Brit. J. Phil. Sci. **71** (2020), 443–469

The Sense of Time

Gerardo Viera

ABSTRACT

It's often claimed in the philosophical and scientific literature on temporal representation that there is no such thing as a genuine sensory system for time. In this article, I argue for the opposite—many animals, including all mammals, possess a genuine sensory system for time based in the circadian system. In arguing for this conclusion, I develop a semantics and meta-semantics for explaining how the endogenous rhythms of the circadian system provide organisms with a direct information link to the temporal structure of their environment. In doing so, I highlight the role of sensory systems in an information processing architecture.

---

1 Introduction

The ability to keep track of time is arguably one of the most widespread psychological capacities in the entire animal kingdom. Every animal that navigates its environment must have some way of coordinating its behaviours with the temporal structure of the events around it. Despite the fact that the ability to keep track of time is so ubiquitous in the animal kingdom, it is a hotly contested question as to how animals keep track of time. Nevertheless, it is widely believed that animals do not possess a genuine sense of time. Instead, any capacity to keep track of time is explained by appealing to the information gathering capacities of the other sensory systems, like vision, audition, touch,

doi:10.1093/bjps/axy019

and so on (Buonomano and Karmarkar [2002]; Aristotle [2004]), or through the operation of post-sensory mechanisms (Gallistel [1996]; Zakay and Block [1997]; Matthen [2014]; Phillips [2008]).

I will argue that a significant portion of how mammals come to keep track of time, and thereby come to coordinate their behaviours with the temporal structure of the world around them, is explained through the operation of a genuine sensory system for time that in many ways parallels the operations of the paradigmatic sensory systems like vision and olfaction. Section 1 begins with a clarification of the notion of a sensory system that I will be discussing. Section 2 describes two arguments found in the literature that aim to show that a wide range of animals fail to have a genuine sense of time. I show that both arguments fail to establish their intended conclusions. As a result, there are no *a priori* or theoretical reasons for why there cannot be a sense of time. In Section 3, I turn to empirical work on the circadian systems of mammals to argue that here we have a genuine sense of time. Section 4 provides an account of the informational content of clocks, both biological and cultural. In Section 5, I provide an information-theoretic account of how clocks represent time. In conclusion, I argue that all mammals possess a genuine sense of time.[1]

2 Sensory Systems and Experience

I will argue that the circadian system constitutes a genuine sensory system for time. However, I will not be arguing that the circadian system contributes to our subjective awareness of time. In fact, the circadian system makes zero direct contribution to our conscious experience.

This might strike some as odd, even contradictory, since according to many philosophers the senses necessarily contribute to our conscious experience. One common reason for thinking this is that the senses are often thought to be individuated according to their introspectively available properties. Some attempt to individuate the senses according to their introspectively available content—for example, vision represents colours while audition represents sounds (Grice [1962]; Dretske [1995]; Aristotle [2004]). Others attempt to individuate the senses according to their distinct phenomenal characters—for example, the phenomenal characters of seeing a square and of feeling a square are simply different (Lopes [2000]). In either case, whatever the sensory systems are, they must contribute to our conscious experience since it is only through making this contribution that the senses can be of a particular type. Therefore, if some system fails to make any contribution to our experience of the world, then it cannot be a sensory system.

---

[1] The conclusion of the article could be extended to cover non-mammals as well; however, this would require a discussion of non-mammalian circadian systems that for reasons of length cannot be done here.

However, as Keeley ([2002]) argued, attaching the notion of sensory systems to phenomenal experience undermines the explanatory role that the senses have in much of cognitive science. Scientists routinely attribute novel sensory systems to animals to explain how those animals coordinate their behaviours with aspects of their environments. Take for instance the case of the elasmobranch fish (sharks, skates, and rays) and their electroreceptive sense (Kalmijn [1982], [2000]; Collin and Whitehead [2004]). Elasmobranch fish often attack creatures hidden under sand. Researchers initially thought there must be some cue picked up by one of the standard sensory systems that allowed these fish to find their prey. However, through systematic experimentation it was discovered that the informational resources of the standard senses could not explain this feeding behaviour. It was then discovered that the elasmobranch fish find their prey by detecting the electromagnetic fields produced by the animals. Researchers attributed to these fish a distinct electroreceptive sense. This attribution was further vindicated by uncovering the biological mechanisms by which this information was gathered. Importantly, in the attribution of this sense it's not merely the case that we do not know what phenomenal properties accompany the electroreceptive sense. The point is that we have no need to appeal to phenomenal properties whatsoever.

The attribution of a novel sensory system arises when the informational resources provided by the existing sensory systems are not up to the task of explaining how organisms coordinate their behaviours with the environment. Even our folk notion of a sensory system operates in this way. The standard plot device of a 'sixth sense' is used to explain how someone has epistemic access to some aspect of the world that they could not have epistemic access to via the standard sensory systems. In both the folk and scientific cases, what makes something a sensory system is that it is a distinct avenue by which information enters the psychological economy of the organism.

How then do we individuate the sensory system if not through introspection? Well, in recent years, various philosophers have argued that there isn't a single answer to this question. Once we realize the explanatory role that sensory systems play in our various theories of the mind, we find that no single set of criteria adequately provides us with the sensory taxonomies required by our theories (Macpherson [2011]; Fulkerson [2014]; Matthen [2015]). In some cases, our explanatory needs compel us to individuate sensory system by using certain criteria (either by their content, biological implementation, evolutionary history, and so on), while in other explanatory contexts we use other criteria. Importantly, no single set of criteria enjoys the distinction of being THE individuating criteria. As a result, we needn't insist that the senses must contribute to experience in order for them to be properly individuated. In fact, cases like that of the elasmobranch fish push us to say that individuation has no need for introspectively available properties.

It is in this context that I will argue that the circadian system constitutes a genuine sense of time. The information gathering capacities of the other sensory systems are not up to the task of explaining how animals coordinate many of their behaviours with the time of day. Furthermore, as we'll see, the biological evidence legitimizes the attribution of this sense by showing us the mechanisms by which organisms gather this information. This is all a matter of empirical discovery and not something learned through introspection. As we consider the arguments against the existence of a sense of time in the subsequent sections, this idea of sensory systems as information gathering mechanisms will be fleshed out.

3 Against the Sense of Time

While it's fairly common to find philosophers and cognitive scientists claiming that animals, including humans, do not possess a sense of time, it is far less common to find these claims backed up with anything like an explicit argument. However, if we look close enough we can find two types of arguments for the conclusion. In this section, I will lay out these arguments and show why they fail to establish their intended conclusions.

3.1 The non-causality argument

A defining characteristic of the senses, which distinguishes them from other aspects of the overall cognitive architecture, is that sensory systems provide animals with a direct information link to the ongoing changes in the environment (the directness of the sensory systems will be described more below). When we consider the classic Aristotelean senses, we find that this information link with the environment is established via the causal influence exerted by the relevant aspects of the environment on the sensory systems. For example, the visual system gathers information about the world through photons impacting the retina. The auditory system gathers information about the world through pressure exerted on the ear drum. Even for non-Aristotelean sensory systems, like proprioception and electroreception, similar causal interactions are at work. These sorts of causal influences from the world to the sensory systems appear all over the place.

It is this causal connection between the world and the sensory systems that forms the basis of the non-causality argument against the existence of a sense of time. As Matthen ([2015], p. 573) puts it:

> Should [the systems responsible for the representation of time] be regarded as transducers for a sense of time? That is, do periods of time cause them to emit a pulse that carries information about these periods of time? Both sides of the question can be argued. A negative answer might be a reason to exclude the sense of time.

Similarly, the psychologist Boroditsky ([2011], p. 333, emphasis added) voices the same concern when she says:

> All of our experience of the world is physical, accomplished through sensory perception and motor action. And yet our internal mental lives go far beyond those things observable through physical experience: we invent sophisticated notions of number and time [. . .] So how is it possible that physical organisms who collect photons through their eyes, respond to physical pressure in their ears, and bend their knees and flex their toes in just the right amount to defy gravity are able to invent and reason about the unperceivable and abstract?

