Word counts

Abstract: 161; Main text: 12,543; References: 897; Total: 13,626

Temporal updating, temporal reasoning and the domain of time

Christoph Hoerl & Teresa McCormack

Abstract: We focus on three main sets of topics emerging from the commentaries on our target article. First, we discuss several types of animal behavior that commentators cite as evidence against our claim that animals are restricted to temporal updating and cannot engage in temporal reasoning. In doing so we illustrate further how explanations of behavior in terms of temporal updating work. Second, we respond to commentators' queries about the developmental process through which children acquire a capacity for temporal reasoning and about the relation between our account and accounts drawing similar distinctions in other domains of cognition. Finally, we address some broader theoretical issues arising from the commentaries, concerning in particular the question as to how our account relates to the phenomenology of experience in time, and the question as to whether our dichotomy between temporal reasoning and temporal updating is exhaustive, or whether there might be other forms of cognition or representation related to time not captured by it.

R1. Introduction

Adult human beings can think of the world they live in as being extended in four dimensions, one of which is the temporal one. Doing so involves a rich variety of reasoning abilities, such as the ability to use *tense* to think about what was the case in the past and what will be the case in the future as well as about what is present, a grasp of events as happening at

particular times, each of which only comes round once, and the ability to represent *temporal order relations* between events, knowledge of which may be crucial to being able to infer what is the case now from information about what happened in the past. In our target article, we contrasted temporal reasoning, thus understood, with a more basic form of cognition, which we call temporal updating. In short, the model of the world that a creature capable only of temporal updating operates with simply lacks a temporal dimension. It concerns the world only ever as it is at present. However, because this model is updated as the creature receives new information and because timing mechanisms can also govern the instantiation of certain elements within the model, a creature equipped only with a temporal updating system can still display behavior that is adapted to how things unfold over time.

Our commentators – even those broadly sympathetic to our account – lay down a number of challenges for us. In this response, we discuss them under three broad headings into which they can be grouped: (i) Explaining animal behavior in time, (ii) Developmental considerations, and (iii) Dual systems, representation and phenomenology. With regard to the first of these, in our target article, we had sided with those who claim that animals are incapable of mental time travel, which, within the context of our model, we elaborated as the claim that animals are capable of temporal updating only. In section R2, we return to this claim and extend our account to three forms of animal behavior that we had not discussed in our target article, which commentators cite as evidence against us. In doing so, we also illustrate some more general features of explanations of behaviour in terms of temporal updating. We address the second set of issues, i.e., developmental ones, in section R3, focusing in particular on commentators' queries about the developmental process through which children acquire a capacity for temporal reasoning and about the relationship between our account and accounts drawing similar distinctions in other domains of cognition. Finally, in section R4, we discuss a set of more general theoretical questions about our account, and

in particular two common themes emerging from a number of different commentaries. The first is how our account relates to questions about the phenomenology of experience in time; and the second is whether there might be other forms of cognition or representation related to time that escape our dichotomy between temporal reasoning and temporal updating. In addressing this second theme we return once more to the issue of animal abilities, and explain some of the reasons underlying our view that animals do not just lack temporal reasoning abilities, as we describe them, but that they also do not possess any other, more basic, ways of representing the domain of times.

R2. Explaining animal behavior in time

Before considering some specific types of animal behavior described by commentators as challenges to our account, it may be useful to illustrate with an example the difference between a temporal reasoning explanation and a temporal updating explanation of behavior that unfolds over time. A dog buries a bone and later returns to dig for it. This is a very simple case in which the question can be raised as to whether a creature's behavior demonstrates a capacity for temporal reasoning or whether it can also be explained in a way that credits the creature only with a capacity for temporal updating. On a temporal updating account, the dog initially operates with a model of the world that does not represent a bone as being at the relevant location, but, after burying, its model of the world gets updated to one that does represent a bone as being at that location, and this representation is what explains subsequent digging. A temporal reasoning account of the dog's behaviour would explain its digging where it does by crediting it with a past tense belief that it buried the bone there. Note, however, that this explanation is not complete unless the dog is credited with another, general, belief, such as 'buried bones stay put'. In order to be motivated to dig at the relevant

location the dog must not just believe that a bone was buried there in the past; it also needs to have a reason for thinking that the bone is still there now (see Bennett, 1964, for discussion).

There are thus two quite different types of explanations of the dog's digging behaviour: Simplifying somewhat,[1] the first appeals to just one *tenseless* belief, and to the idea that the dog acquired that belief in the past and that that belief has persisted over time, whereas the second credits the dog with two beliefs – a past tense belief and a general belief (see also Smith, 1982). On a more abstract level, the first type of explanation is one that appeals to the conditions that determine the instantiation and persistence of a creature's representations over time, rather than appealing to how time itself figures in those representations. Thus described, the buried bone example emerges as just one kind on a much broader spectrum, and the same general type of explanation can also be applied to other, more sophisticated-looking, kinds of behavior. In this section, we will consider three such kinds of behaviour described in the commentaries: hummingbirds' foraging behavior (R2.1), delay of gratification in ravens (R2.2), and orangutans' long calls (R2.3).

R2.1. Temporal updating and questions of explanation – the case of foraging behavior

Pan and Carruthers ask how our account would explain the foraging behavior of hummingbirds, which have been shown to be sensitive to the rate at which flowers replenish. As they rightly assume, our account would assume that the birds possess (something like) a timer that gets entrained with the refresh rate of the flower as the bird visits the flower at different intervals after feeding and finds it full again only at some of those intervals. On a

[1] We frame the distinction here in terms of two explanations both invoking beliefs, but we want to allow that a temporal updating account could also invoke other types of representational states. This is why we say we are simplifying here.

temporal updating account, the function of this timer is to make it the case that at times after feeding that correspond to intervals at which the flower has been found empty, the flower is not represented as full, whereas at times corresponding to intervals at which the flower has been found to be full it is represented as full.

Pan and Carruthers say that they "have no idea how [we] think this is supposed to happen". We think this might be because they don't distinguish clearly enough between two quite different types of questions. Call the first the 'mechanism question'. In the case of the dog digging for the bone, which we mentioned at the start of this section, the mechanism question is the question answered by saying that it is simply the tendency of beliefs to persist over time that explains why, even after some time has elapsed, the dog still believes of a specific location that it contains a buried bone. That is, 'how it happens' that it has that belief at the relevant time is that it had that belief previously, and that belief has persisted over time. Similarly, in the case of the hummingbirds, we suggest that the mechanism question is answered simply by saying that particular states of a timer that gets triggered upon feeding from the flower have become associated with a representation of it as empty and others have become associated with a representation of it as full. As a result, the timer being in the former set of states will cause the bird to represent the flower as empty, and the timer being in the latter set of states will cause it to represent the flower as full. 'How this happens' is explained by the fact that the relevant two types of representations have become associated with the relevant two sets of states of the timer.

Contrast the mechanism question with another type of question that can take the form 'How does it happen…?' Again, going back to our original example of the dog, one can ask why it actually is that the belief about where the bone is buried persists. Why does the dog not

simply lose that belief as soon as the bone is out of sight? We might call this the 'sensitivity question', as it concerns the fact that the persistence conditions of the belief about the bone are somehow sensitive to the persistence conditions of the bone itself. We are happy to grant that such a question can be raised for the simple temporal updating account of the dog's behavior that we have given. However, it should be very clear that exactly the same question also arises for any account that explains the dog's behavior in terms of a belief, on the part of the dog, that it buried the bone in the relevant location in the past. As we said above, such an explanation is not complete unless the dog is credited with the general belief 'buried bones tend to stay put', and clearly there is no reason to ascribe such a belief to the dog either unless it is somehow sensitive to the persistence conditions of buried bones.

