification. At least some of the prudential behaviors studied by Chen involve fairly complex long-term goals imagining the attainment of which is unlikely to have a strong immediate hedonic effect. Rather, we believe that what lies behind these findings is that languages with an obligatory future tense marker encourage a focus on the present, because they encourage thinking of future events as somehow having a different status than present ones. Languages which allow for thinking of future events in the same grammatical categories as present ones, by contrast, facilitate thinking of future events as having the same status as present ones, despite being temporally distant, and therefore as ones that ought to bear a similar weight in one's deliberations. Other existing studies might be seen in a similar light. For instance, a number of studies that have followed up Boyer's (2008) suggestion that episodic future thinking plays a special role in supporting delay of gratification (e.g., Benoit, Gilbert, & Burgess, 2011; Lin & Epstein, 2014; Peters & Büchel, 2010) have demonstrated that simply getting people to imagine the future per se (i.e., not just imagining rewarding events) reduces temporal discounting. Even simple manipulations that vary the way relevant future points in time are characterized, such as presenting them as dates rather than delays, impact on degree of discounting (Read, Frederick, Orsel, & Rahman, 2005).

Crucially, the temporal reasoning system allows individuals to think of their lives as *temporally extended projects*. This requires more than just an ability to trade individual momentary outcomes off each other, but the ability to think of such outcomes as forming a pattern through time that is consistent with the way they want to shape their lives. The broad idea that people's current choices are influenced by how they want their lives to unfold over time has featured centrally in a number of different theoretical frameworks in social psychology (e.g., Markus & Nurius, 1986; Oyserman & James, 2011), and we will not defend this idea here. Rather, we want to emphasize the way of thinking about time that this involves, and the role that way of thinking of time therefore plays in intertemporal choice.

On this picture, in deciding what choice is the best fit with how one wants the temporally extended project of one's life to unfold, times are represented as unique, unrepeating points on a time line of one's life, i.e., as what we have termed above as particular times. This becomes vivid if we consider that, in some circumstances, what makes intertemporal choices important is the lack of an opportunity to revisit them (e.g., deciding whether to spend now rather than save for one's pension, deciding whether to have children). When making intertemporal choices using the temporal reasoning system, choice points are considered as bifurcation points in linear time, with choices determining whether one's life unfolds one way or another. This can indeed sometimes result in the decision to delay gratification, but notably this second mechanism for mitigating temporal discounting is in an important way more flexible than the first one we described involving the temporal updating system. As pointed out by Bulley, Henry, and Suddendorf (2016), depending on how the future is likely to unfold, it may sometimes be the right choice to take the smaller, sooner reward instead of waiting (perhaps in vain) for the larger, later one. Thus, whilst temporal reasoning can facilitate delay gratification, it might also sometimes make it apparent that it is better to ‘seize the day’.

In this section, we have distinguished between two mechanisms facilitating delay of gratification – one that relies on the temporal updating system by bringing the future reward into the current model of world, and one that relies on the temporal reasoning system by allowing people to think about their lives as time lines along which they decide to have a particular temporal profile of reward. This raises important questions about the circumstances in which these mechanisms might be put to use, and about how they are related to each other. With regard to facilitating delay of gratification, we do not view these mechanisms as typically being in conflict with each other. Indeed, we have stressed that the first mechanism draws on the temporal reasoning system to initiate a simulation of a future reward that can then fall under the purview of the temporal updating system. Which mechanism is effective in helping people delay gratification may depend on the nature of the choice. For example, it may be that it is particularly helpful to delay gratification of purely hedonic rewards by imagining those future rewards as actually present right now, whereas delay of gratification for rewards that are meaningful primarily in the larger context of one’s life or individual identity might be best facilitated by thinking about one’s life as a temporally extended project. There also may be group or individual differences; for example, amnesic patients who lack the ability to engage in detailed simulation of future events may base their intertemporal choices primarily on reasoning about the time line of their future lives (this type of reasoning appears to be intact in such patients; Craver, Kwan, Steindam, & Rosenbaum, 2014).