The worry that Matthen ([2015]) and Boroditsky (Boroditsky [2011]; Boroditsky and Prinz [2008]), are pointing to are pointing to is the same: Time is what many call an 'abstract' feature of our world.[2] Time is not something that we can readily point to. Time isn't something that we can clearly manipulate. Some philosophers (notably, Lewis [1973]; Newton-Smith [1980]; Maudlin [2002]) have even gone so far as arguing that time lacks any causal powers whatsoever.[3] If sensory systems gather information about things in the environment through those things causally influencing the sensory systems (call this claim 'the causality constraint'), then, given the abstract nature of time, there could not be a sense of time.

Since the purpose of this article is to understand something about how animals come to keep track of time, and not to understand the fundamental metaphysics of time, we can take the claim that time is causally impotent as a fair assumption. Instead, I will argue that the causality constraint should be rejected.

The causality constraint seems to be most readily supported through an inference to the best explanation. Unlike cognitive systems, sensory systems directly gather information about the ongoing changes in the environment in that their information gathering role is not necessarily mediated by any other information bearing psychological processes or mechanisms. This information can then be used by other downstream systems for further processing. For instance, the accurate deployment of concepts in the thought 'there is a bear on the trail' carries information about the world, but our successful deployment of these concepts relies on the information about the world contained in other psychological systems. The senses, however, carry information about the state of the world, but they do not require for their accurate deployment other psychological systems that already carry information about the world.

Since information enters the cognitive economy somehow, how else might sensory systems carry information about the ongoing changes in the world if not through a causal connection? To deny this would seem to commit us to

---

[2] For similar arguments, see (Harris *et al.* [2010]; Coull [2011]).

[3] See (Benovsky [2012]) for an counterargument.

positing some mysterious connection between sensory systems and the external world. Therefore, we should expect the causality constraint to be true.

However, as Dretske ([1981]) emphasized, for there to be an informational link between a representational system and some state of the world (or in Dretske's terms, between a signal and its source) a causal connection isn't required. All that is required for an information link is that the following conditions be satisfied: 'A [system] *r* carries information that *s* is *F* if and only if the conditional probability of *s*'s being *F* given [the state of] *r* (and *k* [the background channel conditions]), is 1 (but given *k* alone, less than 1)' ([1981], p. 51).

Often what guarantees this relation is a causal connection. However, if the appropriate conditional relationship can be established in the absence of a causal connection between *r* and *s*, then we would still have the appropriate information link. First and foremost, it's the conditional probabilities that matter, not the causal connection.[4]

To show how one system could carry information about another while violating the causality constraint Dretske ([1981]) gave the following example: Imagine two televisions, TV_1 and TV_2, that are isolated from one another but are connected to the same signal so both TVs will have identical images. Even though TV_1 and TV_2 exert no causal influence on one another, there is still an information link between them, since the state of one TV fixes (in a non-causal sense) the state of the other (Figure 1). The channel conditions, that the two TVs have the same source, guarantees this relation.

While Dretske's example shows that there needn't be a causal influence between systems for one to carry information about the other, his example nevertheless relies on a causal connection through a common cause. Despite this connection, the relevant point remains. The causality constraint cannot be motivated by general considerations about what is required for one system to carry information about another. In the following discussion of the circadian system, I will show that the circadian system, and clocks more generally, directly gather information about time while violating the causality constraint (without relying on a common cause). Instead, the endogenous rhythms of clock mechanisms themselves explain how they have this information gathering capacity.

3.2 The integration argument

Kant (in)famously argued that intuitions of space and time serve as preconditions for the possibility of any experience of an objective world

---

[4] A similar point is made in (Dretske [1969], p. 50): A causal connection between perceivers and what is seen is 'something that happens to be true of an enormous number of things we see', but it needn't always be present.

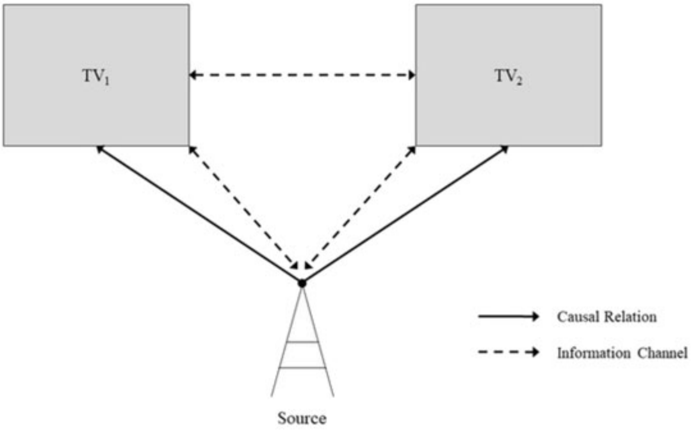

Figure 1. TV$_1$ and TV$_2$ carry information about one another (dashed line) yet there is no causal influence from one TV onto the other.

whatsoever. Only by embedding the deliverances of the individual senses into a spatial and temporal framework can one have experiences of and make sense of the outer world. In this picture, the representations of space and time are in an important sense prior to the individual senses, since it is only by embedding experiences within a spatiotemporal framework that we can have experience of an objective world, therefore, the representation of space and time are not proper parts of any of the individual senses. While it's difficult to evaluate Kant's own arguments for the priority of temporal representations, Matthen ([2014]) has recently argued that while Kant's arguments likely fail, a Kantian-inspired argument can be run where the representation of time is not sensory since it provides a framework for the interpretation of the individual senses.

The starting point for Matthen's argument is to notice that, as a matter of fact, representations of time serve as a 'common measure' for the deliverances of the individual senses. The events detected by the sensory systems are ultimately organized within a temporal framework to create a unified representation of the external world. We do not simply see lightning and hear thunder, but we perceive these events as being simultaneous or as being separated by some interval.

Importantly for Matthen, not only does the representation of time serve as a supramodal common measure for the senses, but how temporal properties are attributed to events in the world gives us reasons for thinking that temporal representations are at one remove from sensory systems. Adapting an account from Phillips ([2008]), Matthen argues that temporal properties are attributed to events in the world through a meta-experience that initially attributes temporal properties to experiences themselves. The individual sensory systems, in

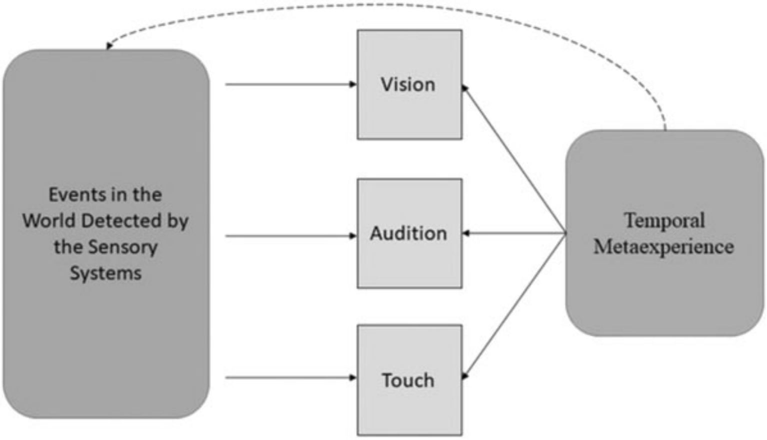

Figure 2. The meta-experience that Matthen proposes is one that monitors the individual sensory systems and exports the temporal properties of those sensory processes to the events in the world. The individual sensory systems have direct connections to the external world, while the temporal meta-experience goes through the individual sensory systems.

tracking the world, provide creatures with modality specific experiences. These experiences, construed as mental events, are themselves temporally structured in that they occur at particular moments, have durations, and stand in temporal relations to other experiences and events in the world. A mechanism tracks the temporal structure of these sensory experiences and then these temporal properties that are initially attributed to the sensory experiences themselves are exported and attributed to those events in the world that those sensory experiences detect.