Similarly, if it is the sensitivity question that **Pan and Carruthers** are after by asking how our account explains the hummingbirds' behavior, the question is how it happens that the birds' timer comes to be entrained with the refresh rate of the flower in the first place. This is indeed an important question, because, for example, such entrainment requires the bird revisiting the flower at different intervals after feeding, which it might be thought to have no reason to do because it has emptied the flower of nectar. Just as in the case of the dog, though, it is wrong to think that our account faces this type of sensitivity question whereas one that ascribes to the bird representations of the time elapsed since its last visit to the flower doesn't. Pan and Carruthers suggest that the birds have a representation with the content "the flower takes between ten and twenty clicks [of the interval timer] to replenish". Presumably the idea is that their having this representation, together with a representation of how many clicks ago they visited the flower, is what explains their behaviour. But then how does the bird form this representation of how many clicks it takes for the flower to replenish, if not again on the basis of repeated visits to the flower at different intervals after feeding?

Thus, we believe **Pan and Carruthers'** impression that our account fails to give an explanation of 'how it happens' (see also **Kaufmann and Cahen**) may be based on a conflation between two different ways in which this question may be understood. It is indeed true that the answer to the mechanism question never on its own explains animal behavior. Yet it is wrong to think that our account therefore faces an extra explanatory burden, because the sensitivity question, which needs to be answered to provide the full story, is one that alternative accounts of animal behavior face in just the same way.

R2.2. Delay of gratification

Another type of animal behavior that we did not discuss in our target paper is behavior that involves delay of gratification. As **Osvath and Kabadayi** point out, the study of theirs that we used to illustrate experiments on animal tool saving included two conditions in which, amongst the items the ravens could choose other than the tool, there was also a food reward, albeit one that was smaller than the one obtainable with the tool. What Kabadayi and Osvath (2017) argued was that the delay after which the birds were given access to the baited apparatus that could be opened with the tool played a significant role in whether they chose the tool or the food reward: the birds were more likely to choose the immediate food reward in a condition in which they had to wait for 15 minutes than one in which they were given immediate access. Osvath and Kabadayi argue that we ignore the role that delay plays in determining the birds' choice, and hence misdescribe the representations underpinning the birds' behavior. We note that, unless we have misunderstood their method, their own results do not force the conclusion that the birds are sensitive to the magnitude of temporal delay per se. The authors described the immediate access condition as one in which the reward in the apparatus was "*spatio*temporally closer" (Kabadayi & Osvath, 2017, p. 203, emphasis added)

– i.e., closer in space as well as time – because in this condition, the birds walked past and saw the apparatus immediately prior to making their selection, whereas in the 15 minute delay condition they had no such prior perceptual access to the apparatus. The confound of spatial and temporal closeness means that it is not possible to be confident that the birds' behavior was sensitive to the magnitude of the delay.

We will set this point aside because there are numerous other studies that suggest that animals are sensitive to the magnitude of delay periods in delay of gratification procedures (Vanderveldt, Oliveira, & Green, 2016). Let us assume that **Osvath and Kabadayi** are correct that the specific delay at which the ravens get re-introduced to the apparatus has a systematic effect on their choices. What is at issue is whether this means that the ravens must be representing both the current situation and a different, future, situation at which they will be reintroduced to the apparatus, and weighing up how far in the future that situation will be. Ravens, unlike humans, cannot be simply informed about the distance in the future of events. Thus, any difference in their choosing behavior must be based on their prior experiences, when they were either reintroduced to the apparatus almost immediately or only after a 15 minute delay. These different experiences can have an impact on how the apparatus and its associated reward is thus represented, and in turn influence tool choice. One straightforward possibility, compatible with the idea that delayed rewards are discounted, is that as a result of their learning experiences the birds accord less value to the tool when it only yields a reward after delay. Indeed, our account leaves open the possibility that animals could make flexible trade-offs that adjust to the magnitude of delay until reward. In describing our updating explanation of Clayton and Dickinson's studies of scrub jays' memory, we suggested that (something akin to) interval timers may determine the representational contents maintained in the updating system, without the birds being able to represent the time of a previous caching

event. Similarly, it is possible that representations of reward magnitude are shaped during learning by the operation of some sort of timing mechanism, without assuming that animals can represent the times at which future rewards will occur. Assuming that some type of timing mechanism is involved here is not controversial – any account of animal sensitivity to delay in temporal discounting studies must appeal to something along these lines. Our argument is that there is no need to, in addition, assume that animals actually represent future points in time. Rather, such mechanisms could directly affect representations of the magnitude of rewards.

R2.3. Orangutans' long calls

We believe a third type of animal behavior mentioned in the commentaries, too, can be explained by our account, and this time even without an appeal to timing mechanisms. These are orangutans' long calls, as discussed by **Kaufmann and Cahen** (van Schaik, Damerius, & Isler, 2013). Van Schaik et al. found that the direction of males' calls emitted in the evening predicted the main direction in which the animals travelled the next day better than chance. Moreover, other orangutans' travel direction was influenced by males' long calls the previous evening, with females remaining at constant distance from the calling male, whereas other males increased their distance. Kaufmann and Cahen say that "the goal of these long calls appears to be to communicate to female orangutans the male's future travel direction" and that the females, as well as other males, "use this information in planning their own travel". This might make it sound as though the conspecifics can determine the direction of the call and then reason from this that the calling male will be found in a certain area in the future. Yet, there is actually no reason for thinking that the conspecifics can determine the direction of the males' call in the first place. Rather, as van Schaik et al. (2013, p. 2) themselves point out, the angular difference between the direction of the call and the line connecting the

conspecific to the calling male is likely simply to determine the perceived distance from the calling male, without the conspecifics being able to distinguish between the calling male being far away from them or the male being closer to them but calling in a direction away from them. Thus, it is more plausible to think that the representation formed upon hearing the call is simply one of how close the calling male is, which is something that could be done within the temporal updating system. If females simply move towards that perceived location if it appears far away, or subordinates move away from it if it appears near, that will be sufficient to yield the observed effect of females staying within earshot of the calling male and subordinates increasing their distance. Yet this would be behavior simply based on a representation of the calling male's perceived location when the call is heard, rather than a representation of its location at some future time different from the present.

Obviously what also needs to be explained is the fact that the calling males do travel in the direction in which they called, and that the conspecifics adjust their own subsequent movements to the perceived location of the call, even though the call was made the previous evening. Van Scheik et al. (2013, p. 8) describe this as being "consistent with the use of some form of episodic memory", but it is unclear why episodic memory, specifically, should be required to explain these behaviors. Consider van Scheik et al.'s (2013, p. 9) own remark that "the important point is that a male orangutan is maintaining an internally generated directional choice towards a distant target out of current sensory range, over a prolonged time period, despite meandering routes." As this makes clear, the crucial explanatory work is being done by the orangutan *maintaining* this choice over time, i.e. forming a representation of a particular distant location as a good location to visit, retaining that representation, and keeping track of its own position relative to that location. We are back to the general type of

explanation that we used to account for how the dog, after burying the bone in the garden, can subsequently maintain a representation of its location.

R2.4. Animal behavior in time: Moving the debate forward

We hope that in addressing some of the specific challenges raised by the commentators, we have clarified the nature and scope of our account. In each instance, we have suggested that the available evidence does not force the conclusion that animals are capable of temporal reasoning. Of course, our account of these behaviors still leaves a variety of unanswered questions about the precise details of the underpinning mechanisms. However, in our view, mental time travel accounts, when considered properly, also face similar – and indeed further and arguably more difficult – questions.