6. Concluding remarks

We have argued that our dual systems approach helps to shed light on a variety of issues ranging from how to characterize animal cognition to the metaphysical assumptions that seem to be part of people's naïve theory of time.

We have sided with those who reject the idea that animals are capable of mental time travel, arguing that animals are not capable of thinking about the past or the future at all. Part of what motivates our arguments is the idea that it is at least not obvious how much use animals would have for the idea of particular past and future times different from the present (on this, see also Campbell, 1996). What we have said in this paper does imply that there are some benefits that come with being able to engage in temporal reasoning, for instance when it comes to dealing with situations in which information about successive changes is received in a different order from the one in which they happened. However, we hypothesize that temporal reasoning abilities have not evolved in animals because opportunities to benefit from knowing that a situation of a particular kind obtained at a unique time in the past are relatively rare, since that time itself will never come around again. By contrast, there are obvious benefits in possessing a general learning system geared towards encoding and retaining information about regular, stable, or reoccurring features of the environment, because such information may be of use on numerous occasions when these features are encountered again.

As humans, we have developed the ability to make time itself an object of thought, to think of the world as extended in four dimensions, one of which is the temporal one.

Given what we have just said, we think it is right that questions as to the primary adaptive function that this ability has evolved to serve in humans' lives have recently started to attract researchers' attention (see, e.g., Mahr & Csibra, 2018). However, as we have tried to argue, there are also reasons for thinking that the more primitive temporal updating system that animals rely on in negotiating the world, in which time is not represented, is still active in humans too, alongside the capacity for temporal reasoning. And this may be part of the explanation as to why time remains to be a phenomenon we can get deeply puzzled by.

References

Aristotle (1930). On memory and reminiscence (J. I. Beare, Trans.). In W. D. Ross (Ed.), *The works of Aristotle (vol. 3)*. Oxford, England: Oxford University Press.

Atance, C. M., Louw, A., & Clayton, N. S. (2015). Thinking ahead about where something is needed: New insights about episodic foresight in preschoolers. *Journal of Experimental Child Psychology, 129*, 98-109. doi:10.1016/j.jecp.2014.09.001

Aveni, A. (1990). *Empires of time: Calendars, clock and cultures.* London, England: Tauris.

Babb, S. J., & Crystal, J. D. (2006). Episodic-like memory in the rat. *Current Biology*, *16*(13), 1317-1321. doi:10.1016/j.cub.2006.05.025

Bardon, A. (2013). *A brief history of the philosophy of time*. Oxford, England: Oxford University Press. doi:10.1093/acprof:oso/9780199976454.001.0001

Barr, R., Dowden, A., & Hayne, H. (1996). Developmental changes in deferred imitation by 6- to 24-month-old infants. *Infant Behavior & Development, 19*(2), 159-170. doi:10.1016/s0163-6383(96)90015-6

Bauer, P. J. (2007). Recall in infancy: A neurodevelopmental account. *Current Directions in Psychological Science*, *16*(3), 142-146. doi: 10.1111/j.1467-8721.2007.00492.x

Bauer, P. J., Hertsgaard, L. A., Dropik, P., & Daly, B. P. (1998). When even arbitrary order becomes important: Developments in reliable temporal sequencing of arbitrarily ordered events. *Memory, 6*(2), 165-198. doi:10.1080/741942074

Bauer, P. J., & Mandler, J. M. (1989). One thing follows another: Effects of temporal structure on 1- to 2-year-olds' recall of events. *Developmental Psychology,* *25*(2), 197-206. doi:10.1037/0012-1649.25.2.197

Bauer, P. J., Wenner, J. A., Dropik, P. L., & Wewerka, S. S. (2000). Parameters of remembering and forgetting in the transition from infancy to early childhood. *Monographs of the Society for Research in Child Development, 65*(4), i-vi + 1-213. doi:10.1111/1540-5834.00104

Bennett, J. (1964). *Rationality: An essay towards an analysis*. London, England: Routledge & Kegan Paul.