It is through these meta-experiences that operate over the individual senses that the various deliverances of the individual sensory systems are coherently integrated into a single unified temporal order. Since the temporal exportation process operates over the individual senses it is not itself a sensory system nor is it a part of any of the senses (Figure 2).[5] In this way, Matthen argues that due to the role of temporal representation in structuring the deliverances of the individual sensory systems we have reasons for denying that the representation of time is due to a genuine sense of time.

However, even if we were to accept everything that Matthen says about the structuring role of representations of time in our conscious sensory experience, and we were to accept the meta-experiential account that Matthen adopts from Phillips, the conclusion that there is no sense of time would not follow. While Matthen describes a mechanism that could explain how we

---

[5] This model of temporal perception is widely accepted in the scientific literature under the guise of the scalar expectancy theory (Gibbon [1977]; Gibbon *et al.* [1984]).

consciously experience certain temporal properties (temporal relations and perhaps duration), what he concludes about that capacity cannot be generalized to other timekeeping capacities organisms possess. Our conscious temporal experience does not exhaust our timekeeping capacities, and as a result, any conclusions about temporal representation drawn from considering conscious temporal experience must be restricted to temporal experience alone. In the following sections, I will argue that the mammalian circadian system constitutes an entirely separate timekeeping capacity that provides animals with information about the approximate time of day and that this system constitutes a genuine sense of time.

4 Circadian Systems

Let's begin with a general point about animal behaviour. For animals to successfully navigate their environments they need to coordinate their behaviours, at various timescales, with the temporal structure of the events around them. From the millisecond timing involved in sensorimotor navigation, to the planning of action over the scales of seconds, minutes, hours, days, and longer, each of these timescales presents temporal properties of the environment that animals take into account to navigate their environment. In some cases, animals must keep track of the temporal structure of events regardless of when they occur in the day (for example, intercepting a falling apple regardless of whether it's morning or evening). However, other behaviours require keeping track of when in the day certain events occur. This latter sort of ability is perhaps most salient in cases of foraging or hunting. Food is available in certain locations at certain times of the day and in other locations at other times of the day. It's this second sort of capacity that we'll focus on.

Throughout the animal kingdom many animals exhibit patterns of behavioural and internal activity that have roughly twenty-four-hour periods (some patterns seem hardwired while others are acquired). These behaviours include things like sleep-wake cycles, eating patterns coordinated with local food sources, hormone regulation, body temperature regulation, and navigation (Moore [1997]), and even patterns pertaining to the effectiveness of memory formation (Ruby *et al.* [2008]).

While the existence of these daily patterns may seem to imply that there is some sort of mechanism (even a single mechanism) that represents the approximate time of day, one could deny that animals need to represent time at all to behave in these ways (see discussions in Gallistel [1990], [1996]; Mulder *et al.* [2013]). Taking just one example, pelicans seem to have an uncanny ability to be around when fishing fleets pull back into the docks. Since fleets come in around the same time of day every day, pelicans seem to have an uncanny ability to be at the right location at the right time of day. There is a temporal pattern to their behaviour. But the explanation of this behaviour needn't make any appeal to the pelican's ability to represent time. Instead, the pelican behaviour is cued by the sound of the incoming boats. Since the arrival of the fleet has a temporal pattern to it, the pelican behaviour cued by their sound also exhibits this temporal pattern. Pelicans only need to track sound, and not time, to do this.

While this explanation of circadian behaviours in terms of environmental cues may succeed in certain cases, it fails to provide a general account of the circadian behaviours since many of these behaviours, even those that are learned, persist in free-running conditions in which all environmental cues (for instance, constant light, temperature, and so on) have been eliminated.[6]

A long line of research, stretching back to the 1700s, has shown circadian patterns persist in free-running conditions.[7] For instance, in a study by Rosenwasser *et al.* ([1984]), rats were trained to anticipate food in particular locations at specific times of day in free running conditions. Since, no external cues were available, it seemed as though they must be relying on an internal timekeeping mechanism to initiate their behaviours.

Furthermore, individual animals kept in free-running conditions for extended periods exhibit circadian patterns that drift away from a standard twenty-four-hour period. Some animals begin to operate on a cycle somewhat shorter, others somewhat longer (Bolles and Moot [1973]). However, once the drift occurs, individual animals enter into reliable patterns that can be used to predict their behaviours. If, for instance, animals were trained to expect food at a location during the final phase of their circadian cycle, whether that cycle has a twenty-five-hour or twenty-three-hour period, then the animal specific circadian period could be used to successfully predict behaviour. Since external cues were controlled for in the free-running conditions, the individual variability in circadian rhythms seemed to result from creature internal factors.

Because of this long line of studies, it is accepted that many circadian behaviours result from the operation of an internal clock. This, however, left it unanswered as to what type of clock mechanism underpinned these behaviours. Specifically, there were two competing models that attempted to

---

[6] An anonymous referee helpfully suggested that another reason for positing an internal clock is that many animal behaviours require a significant amount of time to initiate. If animals were responding to local cues, then many of the resulting behaviours would occur too late. Therefore, animals must possess an internal clock to anticipate when events in the world occur.

[7] In 1729, Jean Jacques d'Ortous de Marain observed that the leaves of *Mimosa pudica* would open and close on an approximately twenty-four-hour cycle in free running conditions. Later studies, in the 1800s, showed that this behaviour could not be explained by appealing to external cues and as a result must be due to something internal to the plant; see discussion in (Sollars and Pickard [2015]).

account for the observed circadian behaviours—interval timers and period timers.[8,9]

According to interval timing models of circadian behaviours the timekeeping device always reports a temporal relation between two events in the world—a marker event that begins the timing mechanism and a target event that signals the end of the measured interval. A helpful way to understand how interval timing mechanisms operate is by analogy to hourglasses. When some relevant 'marker' event is detected the timer is started. As time progresses, something in the interval timer accumulates ticks of the clock (for instance grains of sand). When some second 'target' event is detected, the quantity of the accumulated medium provides information about the amount of time bounded by the two events. If interval timing mechanisms could account for circadian behaviours, then we should be able to specify pairs of reliably detectable events that can be used to time circadian behaviours.

According to the alternative period timing models, circadian behaviours are accounted for by an internal oscillator that is best understood by analogy to a twenty-four-hour analogue clock. Period timers operate by continuously cycling through a fixed sequence of states in a twenty-four-hour period where the state transitions occur at a predictable rate. The specific phases of the oscillator represent specific times of day. Period timer explanations do not require that animals pick out specific events in the world to represent the time of day. Instead, the reliable operation of the internal oscillator is sufficient on its own.

While there are many differences between interval and period timers, their most significant differences lie in their semantics and their mechanical implementations. Let's consider their implementation first. Interval timers are typically understood as representing time through a straightforward accumulation process. There is a monotonic relationship between the number of accumulated ticks and the represented duration and it is the amount of the accumulated substance that directly carries the information about the length of the interval. As Wearden ([2001]) puts it, longer intervals are measured by there being more of something.

Period timers, on the other hand, do not represent time through a straightforward monotonic relationship between the accumulation of a substance and time. Instead, as mentioned above, period timers represent time through an oscillator whose phase states pick out specific moments in time.[10]

---

[8] A similar discussion to this appears in (Mulder *et al.* [2013]).

[9] Circadian researchers have widely rejected interval timing models. However, seeing how interval timing models fail allows us to better develop a semantics and meta-semantics for internal clocks.

[10] Period timers may contain accumulation processes as part of the overall mechanism that gives rise to the oscillation. Similarly, interval timers may contain oscillators as parts that gives rise to the relevant type of accumulation process. Nevertheless, the types of timing mechanisms differ

The semantics of these mechanisms also differ significantly. Interval timers pick out an interval that holds between two specific events—the marker event, e_M, and the target event, e_T. Their content is always of the form $<e_T$ is I since $e_M>$, where I is the duration of the interval between e_T and e_M. Period timers, on the other hand, do not pick out temporal relations between pairs of events. Instead, period timers say of the current moment in time that it has a specific temporal property—their content is always of the form <now is T_P>, where T_P is a temporal property predicated of the current moment in time.