Osvath and Kabadayi's overall criticism is that we are just re-hashing an outdated claim that animals are stuck in time. However, the two authors do seem to agree with us that what the animals do is different from what humans do. At one point, they also describe their view as one according to which animals might "implicitly represent time, without representing time as such." Part of our aim was precisely to try to move forward debate in this area, which has been dogged by the problem that researchers have resorted to concepts and criteria that are difficult to operationalize. We aim to provide a more specific, empirically tractable, characterization of one sense in which animals might be stuck in time. We describe this as their being unable to engage in temporal reasoning, but we also give some operational criteria of when an organism should be described as only capable of temporal updating. We recognize, though, that our account of animal behavior is likely to be controversial, and in section R4.4, we will return to the question of animal abilities to consider also whether there

are types of temporal cognition that do not fall neatly into the two categories we have described.[2]

R3. Developmental considerations

In this section, we turn to issues that commentators raise about our account and human cognitive development. At its most basic, our claim is that being sensitive to temporal features of the world is not the same thing as being able to think about time itself. The idea that there is a distinction between being behaviorally sensitive to a feature of the world and having a proper concept of that feature appears in a wide variety of contexts in developmental psychology. Two challenges that developmental psychologists who make such distinctions typically face are to (i) make clear what the developmental relation is (if any) between the more basic skills and the fully-fledged conceptual grasp of the domain and (ii) explain the process of developmental change. As pointed out by commentators (**DeNigris and Brooks**; **Hamamouche**; **Hohenberger**; **Hudson, Mayhew, and Zhang**; **Tillman**), there is considerably more to be said about how our account meets these challenges (for further detail, see McCormack, 2015; McCormack & Hoerl, 2017).

In what follows, we will focus in particular on the following developmental issues: whether the distinction we draw between two different systems can be subsumed under a more general type of distinction also applicable to other domains (R3.1); whether children's early abilities to process duration information provides evidence against the developmental picture we sketched in the target article (R3.2); what factors might underpin the emergence of temporal reasoning abilities in children (R3.3), in particular the development of an ability to

[2] As will become clear in that section, we reject Osvath and Kabadayi's suggestion that animals might be described as implicitly representing times not just because of the vagueness of that terminology, but also because of more general consideration about representation.

represent particular times (R3.4); and what significance the ability to represent particular past times has for humans (R3.5).

R3.1. Comparison with other domains of cognition

One key issue raised in a number of commentaries concerns the extent to which our account aligns with claims that have been made regarding other domains of cognition, pointing to a more general type of distinction that the one between temporal reasoning and temporal updating is only one instance of. To give examples from the commentaries: **Redshaw** and colleagues suggest that it might be pertinent in this context to distinguish between skills that do or do not require metarepresentation; **Povinelli, Glorioso, Kuznar and Pavlic** identify our account as aligning with a distinction between abilities that do or do not require constructing higher-order analogical-based relations; and **DeNigris and Brooks** suggest that our distinction is one between symbolic and non-symbolic representational formats, with the developmental transition understood in terms of Karmiloff-Smith's (1994) notion of representational redescription (see also **Hohenberger**). A more domain-specific parallel is drawn by **Lohse, Sixtus, and Lonnemann**, who argue that developmental changes in numerical cognition can be seen as paralleling those we have described in temporal cognition (see also **De Nigris and Brooks**; **Roselli**).

We can certainly see the value of considering how our account aligns with more domain-general claims regarding cognitive development. It is also broadly correct to describe the temporal reasoning system as involving a type of 'higher order' cognition that the updating system does not operate with. However, our account differs from those regarding other domains (such as Theory of Mind, number, and weight) that the commentators refer to. In those instances, it is typically assumed that infants have some primitive representation of the

basic domain, on which the more sophisticated conceptual grasp of the domain builds. For example, in the case of Theory of Mind, on a number of theoretical approaches, young children are assumed to be operating with some more primitive way of representing the relations between agents and aspects of the world; the developmental achievement is to reach a proper grasp of the nature of the relations in question. The input to the more primitive process is assumed to be perceptually-based in some important sense (such as observation of the interactions between agents and objects), providing an initial 'way in' to the domain in question that then provides the foundations for more sophisticated understanding.

By contrast, our view is that the temporal updating system does not provide a similar perceptual 'way in' in the case of time, and that the representations of the domain of time with which the temporal reasoning system operates are not simply some sort of more enriched or explicit or 'redescribed' version of representations that the updating system is operating with (see also **Tillman**). Rather, there is an important sense in which the updating system that very young children possess is not operating with temporal representations at all. Indeed, **Oyserman and Dawson** describe the updating system as 'atemporal'; similarly, **Melnikoff and Bargh** say in their commentary: "To our knowledge, this is the first dual-systems theory to define its systems in terms of possessing versus lacking a single representational dimension."

R3.2. Time as a stimulus dimension vs. time as a framework

Hammamouche argues that our account gets the developmental picture completely wrong: in assuming that infants only have the more primitive updating system, we are ignoring evidence that infants are "well-tuned timers" (see also **Viera and Margolis** for similar points regarding animals). The empirical studies unarguably indicate that very young children

possess mechanisms that process temporal duration information. What should be concluded from this? On an alternative developmental picture, one could hold that the representations the updating system deals with are indeed temporal representations of duration, and that the emergence of the reasoning system results from some more explicit or 'redescribed' version of these representations. Such an alternative picture could potentially be seen as aligning with Droit-Volet's description of children as moving from implicit procedural timing to an explicit grasp of the notion of temporal duration (Droit-Volet & Rattat, 1999). The difficulty with this picture is that it deals only with what we might call time as a *stimulus dimension* rather than time as a *framework* (McCormack, 2015). When time is dealt with as a stimulus dimension, what is processed is the magnitude of stimulus duration; the focus of our paper is with time as a framework in which moments in time have locations relative to each other and relative to the present moment. The latter is what the temporal reasoning system operates with, and our claim is that the updating system does not operate over such a domain, because in the updating system other times are not represented at all.

One might argue that it is misleading to describe the temporal updating system as not operating with temporal representations, because we are happy to accept that such a system operates with mechanisms that process temporal duration information. Of course, if it stipulated by definition that the mechanisms that process duration information operate with temporal representations, then, in this sense, the updating system does represent 'time'. However, it is important to point out that this only allows for a narrow sense in which duration can be said to be represented. In particular, because the updating system does not represent time as a framework, while duration could be processed as a stimulus property it could not be represented as the amount of time elapsed between two temporal locations (because temporal locations are not represented in the updating system). And this suggests

that no proper concept of duration is possible for creatures possessing only the updating system, on the assumption that any conceptual grasp of duration must be grounded in the notion of a period of time between a start point and an end point.

Indeed, these considerations raise the issue of what developmental story should be told about the emergence of the concept of duration itself (it was this concept that was the focus of Piaget's, 1969, work on time). Children do acquire such a concept by the time they are five, as evidenced by, e.g., their ability to make explicit judgments regarding the lengths of events (Friedman, 1990). As yet, we do not know how acquisition of the concept of duration is linked to acquisition of the (other) temporal concepts that the temporal reasoning system operates with (such as concepts of the past and future). We are confident, though, that the reasoning system does not simply deal with representations of time that are somehow more explicit or enriched versions of whatever representations are linked to duration sensitivity in infancy. This is not to say that the temporal reasoning system appears out of the blue. In the paper, we argue that 2 to 3-year-olds may not be full-blown temporal reasoners, but they are not solely reliant on the updating system either. Elsewhere (McCormack, 2015; McCormack & Hoerl, 1999, 2017), we have argued that children of this age have *event-based* temporal cognition, in which they represent events and event sequences and can think of some events as unalterable (a primitive version of the past tense) and some potentially alterable (a primitive version of the future tense). This may be sufficient to underpin some types of future-oriented activity, such as those that involve a sensitivity to the fact that aspects of the environment may change (see **Goulding and Friedman**). However, children of this age will struggle when they have to reason ahead about the temporal order in which future changes will or should happen (Martin-Ordas, 2017). There are two key limitations at this stage: children do not represent the times of event occurrences separately from the events

themselves, and, accordingly, they do not represent the systematic temporal relations between events (i.e., they do not have linear time). Nevertheless, these event representations provide the basis on which children begin to understand time itself (see also **Hudson et al**.).