Benoit, R. G., Gilbert, S. J., & Burgess, P. W. (2011). A neural mechanism mediating the impact of episodic prospection on farsighted decisions. *Journal of Neuroscience, 31*(18), 6771-6779. doi:10.1523/jneurosci.6559-10.2011

Bergson, H. (1911). *Matter and memory* (N. M. Paul & W. S. Palmer, Trans.). London, England: George Allen and Unwin. doi:10.1037/13803-000

Blything, L. P., & Cain, K. (2016). Children's processing and comprehension of complex sentences containing temporal connectives: the influence of memory on the time course of accurate responses. *Developmental Psychology, 52*(10), 1517-1529. doi:10.1037/dev0000201

Blything, L. P., Davies, R., & Cain, K. (2015). Young children's comprehension of temporal relations in complex sentences: the influence of memory on performance. *Child Development, 86*(6), 1922-1934. doi:10.1111/cdev.12412

Boroditsky, L. (2011). How languages construct time. In S. Dehaene & E. Brannon (Eds.), Space, time and number in the brain (pp. 333-341). Oxford: Oxford University Press.

Boyer, P. (2008). Evolutionary economics of mental time travel? *Trends in Cognitive Sciences, 12*(6), 219-224. doi:10.1016/j.tics.2008.03.003

Braddon-Mitchell, D., & Miller, K. (2017). On time and the varieties of science. In C. Bouton & P. Huneman (Eds.), *Time of nature and the nature of time* (pp. 67-85). Cham, Switzerland: Springer. doi:10.1007/978-3-319-53725-2_5

Bradshaw, W. E., & Holzapfel, C. M. (2010). What season is it anyway? Circadian tracking vs. photoperiodic anticipation in insects. *Journal of Biological Rhythms, 25*(3),155-165. doi:10.1177/0748730410365656

Buhusi, C. V., & Meck, W. H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nature Reviews Neuroscience, 6*(10), 755-765. doi:10.1038/nrn1764

Bulley, A., Henry, J., & 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. doi:10.1037/gpr0000061

Burns, P., Russell, C., & Russell, J. (2015). Preschool children's proto-episodic memory assessed by deferred imitation. *Memory, 23*(8), 1172-1192. doi:10.1080/09658211.2014.963625

Busby-Grant, J., & Suddendorf, T. (2009). Preschoolers begin to differentiate the times of events from throughout the lifespan. *European Journal of Developmental Psychology, 6*(6), 746-762. doi:10.1080/17405620802102947

Busby-Grant, J., & Suddendorf, T. (2011). Production of temporal terms by 3-, 4-, and 5-year-old children. *Early Childhood Research Quarterly, 26*(1), 87-95. doi: 10.1016/j.ecresq.2010.05.002

Callender, C. (2017). *What makes time special?* Oxford, England: Oxford University Press. doi:10.1093/oso/9780198797302.001.0001

Campbell, J. (1996). Human vs. animal time. In M. A. Pastor & J. Artieda (Eds.), *Time, internal clocks and movement* (pp. 115-126). Amsterdam, The Netherlands: Elsevier. doi:10.1016/S0166-4115(96)80055-0

Carnap, R. (1963). Intellectual autobiography. In: P. A. Schilpp (Ed.), *The philosophy of Rudolf Carnap* (pp. 3-84). La Salle, IL: Open Court.

Carroll, S. (2010). *From eternity to here: The quest for the ultimate theory of time*. Oxford, England: Oneworld Publications.