As we'll see, the behavioural and physiological evidence strongly suggest that we should adopt a period timing model. The first line of evidence is one that we've already seen. The persistence of circadian behaviours in free-running conditions rules out explanations that appeal entirely to local environmental cues. If in free-running conditions there are no external cues that can be used to trigger behaviours, then no external cues can reliably be detected to serve as interval bounds.

Perhaps, however, the very circadian behaviours themselves, like feeding or resulting states like hunger, blood-sugar levels, and so on, could serve as marker events that form interval boundaries.[11] Several problems arise for this suggestion. First, conditioned circadian feeding behaviours survive the occasional 'bad day' in which food is withheld (Biebach *et al.* [1989]; Mulder *et al.* [2013]). If animals were using an interval timer with the previous days feeding, or any resulting state, as a marker event, then we wouldn't be able to explain how animals continued to show the conditioned anticipatory feeding behaviour when that 'marker' event was removed.

Second, no pacemaker or accumulator is perfectly noise-free. As a result, interval timer models all fit Weber's Law according to which the precision of a measurement decreases as the measured magnitude increases (even at the short timescale of milliseconds) (Gallistel and Gelman [1992]; Wearden [2001]; Malapani and Fairhurst [2002]; Dehaene [2003]). Since circadian behaviours require timing over hours, we would expect a large amount of noise in the corresponding representations of time. However, while this sort of noise is found for a number of explicitly interval timing behaviours (Gibbon *et al.* [1984]), this type of noise is not found for circadian behaviours (Ko *et al.* [2003]). Interval timers do not fit the behavioural data.

While the behavioural evidence is telling against the interval timing model, the final nail in the coffin came from evidence concerning the neural mechanism that implements the circadian timer. The first step in uncovering the mechanism underlying circadian behaviours was the localization of the

---

in which properties of the representational vehicle encode temporal information. Thanks to Carl Craver and Lawrence Ward, and an anonymous referee at this journal for raising this point.

[11] Thanks to Lawrence Ward for raising this objection.

mechanism. It was discovered through a series of studies, beginning in the 1970s, that many circadian behaviours in mammals are regulated by the suprachiasmatic nucleus (SCN), a region of hypothalamus, which in turn regulates various peripheral oscillatory systems.[12]

To illustrate the role of the SCN in circadian behaviours, consider a study by Ralph *et al.* ([1990]). Two populations of hamsters were selectively bred with distinct circadian cycles. One group of hamsters possessed normal circadian systems with approximately twenty-four-hour periods, while the second, 'mutant', group had circadian rhythms with approximately twenty-hour periods. Upon the ablation of the SCN in hamsters of either group circadian rhythms halted—sleep/wake cycles, daily activities, temperature regulation, and hormone regulation all became erratic. However, many circadian behaviours were recovered upon the surgical transplantation of foetal SCN tissue.[13] Perhaps most interesting, however, is that when hamsters received donor SCN tissue from hamsters of the other circadian population, the restored circadian rhythms had a period that matched that of the donor population. So, a twenty-four-hour hamster receiving SCN tissue from a twenty-hour donor, would have restored circadian behaviours with a twenty-hour period, not a twenty-four-hour period.

A study by Maruyama *et al.* ([2007]) further illustrates the central role of the SCN. In this study two groups of rats were exposed to elevated temperatures for a fixed five-hour period of the day. One group of rats had bilateral damage to their SCNs whereas a control group had intact SCNs. Both groups exhibited lowered core temperatures for an extended time of the day presumably as a means of combatting the expected heat stress. However, only the control group, with intact SCNs, exhibited lowered core temperature that coincided with the time of day that the training heat stress was administered. The experimental group, while showing periods of lowered body temperature, did not produce this response at a predictable time of the day. They could not coordinate their body temperature with the timing of events in their environment.

Furthermore, the adaptive value in the wild of a properly functioning SCN was shown in a study that involved the long-term monitoring of two chipmunk populations (DeCoursey *et al.* [2000]). One group had their SCNs surgically removed, while the second, control, group underwent a similarly invasive surgery but were left with intact SCNs. After an eighteenth-month

---

[12] For the history of SCN research, see (Weaver [1998]).

[13] A fascinating aspect of these studies is that although some circadian patterns were recovered upon receiving the transplant, it was found that the neural connections between the SCN and the rest of the brain were not properly formed. The SCN was coordinating behaviours through factors released into the bloodstream. If we take this coordination to involve the transmission of information to downstream systems, then it follows that not all information signalling in the brain is the product of action potentials (Silver *et al.* [1996]; Guo *et al.* [2006]).

period in the wild, the chipmunks in the experimental group were significantly more likely to have been eaten by predatory weasels than those in the control group. It was ultimately determined that the SCN-damaged chipmunks were active within their dens at atypical times—times when chipmunks are typically quietly asleep. The improperly timed commotion made by the SCN-damaged chipmunks made them much more noticeable to passing weasels, and as a result, many were eaten by weasels. The loss of the SCN caused the behaviour of the chipmunks to no longer be properly synchronized with their environment.

Knowing that the SCN plays a role in circadian behaviours was a start. However, it was only with a mechanistic understanding of the SCN that the interval timing approach was struck dead. While the details are still being worked out (Gachon *et al.* [2004]; Liu *et al.* [2007]; Lowrey and Takahashi [2011]; Bechtel [2011]; Buhr and Takahashi [2013]; Ye *et al.* [2014]; Bano-Otalora and Piggins [2017]), the general intracellular story is well understood and applies across an enormous range of animals (from insects, birds, mammals, fish, and others). For our purposes, however, a simplified version of the model will suffice. The SCN operates as an endogenously driven molecular clock governed by transcription/translation feedback loops and the cyclical accumulation and breakdown of various proteins that gives rise to electrical oscillations with approximately twenty-four-hour periods. The system can be understood as containing a positive and a negative component. The proteins BMAL1 and CLOCK heterodimerize and translocate to the nucleus where they initiate the transcription of a various genes, including *Per* and *Cry*. As *Per* and *Cry* are expressed and concentrations of PER and CRY in the cell increase, PER and CRY heterodimerize, and are part of a complex that inhibits BMAL1:CLOCK from initiating the transcription of *Per* and *Cry*. PER and CRY are then slowly broken down within the cell, and BMAL1:CLOCK is again able to initiate the expression of *Per* and *Cry*. This back and forth, with BMAL1:CLOCK as the positive component, and PER:CRY as the negative component, give rise to twenty-four-hour oscillations in gene expression that ultimately gives rise to a twenty-four-hour oscillation in the electrical activity of individual cells in the SCN.[14] As a result, the individual cells of the SCN exhibit oscillations in their standing firing rates with approximately twenty-four-hour periods that can even be detected with *in vitro* cell cultures (Beaulé *et al.* [2011]).

This molecular story is an intracellular process that is found in a similar form throughout the animal kingdom. However, individual SCN cells when kept in isolation tend to behave somewhat erratically. It is only at the network

---

[14] How protein levels in the cell influence neuronal firing rate is still under investigation. However, see (Vasalou and Henson [2010]) for a proposal.

level that the mammalian SCN exhibits robust circadian oscillations. The intracellular molecular mechanism determines the period of the oscillations, yet it is the circuitry (for example features like cell density, connectivity and the concentrations of various neurochemicals) that synchronizes the oscillations of SCN subregions and the SCN as a whole (Yamaguchi *et al.* [2003]; Welsh *et al.* [2010]; Brancaccio *et al.* [2014]).[15]

Importantly, the accumulation processes involved in the SCN do not encode temporal information in the way that accumulation processes encode temporal information in interval timing mechanisms. Later times are not represented by there being more of something (there is no single substance that continuously increases through the day). Instead, accumulation processes give rise to the oscillatory behaviour of the system and times are represented by the phase properties of the system.[16]

One important aspect of the SCN's operation that we have yet to discuss is that like any clock the SCN occasionally needs to be calibrated with the time of day. The most significant mechanism for SCN calibration is through the melanopsin visual channel that provides coarse-grained information about ambient light levels through photosensitive retinal ganglion cells (Freedman *et al.* [1999]). Importantly, this input, which bypasses the standard visual system, makes no direct contribution to visual experience.