3.3. Explaining the emergence of temporal reasoning

This still leaves questions about how the temporal reasoning system emerges. A number of the commentators raise the issue of the role of language and enculturation (**DeNigris and Brooks**; **Hohenberger**; **Hudson et al.**; **Tillman**); this issue is also raised by some commentators when querying why we believe animals do not engage in temporal reasoning (**Gentry and Buckner**; **Montemayor**). The idea that concepts of time are tightly linked to language and/or show substantial cross-cultural differences has a controversial history (Aveni, 1990; Gell, 1992). We have yet to come across convincing evidence that what we have described as the temporal reasoning system could not be assumed to be universal, and we suggested that culturally specific constructs may be overlaid on its basic functions. This does not mean, though, that we rule out the possibility that the temporal reasoning system emerges through social interaction. On the contrary, we have a great deal of sympathy for such a view. One approach here is to emphasize a specific role for language itself; as **DeNigris and Brooks** describe, Nelson (1996) has made a persuasive case for this. In our own writings, we have emphasized the role of communication rather than language per se in the acquisition of a linear notion of time. This is because we think that explaining a proper grasp of tense requires making an appeal to distinctive types of shared discourse that allow points in time to be the focus of shared attention with others. This argument is made in detail elsewhere (Hoerl & McCormack, 2010), but the general idea is that it is through joint discussion with other people about events at other times that children start to be able to

separate out locations in time from the events that occur at those locations and get to grips with the systematic relations between temporal locations and hence the time line itself.

Our intuition is that acquisition of temporal concepts is not possible in the absence of such shared discourse about other times, which makes clear, for example, that things in the past were different to how they are right now, or that events happened at past times even if they have left no present traces. Indeed, this is part of what motivates us to suspect that animals are not capable of temporal reasoning (in response to **Gentry and Buckner**; **Montemayor**; **Pan and Carruthers**).

R3.4. The self, narratives, and particular times

The picture of the development of the temporal reasoning system that we have just sketched also raises the interesting question of how it links to the developing conception of the self as temporally extended, not least because children's shared discourse about other times typically involves discussion of events in which the child was or will be a participant. We do not have space here to provide a treatment of this issue (see McCormack, 2015; McCormack & Hoerl, 2001), but in **Oyserman and Dawson's** commentary, they suggest that the temporal reasoning system is put to use in motivating action not just through a recognition that one is temporally extended, but through a (more sophisticated) recognition that one's future identity is not yet fully fixed. Specifically, they describe the way in which one's current actions are guided by how one envisages the identity of one's future self, using temporal reasoning to steer an appropriate path towards that preferred self.

This seems very plausible, and is in line with the idea that temporal reasoning is closely linked to possession of the narrative form (Campbell, 1996; McCormack, 2015; McCormack

& Hoerl, 2011). In the case of Oyserman and Dawson's description, the temporal reasoning system serves in a sophisticated way (most likely in late childhood/adolescence) to shape the narrative that one wants to have regarding how one's life unfolds. There may, though, be a more developmentally basic link between autobiographical narratives and the temporal reasoning system. Campbell (1996) argued that understanding and using autobiographical narratives (which children properly get to grips with much earlier, around 4 or 5) necessitates a grasp of the unique temporal locations of events – the fact that an event occurred at a specific point in one's life has genuine significance in the context of such narratives. It is this way of representing temporal locations, as unique and unrepeated on a time line, that is embodied in the temporal reasoning system.

Gentry and Buckner suggest that, contrary to our account, animals may, at least in principle, be capable of thinking about locations in time this way. They argue that in animals' lives there appear to be relatively few genuinely distinctive 'landmark' events. Thus, some animals might in fact be capable of temporal reasoning but their lives may not be populated with sufficiently unique events to construct a structured time line. While we agree with these commentators that animals have little use to make of the idea of unrepeated events, it is important to distinguish between an event being unique in the sense of being unusual (such as moving to a different enclosure) and being unique in the sense of being represented as happening at a non-recurring time point (an individual occurrence of a repeated event is unique in this sense). Even if an animal did experience a highly salient and rare event, it is difficult to see what use it would make of the fact that this event had occurred at a particular unique time point in its past. Given the nature of animals' lives, this fact does not seem to have any practical significance at all.

R3.5. Temporal reasoning: What is it for?

If animals have no use for the idea that, at a particular time point in its past a certain event took place, what makes it the case that the fact that particular events happened in the past can have significance for humans? **Mahr** links this question to the normative dimension of human interaction, arguing that the benefits of temporal reasoning do not lie solely in its role in preparing for the future; rather, "in the psychological and social domains, representing particular past events has benefits in which the particularity and pastness of events matter for their own sake." His suggestion is that the normative dimensions of social interactions (such as what one's current obligations are) typically hinge on thinking about particular past interactions that explain why the relevant social norm applies. For example, one is entitled to be annoyed if someone fails to keep their promise made yesterday to meet at 5pm today.

One way to think about whether he is right about this claim is to consider whether the temporal updating system could in principle be sufficient for operating according to the relevant social norms. **Mahr** emphasizes the 'particularity' of some instances of events, in the sense that they occur in a specific context. We note, though, that the contextual specificity of the implications of particular event occurrences is not what makes them beyond the scope of the more primitive updating system. Rather, Mahr's interesting claim is that appreciation of the pastness of events per se (which the updating system cannot provide) is necessary in order to fully engage with social norms. The case of promises is a good one to focus on, due to the sort of thought about time that promising seems to involve. If one properly understands what one is doing in making a promise, one is not just anticipating a future point in time at which one will have to deliver on the promise; rather one also simultaneously accepts that at a future point in time, others will be entitled to point back to the current moment in time if the promise is not delivered. Thus, we agree with Mahr that something like the temporal

reasoning system, and not just the updating system, must be involved in properly operating with these sorts of social obligations.

However, such social obligations are not the only example of a case where appreciating the pastness of events matters: this is also true of certain types of complex emotions. For example, the emotion of regret involves appreciating that there was a moment in time that has now passed at which future events could have unfolded differently but did not (Hoerl & McCormack, 2016). This brings us to the issue raised by **Beck and Rafetseder** about the connection between temporal reasoning and counterfactual thought. We agree with their suggestion that developmental research on counterfactual cognition would benefit from a consideration of the "temporal aspects of the demands being made on the child". In fact, we believe that there is an even tighter link between counterfactual and temporal cognition than their commentary explicitly considers. As we mention above, a key achievement associated with the emergence of the temporal reasoning system is the ability to think about locations in time separately from the events that occur in those locations (event-independent thought about time). Elsewhere, we have argued that engaging in counterfactual thought about events in the past (such as the type of counterfactual thought that underpins regret) involves exactly this sort of event-independent thought about time (McCormack, 2015; McCormack & Hoerl, 2008, 2017). Put simply, the idea is that it involves representing past times as locations in time at which a variety of different events could have unfolded. More than that, when one is engaging in genuinely counterfactual thought about a specific past event, one grasps the fact that at an earlier point in time, the event that is in the past now was once in the future before it unfolded as it did. This is what is at the heart of the idea that *things could have been different*. But if this is the right description of counterfactual thought, then it resembles the sort of temporal perspective-taking that we describe in the paper as being a feature of the

temporal reasoning system (i.e., that puts to work the fact that at a different point in time, things that are now in the past were once in the future).[3]

R3.6 The origins of temporal reasoning

We have sketched a developmental picture of the emergence of the temporal reasoning system that is distinctive insofar as we suggested that the temporal reasoning system does not simply operate with representations that are more explicit or enriched than those over which the updating system operates. In this sense, our account contrasts with some accounts of conceptual development in other domains. But if the reasoning system does not straightforwardly emerge from the updating system, what are its developmental origins? Influenced by Katherine Nelson's work, we suspect that the answer to this question will involve considering how children's emerging ability to represent events and event sequences enables them to engage in discourse with others about past and future events, which allows them to grasp the structure and nature of time itself. That is, we believe that the mature temporal concepts employed within the reasoning system are likely to have social origins. Moreover, it is within a broadly social context that the distinctive purposes for which people use the temporal reasoning system seem to be best illustrated, such as in generating and shaping autobiographical narratives, in evaluating past actions by considering alternative choices that one could have made, and in dealing with social norms and obligations. Research on the developmental origins of temporal concepts is sparse, meaning that these suggestions are speculative. We hope that our account at least provides some impetus for addressing what is currently a surprisingly large gap in the literature on cognitive development.