Cheke, L. G., & Clayton, N. S. (2010). Mental time travel in animals. *Wiley Interdisciplinary Reviews: Cognitive Science, 1*(6), 915-930. doi:10.1002/wcs.59

Cheke, L.G. & Clayton N.S. (2012). Eurasian jays (*Garrulus glandarius*) overcome their current desires to anticipate two distinct future needs and plan for them appropriately. *Biology Letters,* 8(2), 171-175. doi:10.1098/rsbl.2011.0909

Chen, M. K. (2013). The effect of language on economic behavior: Evidence from savings rates, health behaviors, and retirement assets. *American Economic Review, 103*(2), 690-731. doi:10.1257/aer.103.2.690

Clayton, N. S., Bussey, T. J., & Dickinson, A. (2003). Can animals recall the past and plan for the future? *Nature Reviews Neuroscience, 4*(8), 685-691. doi:10.1038/nrn1180

Clayton, N. S., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature, 395*(6699), 272-274. doi:10.1038/26216

Clayton, N. S., Yu, K. S., & Dickinson, A. (2003). Interacting cache memories: Evidence for flexible memory use by Western Scrub-Jays (*Aphelocoma californica*). *Journal of Experimental Psychology: Animal Behavior Processes, 29*(1), 14-22. doi:10.1037//0097-7403.29.1.14

Correia, S. P. C., Dickinson, A., & Clayton, N. S. (2007). Western scrub-jays anticipate future needs independently of their current motivational state. *Current Biology, 17*(10), 856-861. doi:10.1016/j.cub.2007.03.063

Craver, C.F. Kwan, D., Steindam, C., & Rosenbaum, R.S. (2014). Individuals with episodic amnesia are not stuck in time. *Neuropsychologia, 57*, 191-195. doi:10.1016/j.neuropsychologia.2014.03.004

Cromer, R. F. (1971). The development of the ability to decenter in time. *British Journal of Psychology, 62*(3), 353-365. doi:10.1111/j.2044-8295.1971.tb02046.x

de Hevia, M. D., Izard, V., Coubart, A., Spelke, E. S., & Streri, A. (2014). Representations of space, time, and number in neonates. *Proceedings of the National Academy of Sciences of the United States of America, 111*(13), 4809-4813. doi:10.1073/pnas.1323628111

Dennett, D. C. (1969). *Content and consciousness*. London, England: Routledge & Kegan Paul.

Dufour, V., & Sterck, E. H. M. (2008). Chimpanzees fail to plan in an exchange task but succeed in a tool-using procedure. *Behavioural Processes, 79*(1), 19-27. doi:10.1016/j.beproc.2008.04.003

Eacott, M. J., Easton, A., & Zinkivskay, A. (2005). Recollection in an episodic-like memory task in the rat. *Learning & Memory, 12*(3), 221-223. doi:10.1101/lm.92505

Evans, J. S. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science, 8*(3), 223-241. doi:10.1177/1745691612460685

Falk, A. (2003). Time plus the whoosh and whiz. In A. Jokić & Q. Smith (Eds.), *Time, tense, and reference* (pp. 211-250). Cambridge, MA: MIT Press.

Fivush, R. (2011). The development of autobiographical memory. *Annual Review of Psychology, 62*, 559-582. doi:10.1146/annurev.psych.121208.131702

Fraisse, P. (1964). *The psychology of time.* London, England: Eyre & Spottiswoode.

Friedman, W. J. (1982). Conventional time concepts and children's structuring of time. In W. J. Friedman (Ed.), *The developmental psychology of time* (pp. 171-208). New York, NY: Academic Press.

Friedman, W. J. (1991). The development of children's memory for the time of past events. *Child Development, 62*(1), 139-155. doi:10.2307/1130710

Friedman, W. J. (2000). The development of children's knowledge of the times of future events. *Child Development, 71*(4), 913-932. doi:10.1111/1467-8624.00199

Friedman, W. J. (2001). Memory processes underlying humans' chronological sense of the past. In C. Hoerl & T. McCormack (Eds.), *Time and memory: Issues in philosophy and psychology* (pp. 139-167). Oxford, England: Oxford University Press.

Friedman, W. J. (2003). The development of a differentiated sense of the past and future. In: R. V. Kail (Ed.), *Advances in child development and behaviour, Volume 31* (pp. 232-272). San Diego, CA: Academic Press.