The sources of calibration, however, should not be overstated. No single source of calibration is required as other sources of circadian calibration come from food, forced sleep, exercise, and various other sources (Stephan [2002]; Schibler *et al.* [2003]). There is no need for any reliable detection of specific events that could serve as marker events for an interval timer. Furthermore, allowing causal influences for the calibration of the system does not undermine the response to the non-causality argument. While the calibration cues are time-related, they are not cases in which the temporal properties of the environment themselves are exerting causal control over the circadian system.[17]

---

[15] It's for these system level features of circadian systems that the conclusion of the article is restricted to mammals.

[16] A referee pointed out that while the SCN may represent time of day, it may not transmit this information. Rather, the system may merely issue GO commands for peripheral systems to initiate various process, in much the way that a conductor signals to musicians to start playing. While this wouldn't change what the SCN represents, it would alter how we understand neural information transmissions.

[17] Noting that the circadian mechanism is involved in the control of various behaviours and that it can be calibrated through a variety of means gives us reasons for thinking that the SCN's semantic content is sufficiently general that it can contribute to these various processes—the content is not behaviour-specific or calibration-cue-specific. Time of day, then seems like a viable interpretation of the common semantic content that the system could contribute to these various behaviours. Thank you to an anonymous referee for suggesting this line of argument.

The case of the SCN provides us with empirical reasons for rejecting Matthen's Integration Argument. Regardless of whether Matthen is correct about how animals perceive temporal relations of earlier than, later than, and simultaneous with, his argument says nothing about how animals represent other temporal properties. With the SCN we see that mammals possess neural systems capable of directly acquiring information about the approximate time of day in a manner that does not rely on information gathered by any of the other sensory systems. Furthermore, in noticing that the SCN represents the time of day through its own internal rhythms, we begin to see how the non-causality argument also fails. In the next sections, we'll see more clearly how non-causality argument fails.

5 The Semantics of (Internal) Clocks

To show how the SCN is capable of directly gathering information about the time of day, while violating the causality constraint, we need to first say something about the informational content of the circadian system. Only then can we explain how the system acquires this content.

Let's begin with a general point. Two types of errors can result in an animal showing up to the right place at the wrong time (for instance, expecting food). First, the animal may misremember when the food is available. Second, the animal may be mistaken about what time it is—that is, its clock may malfunction. It's this second type of error that we'll focus on.

To understand how this type of error is possible, the content of the circadian system must have both referential and predicative components. Only with this sort of structure can the representation be said to accurately or inaccurately describe what time it is right now. That is, clocks say something about (they predicate something of) what time it is right now (the referent). To help our analysis of the content of the circadian clock (and clocks more generally), consider the following sentence

> CLOCK: It is 5 pm.

CLOCK tells us something about what time it is right now, so we should understand the 'it' as an indexical that picks out the present moment or now.[18] While it's a matter of debate in the metaphysics of time as to what particular moments in time amount to and whether there is anything special about the current moment in time (see Meyer [2013]), for our purposes we can remain neutral towards these debates. Whatever metaphysics of time we adopt, there is something that we pick out with the phrase 'this moment in

---

[18] There are clear parallels between the semantics of clocks and the role of indexicals that Perry ([1979]) isolates. In both cases, indexicals are crucial for the individual to situate themselves in the environment.

time' or 'now', and it is to that sort of thing, whatever it is, that clocks attribute some temporal property.

Having pinned down the referential component of CLOCK, we need to understand the overall logical structure of CLOCK. Statements with the surface structure of CLOCK are notorious for having two possible logical forms depending on how we interpret 'is' (Russell [1905]). There is the 'is' of predication, as in '2 is prime' in which we attribute a property to some object. There is also the 'is' of identity, as in '2 is 2', in which we say that one object stands in the relation of being identical to an object (namely, itself). In cases involving the is of identity, we can typically add the phrase 'identical to' after the 'is' in the sentence without any problem (for instance, '2 is 2' can be read as '2 is identical to 2'). Correspondingly, there are two possible interpretations of CLOCK that depend on differing uses of 'is'.

I think there are clear reasons for rejecting the identity reading of CLOCK since it leads to contradiction. If CLOCK is truthfully uttered on a Wednesday, the 'it' will pick out a specific moment in time, $t_{\text{Wednesday}}$, and read as an identity, CLOCK would claim that $t_{\text{Wednesday}} = 5$ pm. Now suppose that CLOCK is also truthfully uttered on a Thursday. In this case the 'it' will pick out t_{Thursday}, and CLOCK would claim that $t_{\text{Thursday}} = 5$ pm. However, it's clearly false that $t_{\text{Thursday}} = t_{\text{Wednesday}}$, since those moments occur on different days. Since the uses of CLOCK on Wednesday and Thursday are both true, CLOCK cannot express an identity (since transitivity leads to a contradiction). Rather, it must be read as predicating some property to the present moment—namely, the property of being 5 pm is truthfully attributed to both $t_{\text{Wednesday}}$ and t_{Thursday}. As a result, we need to find some property that is picked out by 'is 5 pm'.

That a biological clock might express a predicate like is 5 pm might strike some as odd. The predicate is 5 pm is embedded within our culturally constructed timekeeping practices and relies heavily on cultural conventions and scientific technology (Tal [2013], [2016]). Since biological systems, especially those shared with non-human animals, are independent of these cultural contributions, it becomes unclear how the circadian system could have this type of content (especially if one has neo-Whorfian leanings (along the lines of Boroditsky [2011]). To make sense of all of this, we need to notice how the temporal properties picked out by our timekeeping devices are constructed (or imposed upon the world) by our timekeeping mechanisms themselves.[19]

To illustrate the relation between timekeeping mechanisms and the temporal properties they pick out, let's consider an analogous spatial case. Imagine travelling down a long road that lacks any distinctive features

---

[19] An alternative way of putting the point is that the temporal predicates and the properties they pick out are determined by the physiology of the system.

other than rest stops at regular intervals. Unfortunately, your car lacks an odometer, speedometer, or clock, so you're unable to use any of the standard measurements to figure out where you are between the rest stops. However, for some reason, your car has a light system that cycles between four different lights (blue, green, yellow, and red). If you travel at the right speed, the light system cycles through the entire sequence of four lights exactly once between each rest stop. This light system, provided your speed is constant, can then be used to represent your spatial location between each rest stop. The transition of the lights from one colour to the next provides us with a system of divisions that we can use to impose structure on the space between each rest stop. While these divisions are not independent our measurement devices, once we adopt the light system as a means of dividing up the road, and we adopt a specific travelling speed, the lights provide us with information about our objective location along the road.

In this way, the spatial properties picked out by a measurement system and the measurement system itself can come hand in hand. When the light system is showing a blue light, we take the system as representing here is a blue location, and the predicate is a blue location picks out just those spatial regions where the light system will be in a blue state, provided the speed of the car remains constant.

We can think of clocks as imposing a similar sort of structure on time. Consider each passing day as being analogous to the regularly spaced rest stops. We then simply use a system whose states, under appropriate conditions, cycle in such a way that they neatly divide each day into distinct temporal regions. This is how our timekeeping mechanisms impose structure on time that they then go on and represent.

Since the circadian system cycles through its states at a regular rate, these states provide a regular partitioning of the day into distinct time periods. As theorists, we'll naturally describe these temporal divisions in terms of hours, minutes, and seconds, but this doesn't commit us to the claim that the circadian system represents time in terms of those units. Instead, we can use our public language to describe extensionally equivalent temporal categories to those categories established by the circadian system. Importantly, what divisions the circadian system makes are determined the physiology of the circadian system itself and not the result of the representational system latching onto or mirroring mind-independent joints in nature. Instead, the physiology constructs the temporal properties in the world.