[3] If this analysis is correct, though, one issue that comes to the fore in assessing whether we are correct about animal limitations in temporal cognition is whether animals experience regret (Steiner & Redish, 2014; see Hoerl & McCormack, 2017, for discussion).

R4. Dual systems, representation and phenomenology

In this final section, we pick up on a set of broader theoretical issues emerging from across a number of commentaries. The first concerns our characterization of our account as a dual-systems account, and the precise grounds on which we draw a distinction between two systems. In subsection R4.1, we emphasise that our account distinguishes between two cognitive systems for dealing with how things unfold over time based on the *content* of the representations maintained, and we contrast this in particular with the idea that the two systems should be distinguished on the basis of the *representational format* of the representations the two systems operate with. We then turn to questions raised by several commentators about how our account relates to the phenomenology of experience in time, discussing first the special place the present has in experience (R4.2) and then the phenomenology of perceptual experiences of dynamic phenomena (R4.3). Finally, we discuss an issue that can be seen to underlie some of the commentaries discussing perceptual phenomenology, but also emerges from some of the other commentaries, which is whether there are ways of representing times that fall outside our dichotomy between temporal updating and temporal reasoning (R4.4).

R4.1 Questions of demarcation: Systems and representational formats

Some commentators make suggestions about how to extend our account of the temporal updating system and the temporal reasoning system that draw on existing ideas about underlying brain processes. **Nuyens and Griffiths**, in their discussion about the relation between our account and effects of emotion, suggest that the mechanisms underpinning timing utilized by the updating system map on to timing processes described in Scalar Expectancy Theory. **Ueda and Hanakawa** go further and make suggestions about the neural

mechanisms underpinning the two systems, and moreover how these may be impaired in certain clinical populations. We do not wish to make any specific claims about the neural substrates of the two systems, although we acknowledge that some researchers may find this frustrating. Moreover, while we describe our account as a dual systems account, we have sought to make it clear that we understand this label quite broadly. For instance, we are happy to agree with **Isham, Ziskin and Peterson** that our account could also be characterized as a dual process account.

Similarly, our account does not seek to distinguish temporal updating from temporal reasoning on the basis that one involves reasoning (and representations that can underpin the relevant reasoning) and the other one doesn't. Indeed, we agree with **Montemayor** that the representations maintained by the temporal updating system could be used in order to draw inferences (though not necessarily in the particular way he envisages). One basic such inference may be illustrated by an example given by **De Brigard and O'Neill**, who note that rodents can switch to taking a longer path in a maze upon learning that a shorter one that they used to take is closed. On our account, learning that the shorter path is closed causes the rat to update its model of the world so that that path is no longer represented as open, and the rat can then infer which of the remaining paths is now the shortest. We see this as an inference the rat can make on the basis of its present, now updated, representation of its environment alone. De Brigard and O'Neill, by contrast, speak of the animal as being "capable of drawing contrastive inferences between outdated and updated representations". If this is to mean that the animal still has a representation, now outdated, of the situation of the previously shortest path being open, we see no reason for making such an assumption.

Crucially, our account draws a distinction between two cognitive systems for dealing with how things unfold over time based on the *content* of the representations maintained by those systems. Several commentators, by contrast, suggest variations or alternatives to our account that do involve specific claims about representational *formats*. Thus, **Kelly, Prabhakar and Khemlani** suggest that our distinction between temporal updating and temporal reasoning should be seen to map on to two different kinds of iconic representations, which they call 'perceptual models' and 'event models', respectively. Similarly, **Keven** suggests that we should see the temporal updating system as operating with "snapshot like" representations that he calls event memories, which are based on automatic perceptual processes.

As we are not committed to any specific claim about the representational format of the representations maintained by the temporal updating system, these other accounts may in principle be compatible with ours, but there are also potential pitfalls if questions of content and questions of format are not clearly distinguished from one another. For instance, as **Kelly et al.** themselves point out, mature thinking about time does not just consist in drawing inferences on the basis of recalled events. It also involves more mundane abilities such as the capacity to simply recount a past sequence of events. Thus, it may involve mental states involving quite different representational formats, e.g., episodic memories as well as mere beliefs about the time when events happened.

Keven makes a number of claims that we, too, would endorse. For instance, he claims that the birds in Clayton and Dickinson's (1998) study need not be able to remember "the actual experience of caching the food items", but might instead "remember in the same way I can remember *where my keys are* without remembering the actual experience of putting them there" (our emphasis). In line with Keven's use of the present tense in drawing this analogy,

we similarly believe that the behavior in Clayton and Dickinson's study can be explained on the basis of a representation birds entertain that concerns their current environment only, though of course one that originates in past experience.

Is there anything more than just a terminological difference between our view and **Keven's**? The main reason he seems to prefer to frame his view in terms of the idea of 'event memories' is that he thinks that information from past experience is retained in a quasi-perceptual format. As he quite rightly says, our view is not committed to any specific format in which information is retained by the temporal updating system. We regard this as a further, empirical, issue. But even supposing that the representations that make up the model maintained by the updating system are indeed quasi-perceptual, does it follow that the term 'event memories' is an apt one? We believe that questions can be asked about both words in that term. First, with regard to the word 'event', consider Keven's description of those representations as 'snapshots'. Suppose that, when I put down my keys, I actually take a photo of them in the location where I put them down, which I then carry around to remind me of their whereabouts. If the only use I subsequently make of this snapshot is as a representation of where my keys are, it seems far from obvious why it should be counted as serving the role of representing *an event*, rather than simply a state of affairs.

Second, turning to the word 'memory', **Keven** exhorts us to "call a memory a memory", and characterizes our description of the model of the world maintained in mere temporal updating as one made up of representations that are at least "memory-like". We believe that this terminology, too, is unhelpful, as the term 'memory' is first and foremost a name of a capacity, rather than singling out a particular type of representation. Indeed, on our view, when a creature uses the temporal updating system, there is nothing merely 'memory-like'

about the capacities that allow it to do so. It clearly uses its memory in so far as it acquires, encodes and stores information. Thus, on that score, we are more than happy to 'call memory memory'. The reasons for describing the creature as having *a* memory are much less obvious to us.

R4.2. The phenomenology of the present

An existing claim regarding dual systems is that they give rise to what Sloman (1996) calls "simultaneous contradictory beliefs". In section 4 of our target paper, we draw on this idea to connect our dual systems approach with a specific aspect of human adults' everyday picture of time, viz. the fact that it accords a special status to the present moment in time, whilst also granting that every time has the property of being present when it is that time. Commentators raise a number of questions about this, including: (i) How exactly should this aspect of people's everyday picture of time be characterized (and does it indeed involve a contradiction)? (ii) Even if it does involve a contradiction, is this good evidence for the existence of dual systems? (iii) How are our claims related to questions about the phenomenology of people's experience of time, which is often held to underpin their everyday picture of time?