Friedman, W. J., Gardner, A. G., & Zubin, N. R. E. (1995). Children's comparisons of the recency of two events from the past year. *Child Development, 66*(4), 970-983. doi:10.1111/j.1467-8624.1995.tb00916.x

Friedman, W. J., & Kemp, S. (1998). The effects of elapsed time and retrieval on young children's judgments of the temporal distances of past events. *Cognitive Development, 13*(3), 335-367. doi:10.1016/s0885-2014(98)90015-6

Galton, A. (2011). Time flies but space does not: Limits to the spatialisation of time. *Journal of Pragmatics, 43*(3), 695-703. doi:10.1016/j.pragma.2010.07.002

Gell, A. (1992). *The anthropology of time: Cultural constructions of temporal maps and images.* Oxford, England: Berg.

Gelman, S. A. (2006). Naive theories, development of. In L. Nadel (Ed.), *Encyclopedia of Cognitive Science*. Chichester, England: Wiley doi:10.1002/0470018860.s00520

Grondin, S. (2010). Timing and time perception: a review of recent behavioral and neuroscience findings and theoretical directions. *Attention, Perception, & Psychophysics*, *72*(3), 561-582. doi: 10.3758/APP.72.3.561

Harner, L. (1980). Comprehension of past and future reference revisited. *Journal of Experimental Child Psychology, 29*(1), 170-182. doi:10.1016/0022-0965(80)90099-5

Harner, L. (1982). Talking about the past and future. In W. J. Friedman (Ed.), *The developmental psychology of time* (pp. 141-170). New York, NY: Academic Press.

Hayne, H., Gross, J., McNamee, S., Fitzgibbon, O., & Tustin, K. (2011). Episodic memory and episodic foresight in 3-and 5-year-old children. *Cognitive Development*, *26*(4), 343-355. doi:10.1016/j.cogdev.2011.09.006

Hoerl, C. (2008). On being stuck in time. *Phenomenology and the Cognitive Sciences, 7*(4), 485-500. doi:10.1007/s11097-008-9089-z

Hoerl, C., & McCormack, T. (2005). Joint reminiscing as joint attention to the past. In N. Eilan, C. Hoerl, T. McCormack, & J. Roessler (Eds.), *Joint Attention: Communication and Other Minds: Issues in Philosophy and Psychology* (pp. 260-286). Oxford, England: Clarendon Press.

doi:10.1093/acprof:oso/9780199245635.003.0012

Howe, M. L., & Courage, M. L. (1997). The emergence and early development of autobiographical memory. *Psychological Review, 104*(3), 499-523. doi:10.1037/0033-295x.104.3.499

Hudson, J. A., & Mayhew, E. M. Y. (2011). Children’s temporal judgments for autobiographical past and future events. *Cognitive Development, 26*(4), 331-342. doi:10.1016/j.cogdev.2011.09.005

Ismael, J. (2017). Passage, flow and the logic of temporal perspectives. In C. Bouton & P. Huneman (Eds.), *Time of nature and the nature of time* (pp. 23-38). Cham, Switzerland: Springer. doi:10.1007/978-3-319-53725-2_2

Kabadayi, C., & Osvath, M. (2017). Ravens parallel great apes in flexible planning for tool-use and bartering. *Science, 357*(6347), 202-204. doi:10.1126/science.aam8138

Kaller, C. P., Rahm, B., Spreer, J., Mader, I., & Unterrainer, J. M. (2008). Thinking around the corner: The development of planning abilities. *Brain and Cognition, 67*(3), 360-370. doi:10.1016/j.bandc.2008.02.003

Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus and Giroux.