What we have is an account of the content of the circadian system. It represents the time of day in the format of <now is T_{cx}>, where the predicate 'is T_{cx}' picks out a temporal category or property that is defined relative to the operation of the circadian system itself. With this account in hand, in the next section we'll describe how the circadian system more precisely comes to have this content.

6 An Information-Theoretic Account

To finally show how the non-causality argument fails, we'll see in this section how the circadian system comes to have its content while violating the causality constraint. The key to understanding how the appropriate information link is established is to notice what Smith ([1988]) called 'the participatory nature of clocks'. Clocks don't only represent time, but they do so in virtue of how their underlying mechanisms evolve or unfold in time.

Consider how a mechanical watch tells time. Energy is stored in a compression spring and is slowly expended through a series of gears and springs that are finely tuned to move the hands of the clock in a predictable manner. If we were to lay out the causal story of how this sort of mechanism works, the entire story could be told step by step through the transference of energy from one part of the mechanism to the next. In no place in the story do we have to appeal to time itself being a causal force—time itself is not another cog in the machinery of the clock. Instead, the specific causal interactions between the parts of the watch are governed by physical law such that they evolve in a temporally predictable manner. It is this lawful unfolding of clock processes that allows clocks to represent time.

A similar story explains how the circadian system provides animals with a direct information link with the approximate time of day. To make our discussion a little clearer, let's introduce some additional terminology. Let's label the states of the circadian system s_1, s_2, [. . .], s_n. Each of the states of the circadian system will represent, or carry the information, a specific time of day. For example, s_1 is semantically interpreted as carrying the information that <now is T_{c1}>, where T_{c1} is the predicate that picks out the first division of circadian time.

As mentioned earlier, for a system to carry information about some aspect of the world is for the conditional probability of the world's being a certain way, given the state of the information carrying system, is 1 (provided the background channel conditions are intact). In this case, to show that the circadian system provides the appropriate information link, all that needs to be shown is that provided the appropriate background channel conditions (a healthy and calibrated SCN), the conditional probability of the time being T_{cx}, given that the circadian system is in state s_x, is 1.

To see that this sort of information link is established is fairly easy given everything described so far. Instead of the transfer of energy through a series of gears, the circadian clock depends on the timely evolution of its biochemical and electrical states. Each transition of the system, from one state to the next,

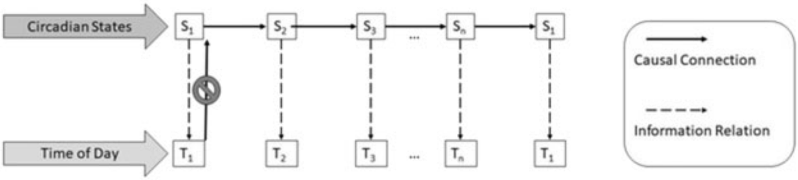

Figure 3. The states of the circadian system carry information about the time of day, despite the fact that there is no causal influence on the states of the circadian systems from the time of day.

is the result of causal interactions between the various biochemical components of the SCN (as described in Section 3). Furthermore, due to the laws that govern the physiological processes of the brain, these transitions evolve in a temporally predictable manner. If the circadian system is calibrated with the local time, through any means, to be in state s_1 at the beginning of the day, then the evolution through time of the circadian mechanisms guarantees that each subsequent state of the system will correspond with the changing time of day (Figure 3).

Taking this correspondence on board, we can understand why the informational relation holds. The conditional relation between the states of the circadian system and the time of day is due to how the progression of circadian states through time is under the nomic control of time-dependent processes. Notice that the explanation of how the circadian system acquires information about time while violating the causality constraint is plausibly only available for something like temporal representation. Clocks can violate the causality constraint because of their participatory nature. Similar stories cannot be given for other sensory systems like vision—the visual system represents many things without participating in them.

We have an account of how the circadian system successfully represents time, but the true test is whether the circadian system is capable of misrepresentation. This brings up a general problem for information-theoretic accounts of content. Pure information theories, especially ones like Dretske's where the conditional probability of the world being a certain way given a particular system is 1, have difficulties explaining misrepresentation. If the world is not the way that the representational system 'describes', then technically, no information link is established, so the system does not carry any information about the world. For Dretske ([1981]), there is no such thing as misinformation.

Since we want an account of misrepresentation, we need to augment the simple informational account, and this is what Dretske ([1991], [1995]) himself does in several places. Cases of misrepresentation are handled through the addition of a teleological component. A system that carries information about the world when the appropriate background channel conditions are in place can misrepresent the world provided that the system has acquired the function of carrying information such that the system is treated by consumer systems as though the background channel conditions were satisfied.[20]

A clock has the function of telling the time, because when everything is functioning properly, the clock carries this information and we use the clock on the assumption that it is successfully carrying this information. Misrepresentation occurs when the background channel conditions fail to hold either through a breakdown in the clock mechanism itself (the gears becoming worn) or through a failure of proper calibration (you forget to reset your watch after a flight). In either case, without any reason to believe that the channel conditions are violated, we treat the clock as though it is operating properly and our behaviours become poorly coordinated with the environment.

The same story applies to circadian systems. When the background channel conditions are violated, either through failures of calibration (for instance, long distance air travel, shift work, or artificial lights) or through failures of the circadian system (due to changes in biochemistry), consumer systems still treat the circadian system as though it were properly conveying information about the time of day, and as a result we have cases of misrepresentation.

The appeal to function also provides the resources to respond to a possible objection.[21] Consider the case in which the two TVs are receiving signals form a common source. Now, consider Stan and Stu, who are each watching a TV. Due to the set-up of the scenario, Stan's sensory states not only carry information about the TV in front of him, but they also carry information about the state of the TV's internal circuitry, the signals coming through the wires, the state of the other TV, and even the state of Stu's sensory system. Yet it would be an unfortunate consequence if we were committed to saying that Stan can sense the state of Stu's mind. The point generalizes. Information is easy to come by, and any given sensory system will carry information about many unrelated aspects of the world. In the case of the sense of time, the circadian system also carries information about the probable location of the sun, the proximal causes of the current state of the circadian clock, and so on. What then distinguishes what we would typically describe as the sensory content of the system from all this informational mess?

The causality constraint would seem to provide an answer. Stan senses the TV in front of him, but not Stu's mind, because only the TV exerts causal influence over Stan's sensory system. However, this only shifts the problem (Dretske [2000]). The internal state of the TV also counts as having causal

---

[20] Dretske ([1981]) argued that these functions are determined by a privileged learning period. However, see (Nanay [2014]) for an alternate modal account of function determination.

[21] Thank you to an anonymous referee for raising this objection.

control over Stan's sensory system (although less proximal than that of the TV), but the state of the retinal photoreceptors in Stan's eye count as an even more proximal cause of the state of Stan's sensory system. Yet, in the same way we hesitate to say that Stan senses Stu's mind, we also hesitate to say that Stan senses his own retina or that he senses the internal circuitry of the TV. The causality constraint only rules out sensing Stu's mind, but not any of the other possibilities.

Instead, consumer systems can help. Sensory states have a rich informational content due to how they relate to the world. However, our attributions of sensory content track those pieces of information typically contained in the informational content of the system that downstream consumer systems treat the sensory system as possessing. Stan senses the TV, and not the internal circuitry of the TV, since consumer systems make use of the information about the surface properties of the TV in ways that they do not make use of other information.[22] That other information simply isn't made available in this way.

Another way to put the point that the system is representing time of day, and not something else, is the following: Consider the causal story of a consumer system that is making use of the circadian signals themselves, and not any information they contain, and a consumer system that is making use of the temporal information encoded in the circadian signals. In both cases, the causal story about consumption is the same. The signals produced by the circadian system are those things that have a causal effect on consumer systems. The causal story alone will not distinguish between a case in which the system is making use of a particular piece of information versus merely making use of semantically un-interpreted signals. However, there is a non-causal asymmetry that gives us a handle on this notion of making use of.