The contradiction in people's everyday picture of time that we suggest might be explained by the existence of the two systems arises from the idea that time flows or passes. But philosophers have given two somewhat different characterizations of this idea. As pointed out by **Miller, Holcombe and Latham** as well as **Prosser**, there are ways of spelling out the idea that time flows that do not obviously involve a contradiction. According to *presentism*, only present things exist, and what is called the 'flow of time' is simply those things changing. On such a view time itself is an abstraction, rather than a dimension of reality, and moments in

time are considered logical constructions (Prior, 1972, 1996; see Hoerl, 2015, for discussion). The presentist view of the flow of time does not obviously involve a contradiction. A contradiction only arises on what we might call a *moments-based* understanding of the flow of time. Contrary to presentism, this understanding does operate with a view of time as a dimension made up of different moments of time. According to it, the flow of time is a change over time in which of these moments is present (or, alternatively, a change that consists in events occupying future, present and past moments in time in succession).

Does human's everyday conception of the flow of time correspond more closely to the presentist understanding or to the moments-based understanding of the flow of time? In our target paper, we assumed the latter. In assuming this, we are sympathetic to **Prosser's** suggestion that people's everyday picture of time is in fact an amalgam of presentism, on the one hand, and *eternalism*, on the other – where the latter is the view that reality is extended in time and that all times are equally real. As Prosser suggests, this might be because the temporal updating system embodies something like a presentist view of reality, whereas the temporal reasoning system embodies an eternalist metaphysics.

Even if this suggestion is along the right lines, it will only explain one specific ingredient of people's everyday picture of time. It is focussed specifically on the privileged status that our everyday picture of time accords the present moment in time, as one ingredient in the idea that time flows or passes. As should be clear from what we have just said, and as also pointed out by **Callender**, **Prosser** and **Kenward and Pelling**, another aspect of the idea of time as flowing or passing is that of a constant change of a certain (perhaps hard to define) type – in Callender's words, the idea that "future events draw nearer and past events recede". We are not claiming to have given an account of the source of this further idea (for one suggested

account, see Hoerl, 2014). Yet, even if it is thus limited in its scope, our proposed explanation might help address an important question **Miller et al.** press in their commentary. As they point out, what really needs to be explained is the sense in which it still *seems* to people that the present has an objectively privileged status, even once they have come to the conclusion that it does not. This seeming, Miller et al. say, is different from belief – it is a matter of phenomenology. But how should we think of the type of phenomenology at issue here?

There are clearly aspects of *perceptual phenomenology* linked to time, such as the experience of motion (**Hayman and Huebner**) or sound (**Kenward and Pilling**). We will discuss these in more detail in the next section. For present purposes, though, it is enough to note that the present moment in time seeming privileged is not a matter of perceptual phenomenology. It is not something that simply figures in perceptual experience alongside, say, motion or sound.

In drawing on Sloman (1996) in our target paper, we aimed to bring out that there is also another, distinct, kind of phenomenology – one which dual systems can give rise to. Whilst Sloman frames the point in terms of the idea that dual systems can generate "simultaneous contradictory beliefs", we pointed out that he actually uses the term 'belief' in this context in a non-standard way. As **Melnikoff and Bargh** point out in their commentary, contradictory beliefs, as such, might also arise in a single system. What we take Sloman to be arguing, rather, is that the existence of dual systems can explain in particular how situations can arise in which a belief-like state can persist – a seeming as though something is true – even though, at the same time, the content of that belief-like state has been dismissed as untrue by the thinker. This is the result of the more basic system still continuing to deliver its verdict, even though this is not endorsed by the more sophisticated system. This would give us a way of understanding how its seeming that the present moment in time has a privileged status can be

an aspect of people's phenomenology, even though it is clearly not a matter of perceptual phenomenology.

R4.3. Temporal updating and perceptual experience

Amongst the commentaries there are also some that do focus specifically on perceptual phenomenology – a topic that we did not touch on. A common theme is that our account takes insufficient account of the dynamic character of perceptual experiences. For example, **Kenward and Pilling** mention the ability to perceive sounds and argue that this ability is evidence that perceptual experience "encompasses happening events rather than just a millisecond snapshot" (compare also **Roselli**). Similarly, **Hayman and Huebner** write: "If a moving object were experienced only as present, each momentary state would feature a static object, with nothing to bind these states together as an experience of ongoing motion". And **Elliot** uses an example from a study with fighting fish conducted by Brecher (1932) to argue that "events separated by time might, given short intervals between their presentations, be combined to form meaningful […] experiential content."

How exactly do these claims bear on our account? On our account, the perception of dynamic phenomena plausibly causes a rapid updating of the model of the world maintained by the temporal updating system. Thus, for instance, as a creature watches another creature move, new information about the location of the other creature constantly overrides previous information about its location. This is something our account can clearly accommodate. But the commentators seem to think that our account nevertheless seems to leave out something important about the phenomenology of perceptual experience.

We take it that the authors in question think that we face a dilemma. Either we deny that perceptual experience can encompass more than what **Kenward and Pilling** call a "millisecond snapshot". But that would be to deny the obvious phenomenological truth that we can perceive such things as sounds and movements. Or we allow that perceptual experience can encompass things that happen over an extended period of time. But then perceptual experience might be sufficient, even in animals, to ground representations of different things happening at different times. As Kenward and Pilling put it, we might be right in claiming that animals are incapable of mental time travel and therefore have no way of representing change over longer periods of time. Yet, this does not rule out that they are able to do so within the section of the world they are perceptually conscious of. Thus, contrary to what we say, animals' model of the world does contain a temporal dimension, it is just that "their representational time-line is very short".

We have no intention of trying to attack the first horn of this dilemma. Nothing of what we say is meant to imply any claim to the effect that perceptual experience only encompasses a "millisecond snapshot" (cf. Hoerl, 2009, 2013). However, things are different with the second horn of the dilemma. Does the fact that perceptual experience encompasses things that in fact happen over a period of time imply that it gives rise to representations of different things happening at different times? We are not convinced it does.

Take, for instance, **Elliot's** own example of Brecher's fighting fish. Brecher presented the fish with a tachistoscopic display in which an image of a conspecific repeatedly flashed up briefly in succession. What Brecher found was that at a frequency of around 30 Hz the fish would attempt to attack the image, even though they would not do so at lower frequencies. Arguably, what this means is that it is only at that or higher frequencies that the fish

recognized a conspecific in the display, which in turn means that their visual system must be integrating information over a period of time in which what are in fact several different images are flashed up in succession. Yet this does not provide a good reason for thinking that the resulting representation is itself one of succession, or some other form of representation in which times or temporal relations figure. Rather, the relevant representation formed is just that of another fish being present.

We can extract from this a general moral that can also be extended to some other examples mentioned in the commentaries. Even if it is necessary for an organism to sample information over time to detect a certain feature of its environment, this does not imply that the resulting representation of that feature is one in which temporal properties or relations figure (for a somewhat related argument, see also Prosser, 2016, ch. 6). This might even be applied to **Kenward and Pilling's** example of the experience of sound, for instance, which they describe as an "inherently temporal phenomenon". It is of course true that sounds themselves are temporal phenomena in that they have an onset, and offset and a duration. But it is ultimately not clear that this marks any deep difference between the way sounds are related to time and the way, e.g., colors are. (On this issue, see also Cohen, 2010.) Thus, like the latter, they might simply be represented as qualities instantiated at a time.