Keven, N. (2016). Events, narratives and memory. *Synthese, 193*(8), 2497-2517. doi:10.1007/s11229-015-0862-6

Leininger, L. (2015). Presentism and the myth of passage. *Australasian Journal of Philosophy, 93*(4), 724-739. doi:10.1080/00048402.2015.1007463

Lin, H., & Epstein, L. H. (2014). Living in the moment: effects of time perspective and emotional valence of episodic thinking on delay discounting. *Behavioral Neuroscience, 128*(1), 12-19. doi:10.1037/a0035705

Lohse, K., Kalitschke, T., Ruthmann, K., & Rakoczy, H. (2015). The development of reasoning about the temporal and causal relations among past, present, and future events. *Journal of Experimental Child Psychology, 138*, 54-70. doi:10.1016/j.jecp.2015.04.008

Mahr, J. B., & Csibra, G. (2018): Why do we remember? The communicative function of episodic memory. *Behavioral and Brain Sciences 41*, e1. doi:10.1017/S0140525X17000012

Markus, H., & Nurius, P. (1986). Possible selves. *American Psychologist, 41*(9), 954-969. doi:10.1037//0003-066x.41.9.954

Martin-Ordas, G. (2017). 'Will I want these stickers tomorrow?' Preschoolers' ability to predict current and future needs. *British Journal of Developmental Psychology, 35*(4), 568-581. doi:10.1111/bjdp.12195

Martin-Ordas, G., Atance, C. M., & Caza, J. (2017). Did the popsicle melt? Preschoolers' performance in an episodic-like memory task. *Memory, 25*(9), 1260-1271. doi:10.1080/09658211.2017.1285940

McCarty, M. E., Clifton, R. K., & Collard, R. R. (1999). Problem solving in infancy: The emergence of an action plan. *Developmental Psychology, 35*(4), 1091-1101. doi:10.1037/0012-1649.35.4.1091

McColgan, K. L., & McCormack, T. (2008). Searching and planning: Young children's reasoning about past and future event sequences. *Child Development*, *79*(5), 1477-1497. doi:10.1111/j.1467-8624.2008.01200.x

McCormack, T. (2001). Attributing episodic memory to animals and children. In C. Hoerl & T. McCormack (Eds.), *Time and memory: Issues in philosophy and psychology* (pp. 285-313). Oxford, England: Clarendon Press.

McCormack, T. (2015). The development of temporal cognition. In R.M. Lerner, L. S. Liben, & U. Mueller (Eds.), *Handbook of child psychology and developmental science, Volume 2: Cognitive processes* (7th ed.) (pp. 624-670). Malden, MA: Wiley-Blackwell.

McCormack, T., & Hanley, M. (2011). Children's reasoning about the temporal order of past and future events. *Cognitive Development*, *26*(4), 299-314. doi:10.1016/j.cogdev.2011.10.001

McCormack, T. & Hoerl, C. (1999). Memory and temporal perspective: The role of temporal frameworks in memory development. *Developmental Review 19*, 154-182. doi:10.1006/drev.1998.0476

McCormack, T. & Hoerl, C. (2001). The child in time: temporal concepts and self-consciousness in the development of episodic memory. In: C. Moore and K. Lemmon (eds.), *The self in time: developmental perspectives* (pp. 203-227). Mahwah, NJ: Lawrence Erlbaum.

McCormack, T., & Hoerl, C. (2005). Children's reasoning about the causal significance of the temporal order of events. *Developmental Psychology*, *41*(1), 54-63. doi:10.1037/0012-1649.41.1.54

McCormack, T., & Hoerl, C. (2007). Young children's reasoning about the order of past events. *Journal of Experimental Child Psychology, 98*(3), 168-183. doi:10.1016/j.jecp.2007.06.001

McCormack, T., & Hoerl, C. (2017). The development of temporal concepts: Learning to locate events in time. *Timing and Time Perception, 5*(3-4): 297-327. doi:10.1163/22134468-00002094

McTaggart, J.E. (1908). The unreality of time. *Mind, 17*(4), 457-474. doi:10.1093/mind/XVII.4.457

Metcalfe, J., & Mischel, W. (1999). A hot/cool-system analysis of delay of gratification: dynamics of willpower. *Psychological Review, 106*(1), 3-19. doi:10.1037/0033-295X.106.1.3

Mulcahy, N. J., & Call, J. (2006). Apes save tools for future use. *Science, 312*(5776), 1038-1040. doi:10.1126/science.1125456

Nelson, K. (1996). *Language in cognitive development: The emergence of the mediated mind.* Cambridge, England: Cambridge University Press.