A sensitivity to the signals of the circadian system only has an adaptive value to the organism provided that they carry information about the time of day. If, due to mis-calibration or some genetic defect, the signals no longer carry information about the time of day, then the organism will find no adaptive advantage in coordinating its activities with the circadian signals. However, information about the time of day, regardless of how it is conveyed to downstream systems does confer an adaptive advantage. It's in this way, while the causal story of how the circadian system is being used by consumer systems does not distinguish what is being used, we can still understand that it is time of day and not something else that is being used, since all uses of the circadian system asymmetrically depend on the use of the system in which its information about the time of day is being exploited.[23]

---

[22] Importantly, the sort of use that I am appealing to here does not involve any person-level inference.

[23] While similar to Fodor's ([1987]) asymmetric dependency theory, the proposal here puts the asymmetry in the consumption and not the production of representational states.

While the circadian system may carry information about more proximal causes of the operation of the circadian system it does not have this information as its sensory content. The circadian system produces sensory representations of time, while violating the causality constraint, because when the channel conditions are intact the circadian system carries information about the time of day, and consumer systems treat the system as though it is carrying this information.

7 Conclusion

Here is a description of a psychological mechanism that provides organisms with a direct informational link with the approximate time of day. That is, we have a description of a genuine sensory system for time. Furthermore, given the empirically uncovered physiology that underpins the circadian system, we have an explanation for how this direct information link is established while nevertheless violating the causality constraint. Importantly, the discovery that we have a sense of time is not an introspective discovery but relies on closely understanding the various mechanisms that animals use to keep track of the temporal structure of their environment. It is only through these details that we see how the non-causality and integration arguments fail.

I have not addressed whether the other timekeeping capacities possessed by animals are sensory or not. However, I am doubtful that any other system could qualify as a genuine sense of time. Our other timekeeping capacities seem to be distributed across supramodal and modality specific timekeeping mechanisms (Lewis *et al.* [2003]; Buhusi and Meck [2005]; Ivry and Schlerf [2008]; Wittmann and van Wassenhove [2009]). As a result, they would not stand on their own as independent means of directly providing organisms with information about time.

Acknowledgments

I would like to thank Murat Aydede, Daniel Booth, Carl Craver, Emma Esmaili, Eric Margolis, Christopher Mole, David Sackris, Evan Thompson, Lawrence Ward, and two anonymous referees for this journal for comments on earlier drafts. This article was written in part under the funding of ERC Consolidator Grant STYDS (grant number 726251) to Bence Nanay.

Centre for Philosophical Psychology
University of Antwerp
Antwerp, Belgium
gerardo.viera@gmail.com

References

Aristotle. [2004]: *De Anima*, London: Penguin.

Bano-Otalora, B. and Piggins, H. [2017]: 'The Mammalian Neural Circadian System: From Molecules to Behavior', in V. Kumar (*ed.*), *Biological Timekeeping: Clocks, Rhythms, and Behavior*, New Dehli: Springer.

Beaulé, C., Granados-Fuentes, D., Marpegan, L. and Herzog, E. D. [2011]: '*In Vitro* Circadian Rhythms: Imaging and Electrophysiology', *Essays in Biochemistry*, **49**, pp. 103–17.

Bechtel, W. [2011]: 'Representing Time of Day in Circadian Clocks', in A. Newen, A. Bartels and E.-M. Jung (*eds*), *Knowledge and Representation*, Palo Alto, CA: CSLI Publications and Paderborn.

Benovsky, J. [2012]: 'The Causal Efficiency of the Passage of Time', *Philosophia*, **40**, pp. 763–9.

Biebach, H., Gordijn, M. and Krebs, J. R. [1989]: 'Time-and-Place Learning by Garden Warblers, Sylvia Borin', *Animal Behaviour*, **37**, pp. 353–60.

Bolles, R. C. and Moot, S. A. [1973]: 'The Rat's Anticipation of Two Meals a Day', *Journal of Comparative and Physiological Psychology*, **83**, pp. 510–14.

Boroditsky, L. [2011]: 'How Languages Construct Time', in S. Dehaene and E. Brannon (*eds*), *Space, Time, and Number in the Brain*, London: Academic Press, pp. 333–41.

Boroditsky, L. and Prinz, J. [2008]: 'What Thoughts Are Made of', in G. R. Semin and E. R. Smith (*eds*), *Embodied Grounding: Social, Cognitive, Affective, and Neuroscientific Approaches*, Cambridge: Cambridge University Press, pp. 98–115.

Brancaccio, M., Enoki, R., Mazuski, C. N., Jones, J., Evans, J. A. and Azzi, A. [2014]: 'Network-Mediated Encoding of Circadian Time: The Suprachiasmatic Nucleus (SCN) from Genes to Neurons to Circuits, and Back', *Journal of Neuroscience*, **34**, pp. 15192–9.

Buhr, E. D. and Takahashi, J. S. [2013]: 'Molecular Components of the Mammalian Circadian Clock', *Handbook of Experimental Pharmacology*, **217**, pp. 3–27.

Buhusi, C. V. and Meck, W. H. [2005]: 'What Makes Us Tick? Functional and Neural Mechanisms of Interval Timing', *Nature Reviews: Neuroscience*, **6**, pp. 755–65.

Buonomano, D. V. and Karmarkar, U. R. [2002]: 'How Do We Tell Time?', *The Neuroscientist*, **8**, pp. 42–51.

Collin, S. P. and Whitehead, D. [2004]: 'The Functional Roles of Passive Electroreception in Non-Electric Fishes', *Animal Biology*, **54**, pp. 1–25.

Coull, J. [2011]: 'Discrete Neuroanatomical Substrates for Generating and Updating Temporal Expectation', in S. Dehaene and E. Brannon (*eds*), *Space, Time, and Number in the Brain*, Elsevier.

DeCoursey, P. J., Walker, J. K. and Smith, S. A. [2000]: 'A Circadian Pacemaker in Free-Living Chipmunks: Essential for Survival?', *Journal of Comparative Physiology*, **186**, pp. 169–80.

Dehaene, S. [2003]: 'The Neural Basis of the Weber–Fechner Law: A Logarithmic Mental Number Line', *Trends in Cognitive Sciences*, **7**, pp. 145–7.

Dretske, F. [1969]: *Seeing and Knowing*, Chicago: University of Chicago Press.

Dretske, F. [1981]: *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press.

Dretske, F. [1991]: *Explaining Behavior: Reasons in a World of Causes*, Cambridge, MA: Bradford Books.

Dretske, F. [1995]: *Naturalizing the Mind*, Cambridge, MA: MIT Press.

Dretske, F. [2000]: 'Simple Seeing', in his *Perception, Knowledge, and Belief: Selected Essay*, Cambridge: Cambridge University Press.

Fodor, J. [1987]: *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, Cambridge, MA: MIT Press.

Freedman, M. S., Lucas, R. J., Soni, B., Schantz, M. v., Muñoz, M., David-Gray, Z. and Foster, R. [1999]: 'Regulation of Mammalian Circadian Behavior by Non-Rod, Non-Cone, Ocular Photoreceptors', *Science*, **284**, pp. 502–4.

Fulkerson, M. [2014]: 'Rethinking the Senses and Their Interactions: The Case for Sensory Pluralism', *Frontiers in Psychology*, **5**, available at <doi.org/10.3389/fpsyg.2014.01426>.

Gachon, F., Nagoshi, E., Brown, S. A., Ripperger, J. and Schibler, U. [2004]: 'The Mammalian Circadian Timing System: From Gene Expression to Physiology', *Chromosoma*, **113**, pp. 103–12.

Gallistel, C. R. [1990]: *The Organization of Learning*, Cambridge, MA: MIT Press.

Gallistel, C. R. [1996]: 'The Perception of Time', in K. Akins (*ed.*), *Perception: Vancouver Studies in Cognitive Science*, Oxofrd: Oxford University Press.

Gallistel, C. R. and Gelman, R. [1992]: 'Preverbal and Verbal Counting and Computation', *Cognition*, **44**, pp. 43–74.

Gibbon, J. [1977]: 'Scalar Expectancy Theory and Weber's Law in Animal Timing', *Psychological Review*, **84**, pp. 279–325.