The same general moral also carries over to other considerations not involving perceptual phenomenology. For instance, **Pan and Carruthers** argue that foraging animals must have a representation of time because they can compare the rate of reward at the current location with the average rate experienced at other locations and take this into account in calculating whether they should stay where they are or make the effort to move elsewhere. Clearly, as Pan and Carruthers say, "a rate is a measure of a quantity per unit of time". But does that

mean that sensitivity to a rate requires an ability to entertain representations in which other times figure?

To adapt an analogy used by Prosser (2016), think of the dial on a car speedometer. This acts as a representation of a rate, a rate related to what locations the car occupies at what times, but it does not act as a representation of a succession of events happening at different times. Arguably, people also use it in ways that don't need to involve any reasoning about other times, such as when they simply look at the speedometer, notice that they are breaking the speed limit, and take their foot off the gas. In a similar way, we want to suggest that animals may be sensitive to the rate of reward at different locations, but in a way that does not involve any reasoning about different moments in time at all.

R4.4. Representing space and representing time

In the preceding section, we mentioned **Kenward and Pilling's** suggestion that animals' model of the world does contain a temporal dimension, but that "their representational time-line is very short". The general idea is that there might be ways of representing time that escape our dichotomy between temporal updating and temporal reasoning. A similar idea can also be seen to be at work in some of the commentaries that raise issues regarding the relation between temporal and spatial representation.

On our account space and time are given a very disparate treatment by the temporal updating system, which can represent spatial relationships, but ignores altogether that reality also has a temporal dimension. **Callender** questions this, and speaks of "how tightly linked [time and space] are in physics, biology and psychology". There are clearly a number of ways in which temporal and spatial processing are conceptualized similarly in biology and psychology.

Callender mentions the idea that there are ‘time cells’ as well as ‘space cells’ in the brain. Similarly, **De Corte and Wasserman** draw attention to work that has been trying to establish the existence of ‘temporal maps’ in animals. And **Gentry and Buckner** ask why mechanisms for spatial representation that have been demonstrated in animals “could not be bootstrapped to represent an additional temporal dimension as well”. To what extent might findings in these areas challenge our view? We will focus in particular on the research on ‘temporal maps’.

In a typical backward conditioning study of ‘temporal maps’, the animal completes two training phases. In the first of these it is repeatedly presented with two different cues (Cue A and Cue B), separated by a time interval. In the second training phase, the animal is then presented with a reward, followed soon after by the second cue from the first training phase (Cue B), where the interval between the reward and Cue B is shorter than the interval between Cue A and Cue B in the first training phase. What the research seems to show is that animals integrate the associations between Cue A and Cue B and between the reward and Cue B, including the intervals between them, such that, when given Cue A, they now expect the reward, despite never having experienced Cue A followed by the reward.

De Corte and Wasserman, who mention this research in their commentary, seem to primarily want to question our characterization of the signature limit that we claim the temporal updating system is subject to. In describing the behavioral implications of a creature operating only with a temporal updating system, we identify as one signature limit of such a system that it cannot deal with situations in which information about changes is received in an order that differs from the one in which these changes happened, and we provide some empirical evidence that children do indeed have problems in situations of this type. We also

suggest, though, that even a creature capable only of temporal updating might nevertheless be able to engage in a form of sequential learning, by coming to acquire a routine for updating its model of the world in a particular order.

The type of sequential learning described by **De Corte and Wasserman** is indeed more complex than anything we describe in the target paper. But does it contradict our characterization of the signature processing limit that the temporal updating system is subject to? Note that, whilst the animal at first learns about two sequences in isolation from each other, it is still the case that it learns about each of those sequences in the order in which that sequence happens. This is different from the kind of case we had in mind when describing the limitations the temporal updating is subject to. These have to do specifically with the fact that the overall outcome of a sequence of changes can sometimes depend on the order in which those changes happen. Now consider a case in which the animal does not directly witness the each of the relevant changes in turn, but receives information about their occurrence in an order that does not correspond to the order in which they actually occurred. The only way in which it can arrive at a correct representation of the overall outcome is if it can reason about temporal order to infer that overall outcome.

As we read them, **De Corte and Wasserman** do not necessarily want to deny that the abilities that have led researchers to credit animals with ‘temporal maps’ can be explained without ascribing to them a capacity for temporal reasoning in our sense. The paradigms used in this research are Pavlovian ones, and they are typically interpreted in associationist terms – the idea being that, when an association between two events is being formed, the temporal relationship between them is also encoded as part of the association between them. This can then explain why the animal forms certain expectations at certain times.

However, the use of the term ‘temporal map’ might nevertheless prompt a more general query for our approach. We are making two crucial claims about animals: The first is that they cannot engage in temporal reasoning as we describe it, including mental time travel. The second is that the representations they can entertain are ones from which the dimension of time is simply absent. As can be seen in several places in this response, across a number of different commentaries commentators seem prepared to agree with a version of the first claim, but disagree with the second one (see, e.g., **Hohenberger**; **Osvath and Kabadayi**; **Pan and Carruthers**; **Roselli**). Moreover, the idea that there might be ways of representing things in time that escape our dichotomy of temporal reasoning vs. temporal updating is also a theme in some comments not concerned with animal cognition (see, e.g., **Viera and Margolis**; **Elliot; Hayman and Huebner; Kenward and Pilling**). We want to end with some remarks on this general issue, addressing in particular the issue of whether it is correct to think of animals’ ‘temporal maps’ as closely analogous to the ‘spatial maps’ commonly ascribed to them in the literature on animal navigation.

A picture of animal cognition on which animals have basic capacities for temporal representation that closely parallel their capacities for spatial representation has been developed by John Campbell (1994, 2006). On Campbell’s account, animals operate with both spatial and temporal frameworks for orienting themselves in their environment, but ones that are crucially different from the kinds of representations of space and time that adult humans can frame. This is because the animal gives causal significance to those frameworks only through its own practical engagement with its environment. In the case of the spatial frameworks used by animals, the consequence of this is that, whilst the animal may be able to represent the direct route from its current location to any other place represented on its

cognitive map, "it cannot represent the spatial relations between any two arbitrary places in its environment" (Campbell, 1993, p. 87), irrespective of whether it is located at them. The animal's orientation in time is also limited, according to Campbell, in that it is only an orientation with respect to phase, rather than with respect to particular times. Thus, while an animal with a circadian timer might be said to be able to represent an event as happening at noon, on Campbell's account, "it has no use for the distinction between noon on one day and noon on another" (Campbell, 1996, p. 118). On Campbell's picture, animals do have temporal frameworks, but these are more primitive and do not allow for times to be represented as unique unrepeating locations on a time line.

We are happy to go along with Campbell's characterization of the spatial representational abilities of animals. Given this, and given that it is possible to construct a structurally parallel model of their temporal representational abilities, which constitutes something like a halfway house between what we describe as temporal updating and temporal reasoning, respectively, what reasons could there be for being sceptical that such a halfway house exists? We note that, in the literature mentioned by **Gentry and Buckner** on ways in which humans recruit spatial representations to think about time, a common claim is also that the reverse does not hold – where this is often combined with the claim that time is 'abstract' in a way that space is not (Boroditsky, 2000; Casasanto & Boroditsky, 2008). We think that this might point to an important distinction in how talk about 'temporal maps' and 'spatial maps' in animals should be understood.

In the case of spatial maps, we have reason to think that possession of such a map by an animal involves the ability to represent places in its environment, because the animal uses the map for navigation, i.e. for actively moving from one such place to another. This provides for

a concrete sense in which the animal can give significance to the fact that there are places other than the one at which it is located. Yet, there is no analogue to this in the case of time. Even if we ascribe to the animal a 'temporal map' such as the ones postulated in the literature on backward conditioning, there is nothing the animal itself can or needs to do to exploit the information encoded in the map other than form the right representations (i.e., first of Cue A, then of the reward, and then of Cue B) at the appropriate moments. But if all an animal can do is generate the right representations at the right moments, then it is not clear that it is useful to describe it as possessing a 'map' of time at all, in the way the term 'map' is usually understood, i.e., as something that gives the topography of a domain. As we have argued, it could generate the appropriate representation at any given moment without representing other moments in time at all. Thinking of the issue in this way gives the term 'mental time travel' its proper import: what is significant about the term is the idea of mentally navigating the temporal domain, i.e., making use of a mental map of the structure of relations between points in time. What we are seeking to do is re-frame the debate over whether animals can engage in mental time travel as a debate about whether they possess maps of time in this sense.