Osvath, M., & Osvath, H. (2008). Chimpanzee (*Pan troglodytes*) and orangutan (*Pongo abelii*) forethought: self-control and pre-experience in the face of future tool use. *Animal Cognition, 11*(4), 661-674. doi:10.1007/s10071-008-0157-0

Oyserman, D., & James, L. (2011). Possible identities. In S. Schwartz, K. Luyckx, & V. Vignoles (Eds.), *Handbook of identity theory and research* (pp. 117-145). New York, NY: Springer.

Pathman, T., Larkina, M., Burch, M. M., & Bauer, P. J. (2013). Young children's memory for the times of personal past events. *Journal of Cognition and Development, 14*(1), 120-140. doi:10.1080/15248372.2011.641185

Peacocke, C. (2017). Temporal perception, magnitudes and phenomenal externalism. In I. Phillips (Ed.), *The Routledge handbook of philosophy of temporal experience* (pp. 213-224). Abingdon, England: Routledge.

Perner, J. (1991). *Understanding the representational mind*. Cambridge, MA: MIT Press.

Peters, J., & Büchel, C. (2010). Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-mediotemporal interactions. *Neuron, 66*(1), 138-148. doi:10.1016/j.neuron.2010.03.026

Peterson, C. (2002). Children's long-term memory for autobiographical events. *Developmental Review, 22*(3), 370-402. doi:10.1016/s0273-2297(02)00007-2

Piaget, J. (1969). *The child's conception of time.* London, England: Routledge & Kegan Paul.

Povinelli, D. J., Landry, A. M., Theall, L. A., Clark, B. R., & Castille, C. M. (1999). Development of young children's understanding that the recent past is causally bound to the present. *Developmental Psychology, 35*(6), 1426-1439. doi: 10.1037 /0012-1649.35.6.1426

Price, H. (2011). The flow of time. In C. Callender (Ed.), *The Oxford Handbook of Philosophy of Time* (pp. 276-311). Oxford, England: Oxford University Press. doi:10.1093/oxfordhb/9780199298204.003.0010

Prosser, S. (2006). Temporal metaphysics in Z-land. *Synthese, 149*(1), 77-96. doi: 10.1007/s11229-004-6249-8

Pyykkönen, P., & Järvikivi, J. (2012). Children and situation models of multiple events. *Developmental Psychology, 48*(2), 521-529. doi:10.1037/a0025526

Raby, C.R., Alexis, D.M., Dickinson, A., & Clayton N. S. (2007). Planning for the future by western scrub-jays. *Nature, 445*(7130), 919-921. doi:10.1038/nature05575

Read, D., Frederick, S., Orsel, B., & Rahman, J. (2005). Four score and seven years from now: The date/delay effect in temporal discounting. *Management Science, 51*(9), 1326-1335. doi:10.1287/mnsc.1050.0412

Redshaw, J. (2014). Does metarepresentation make human mental time travel unique? *Wiley Interdisciplinary Reviews: Cognitive Science, 5*(5), 519-531. doi:10.1002/wcs.1308

Redshaw, J., & Suddendorf, T. (2016). Children's and apes' preparatory responses to two mutually exclusive possibilities. *Current Biology, 26*(13), 1758-1762. doi: 10.1016/j.cub.2016.04.062

Redshaw, J., Taylor, A. H., & Suddendorf, T. (2017). Flexible planning in ravens? *Trends in Cognitive Sciences, 21*(11), 821-822. doi:10.1016/j.tics.2017.09.001

Roberts, W. A. (2002). Are animals stuck in time? *Psychological Bulletin, 128*(3), 473-489. doi:10.1037//0033-2909.128.3.473

Rovee-Collier, C. (1999). The development of infant memory. *Current Directions in Psychological Science, 8*(3), 80-85. doi:10.1111/1467-8721.00019

Rovelli, C. (2018). *The order of time*. New York: Riverhead Books.