R5. Conclusion

We are very grateful to all commentators for taking the time to write their commentaries, which have prompted us to both get and be clearer about a number of aspects of the account we are proposing. Given the richness and variety of the commentaries, it has not been possible to provide a detailed discussion of all the points that were raised by commentators. Instead, one thing we have focused on in particular in response to them is bringing out ways in which time appears special as a domain of cognition. In all three main sections of this response, we discussed commentators' suggestions to bring ideas already familiar from other

domains to bear on our account: for instance, by allowing that animals might be able to implicitly represent the dimension of time, but without representing it as such; by subsuming our account under a broader type of developmental distinction; or by assimilating temporal representation to spatial representation. In each case, while it is certainly helpful to consider such parallels, we suspect that these suggestions do not account sufficiently for the difference between time and other cognitive domains. Specifically, we have argued that whilst the temporal updating system can explain a range of behaviors that are adapted to how things unfold over time, this does not mean that it provides some primitive way of representing the domain of times. Conversely, the representations of time that the temporal reasoning system operates with cannot be understood as some sort of more explicit or enriched versions of representations present already in the temporal updating system. Time thus raises its own distinctive issues when it comes to the question as to what it takes to represent it, and how the capacity to do so emerges phylogenetically and ontogenetically.

References

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

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

Boroditsky, L. (2000). Metaphoric structuring: understanding time through spatial metaphors. *Cognition, 75*(1), 1-28. doi:10.1016/S0010-0277(99)00073-6

Brecher, G. A. (1932). Die Entstehung und biologische Bedeutung der subjektiven Zeiteinheit, — des Momentes. *Zeitschrift für Vergleichende Physiologie, 18*(1), 204-243. doi:10.1007/bf00338160

Campbell, J. (1993). The role of physical objects in spatial thinking. In N. Eilan, R. McCarthy, & B. Brewer (Eds.), *Spatial representation* (pp. 65-95). Oxford, England: Blackwell.

Campbell, J. (1994). *Past, space and self*. Cambridge, MA: MIT Press.

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

Campbell, J. (2006). Ordinary thinking about time. In M. Stöltzner & F. Stadler (Eds.), *Time and History: Proceedings of the 28. International Ludwig Wittgenstein Symposium, Kirchberg Am Wechsel, Austria 2005* (pp. 1-12). Heusenstamm, Germany: ontos verlag. doi:10.1515/9783110333213.1

Casasanto, D., & Boroditsky, L. (2008). Time in the mind: using space to think about time. *Cognition, 106*(2), 579-593. doi:10.1016/j.cognition.2007.03.004

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

Cohen, J. (2010). Sounds and temporality. *Oxford Studies in Metaphysics, 5*, 303-320.

Droit-Volet, S., & Rattat, A.-C. (1999). Are time and action dissociated in young children's time estimation? *Cognitive Development, 14*(4), 573-595. doi:10.1016/S0885-2014(99)00020-9

Friedman, W. (1990). *About time: Inventing the fourth dimension*. Cambridge, MA: MIT Press.

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

Hoerl, C. (2009). Time and tense in perceptual experience. *Philosophers Imprint, 9*(12), 1-18. URL: <www.philosophersimprint.org/009012/>

Hoerl, C. (2013). 'A succession of feelings, in and of itself, is not a feeling of succession'. *Mind, 122*(486), 373-417. doi:10.1093/mind/fzt070

Hoerl, C. (2014). Time and the domain of consciousness. *Annals of the New York Academy of Sciences, 1326*(1), 90-96. doi:10.1111/nyas.12471

Hoerl, C. (2015). Tense and the psychology of relief. *Topoi, 34*(1), 217-231. doi:10.1007/s11245-013-9226-3

Hoerl, C., & McCormack, T. (2010). 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: Oxford University Press. doi:10.1093/acprof:oso/9780199245635.001.0001

Hoerl, C., & McCormack, T. (2016). Making decisions about the future: Regret and the cognitive function of episodic memory. In K. Michaelian, S. Klein, & K. Szpunar (Eds.), *Seeing the future: Theoretical perspectives on future-oriented mental time travel* (pp. 241-266). Oxford, England: Oxford University Press. doi:10.1093/acprof:oso/9780190241537.001.0001

Hoerl, C., & McCormack, T. (2017). Animal minds in time: the case of episodic memory. In K. Andrews & J. Beck (Eds.), *The Routledge Handbook of Philosophy of Animal Minds* (pp. 56-64). Oxford, England: Routledge.

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

Karmiloff-Smith, A. (1994). Précis of Beyond modularity: A developmental perspective on cognitive science. *Behavioral and Brain Sciences, 17*(4), 693. doi:10.1017/S0140525X00036621

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

McCormack, T. (2015). The development of temporal cognition. In R. Lerner (Ed.), *Handbook of Child Psychology and Developmental Science, 7th Edition, Volume 2: Cognitive Processes* (pp. 624-670). Hoboken, N.J.: Wiley-Blackwell. doi:10.1002/9781118963418.childpsy215

McCormack, T., & Hoerl, C. (1999). Memory and temporal perspective: The role of temporal frameworks in memory development. *Developmental Review, 19*(1), 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 & K. Lemmon (Eds.), *The self in time: Developmental perspectives* (pp. 203-227). Mahwah, NJ: Lawrence Erlbaum.

McCormack, T., & Hoerl, C. (2008). Temporal Decentering and the Development of Temporal Concepts. *Language Learning, 58*, 89-113. doi:10.1111/j.1467-9922.2008.00464.x

McCormack, T., & Hoerl, C. (2017). The development of temporal concepts: Learning to locate events in time. *Timing & Time Perception Reviews, 2*(3), 1-18. doi:10.1163/22134468-00002094

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

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

Prior, A. N. (1972). The notion of the present. In J. T. Fraser, F. C. Haber, & G. H. Müller (Eds.), *The Study of Time: Proceedings of the First Conference of the International Society for the Study of Time Oberwolfach (Black Forest) — West Germany* (pp. 320-323). Berlin, Germany: Springer. doi:10.1007/978-3-642-65387-2_22

Prior, A. N. (1996). Two essays on temporal realism. In B. J. Copeland (Ed.), *Logic and Reality: Essays on the Legacy of Arthur Prior* (pp. 43-51). Oxford, England: Oxford University Press.

Prosser, S. (2016). *Experiencing time*. Oxford, England: Oxford University Press.

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, P. (1982). Bennett's beliefs. *Philosophical Studies, 41*(3), 431-442. doi:10.1007/BF00353891

Steiner, A. P., & Redish, A. D. (2014). Behavioral and neurophysiological correlates of regret in rat decision-making on a neuroeconomic task. *Nature Neuroscience, 17*(7), 995-1002. doi:10.1038/nn.3740

van Schaik, C. P., Damerius, L., & Isler, K. (2013). Wild orangutan males plan and communicate their travel direction one day in advance. *PloS One, 8*(9), e74896. doi:10.1371/journal.pone.0074896

Vanderveldt, A., Oliveira, L., & Green, L. (2016). Delay discounting: Pigeon, rat, human – does it matter? *J Exp Psychol Anim Learn Cogn, 42*(2), 141-162. doi:10.1037/xan0000097