Russell, J., Alexis, D., & Clayton, N. (2010). Episodic future thinking in 3-to 5-year-old children: The ability to think of what will be needed from a different point of view. *Cognition, 114*(1), 56-71. doi:10.1016/j.cognition.2009.08.013

Russell, J., Cheke, L. G., Clayton, N. S., & Meltzoff, A. N. (2011). What can What-When-Where (WWW) binding tasks tell us about young children's episodic foresight? Theory and two experiments. *Cognitive Development, 26*(4), 356-370. doi:10.1016/j.cogdev.2011.09.002

Sachs, J. (1983). Talking about the there and then: The emergence of displaced reference in parent-child discourse. In K. E. Nelson (Ed.), *Children's language* (Vol. 4, pp. 1-28). Hillsdale, NJ: Erlbaum.

Salwiczek, L. H., Watanabe, A., & Clayton, N. S. (2010). Ten years of research into avian models of episodic-like memory and its implications for developmental and comparative cognition. *Behavioural Brain Research, 215*(2), 221-234. doi:10.1016/j.bbr.2010.06.011

Skow, B. (2011). Experience and the passage of time. *Philosophical Perspectives, 25*(1), 359-387. doi:10.1111/j.1520-8583.2011.00220.x

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119*(1), 3-22. doi:10.1037/0033-2909.119.1.3

Smith, E. R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review, 4*(2), 108-131. doi:10.1207/s15327957pspr0402_01

Stich, S. P. (1978). Beliefs and subdoxastic states. *Philosophy of Science, 45*(4), 499-518. doi:10.1086/288832

Suddendorf, T., & 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. doi:10.1017/s0140525x07001975

Suddendorf, T., & Corballis, M. C. (2010). Behavioural evidence for mental time travel in nonhuman animals. *Behavioural Brain Research, 215*(2), 292-298. doi:10.1016/j.bbr.2009.11.044

Tecwyn, E. C., Thorpe, S. K. S., & Chappell, J. (2014). Development of planning in 4- to 10-year-old children: Reducing inhibitory demands does not improve performance. *Journal of Experimental Child Psychology, 125*, 85-101. doi:10.1016/j.jecp.2014.02.006

Tillman, K. A., Marghetis, T., Barner, D., & Srinivasan, M. (2017). Today is tomorrow's yesterday: Children's acquisition of deictic time words. *Cognitive Psychology, 92*, 87-100. doi:10.1016/j.cogpsych.2016.10.003

Tulving, E. (2005) Episodic memory and autonoesis: Uniquely human? In H. S. Terrace & J. Metcalfe (Eds.), *The missing link in cognition* (pp. 3-56). Oxford, England: Oxford University Press.

Welch-Ross, M. K. (2001). Personalizing the temporally extended self: Evaluative self-awareness and the development of autobiographical memory. In C. Moore & K. Lemmon (Eds.), *The self in time: Developmental perspectives* (97-120). Mahwah, NJ: Lawrence Erlbaum.

Weist, R. M. (1989). Time concepts in language and thought: Filling the Piagetian void between two and five years. In I. Levin & D. Zakay (Eds.), *Time and human cognition: A life-span perspective* (pp. 63-118). Amsterdam, The Netherlands: Elsevier.

Weist, R. M., & Buczowska, E. (1987). The emergence of temporal adverbs in child Polish. *First language*, *7*, 217-229. doi: 10.1177/014272378700702105

Weist, R. M., & Zevenbergen, A. A. (2008). Autobiographical memory and past time reference. *Language Learning and Development*, *4*(4), 291-308. 10.1080/15475440802293490

Zentall, T. R. (2005). Animals may not be stuck in time. *Learning and Motivation, 36*(2), 208-225. doi:10.1016/j.lmot.2005.03.001

Zimmerman, D. W. (2005). The A-theory of time, the B-theory of time, and 'taking tense seriously'. *Dialectica*, *59*(4), 401-457. doi:10.1111/j.1746-8361.2005.01041.x