Gibbon, J., Church, R. M. and Meck, W. H. [1984]: 'Scalar Timing in Memory', *Annals of the New York Academy of Sciences*, **423**, pp. 52–77.

Grice, H. P. [1962]: 'Some Remarks about the Senses', in R. J. Butler (*ed.*), *Analytical Philosophy, First Series*, Oxford: Oxford University Press.

Guo, H., Brewer, J. M., Lehman, M. N. and Bittman, E. L. [2006]: 'Suprachiasmatic Regulation of Circadian Rhythms of Gene Expression in Hamster Peripheral Organs: Effects of Transplanting the Pacemaker', *Journal of Neuroscience*, **26**, pp. 6406–12.

Harris, L., Harrar, V., Jaekl, P. and Kopinska, A. [2010]: 'Mechanisms of Simultaneity Constancy', in R. Nijhawan and B. Khurana (*eds*), *Space and Time in Perception and Action*, Cambridge: Cambridge University Press.

Ivry, R. B. and Schlerf, J. E. [2008]: 'Dedicated and Intrinsic Models of Time Perception', *Trends in Cognitive Sciences*, **12**, pp. 273–80.

Kalmijn, A. J. [1982]: 'Electric and Magnetic Field Detection in ElasmobranchFishes', *Science*, **218**, pp. 916–18.

Kalmijn, A. J. [2000]: 'Detection and Processing of Electromagnetic and Near–Field Acoustic Signals in Elasmobranch Fishes', *Philosophical Transactions of the Royal Society B*, **355**, pp. 1135–41.

Keeley, B. L. [2002]: 'Making Sense of the Senses: Individuating Modalities in Humans and Other Animals', *Journal of Philosophy*, **99**, pp. 5–28.

Ko, C. H., McDonald, R. J. and Ralph, M. R. [2003]: 'The Suprachiasmatic Nucleus Is not Required for Temporal Gating of Performance on a Reward-Based Learning and Memory Task', *Biological Rhythm Research*, **34**, pp. 177–92.

Lewis, D. K. [1973]: *Counterfactuals*, Oxford: Blackwell.

Lewis, P. A., Miall, R. C., Daan, S. and Kacelnik, A. [2003]: 'Interval Timing in Mice Does not Rely upon the Circadian Pacemaker', *Neuroscience Letters*, **348**, pp. 131–4.

Liu, A. C., Welsh, D. K., Ko, C. H., Tran, H. G., Zhang, E. E., Priest, A. A., Buhr, E. D., Singer, O., Meeker, K., Verma, I. M., Doyle, F. J., Takahashi, J. S. and Kay, S. A. [2007]: 'Intercellular Coupling Confers Robustness against Mutations in the SCN Circadian Clock Network', *Cell*, **129**, pp. 605–16.

Lopes, D. M. M. [2000]: 'What Is It like to See with Your Ears? The Representational Theory of Mind', *Philosophy and Phenomenological Research*, **60**, pp. 439–53.

Lowrey, P. L. and Takahashi, J. S. [2011]: 'Genetics of Circadian Rhythms in Mammalian Model Organisms', *Advances in Genetics*, **74**, pp. 175–230.

Macpherson, F. [2011]: 'Taxonomising the Senses', *Philosophical Studies*, **153**, pp. 123–42.

Malapani, C. and Fairhurst, S. [2002]: 'Scalar Timing in Animals and Humans', *Learning and Motivation*, **33**, pp. 156–76.

Maruyama, M., Hara, T., Katakura, M., Hashimoto, M., Haque, A., Li, G. and Shido, O. [2007]: 'Contribution of the Suprachiasmatic Nucleus to the Formation of a Time Memory for Heat Exposure in Rats', *ResearchGate*, **57**, pp. 107–14.

Matthen, M. [2014]: 'Active Perception and the Representation of Space', in D. Stokes, M. Matthen and S. Biggs (*eds*), *Perception and Its Modalities*, Oxford: Oxford University Press, pp. 44–72.

Matthen, M. [2015]: 'The Individuation of the Senses', in M. Matthen (*ed.*), *Oxford Handbook of the Philosophy of Perception*, Oxford: Oxford University Press, pp. 567–86.

Maudlin, T. [2002]: 'Remarks on the Passing of Time', *Proceedings of the Aristotelian Society*, **102**, pp. 237–52.

Meyer, U. [2013]: *The Nature of Time*, Oxford: Clarendon Press.

Moore, R. Y. [1997]: 'Circadian Rhythms: Basic Neurobiology and Clinical Applications', *Annual Review of Medicine*, **48**, pp. 253–66.

Mulder, C. K., Gerkema, M. P. and Van der Zee, E. A. [2013]: 'Circadian Clocks and Memory: Time-Place Learning', *Frontiers in Molecular Neuroscience*, **6**, available at <doi.org/10.3389/fnmol.2013.00008>.

Nanay, B. [2014]: 'Teleosemantics without Etiology', *Philosophy of Science*, **81**, pp. 798–810.

Newton-Smith, W. [1980]: *The Structure of Time*, London: Routledge and Kegan Paul.

Perry, J. [1979]: 'The Problem of the Essential Indexical', *Noûs*, **13**, pp. 3–21.

Phillips, I. [2008]: 'Perceiving Temporal Properties', *European Journal of Philosophy*, **18**, pp. 176–202.

Ralph, M. R., Foster, R. G., Davis, F. C. and Menaker, M. [1990]: 'Transplanted Suprachiasmatic Nucleus Determines Circadian Period', *Science*, **247**, pp. 975–8.

Rosenwasser, A. M., Pelchat, R. J. and Adler, N. T. [1984]: 'Memory for Feeding Time: Possible Dependence on Coupled Circadian Oscillators', *Physiology and Behavior*, **32**, pp. 25–30.

Ruby, N. F., Hwang, C. E., Wessells, C., Fernandez, F., Zhang, P., Sapolsky, R. and Heller, H. C. [2008]: 'Hippocampal-Dependent Learning Requires a Functional Circadian System', *Proceedings of the National Academy of Sciences*, **105**, pp. 15593–8.

Russell, B. [1905]: 'On Denoting', *Mind*, **14**, pp. 479–93.

Schibler, U., Ripperger, J. and Brown, S. A. [2003]: 'Peripheral Circadian Oscillators in Mammals: Time and Food', *Journal of Biological Rhythms*, **18**, pp. 250–60.

Silver, R., LeSauter, J., Tresco, P. A. and Lehman, M. N. [1996]: 'A Diffusible Coupling Signal from the Transplanted Suprachiasmatic Nucleus Controlling Circadian Locomotor Rhythms', *Nature*, **382**, pp. 810–13.

Smith, B. [1988]: 'On the Semantics of Clocks', in J. H. Fetzer (*ed.*), *Aspects of AI*, Dordrecht: Kluwer.

Sollars, P. J. and Pickard, G. E. [2015]: 'The Neurobiology of Circadian Rhythms', *The Psychiatric Clinics of North America*, **38**, pp. 645–65.

Stephan, F. K. [2002]: 'The "Other" Circadian System: Food as a Zeitgeber', *Journal of Biological Rhythms*, **17**, pp. 284–92.

Tal, E. [2013]: 'Old and New Problems in Philosophy of Measurement', *Philosophy Compass*, **8**, pp. 1159–73.

Tal, E. [2016]: 'Making Time: A Study in the Epistemology of Measurement', *British Journal for the Philosophy of Science*, **67**, pp. 297–335.

Vasalou, C. and Henson, M. A. [2010]: 'A Multiscale Model to Investigate Circadian Rhythmicity of Pacemaker Neurons in the Suprachiasmatic Nucleus', *PLoS Computational Biology*, **6**, pp. e1000706, available at <doi.org/10.1371/journal.pcbi.1000706>.

Wearden, J. H. [2001]: 'Internal Clocks and the Representation of Time', in C. Hoerl and T. McCormark (*eds*), *Time and Memory*, Oxford: Oxford University Press, pp. 37–58.

Weaver, D. R. [1998]: 'The Suprachiasmatic Nucleus: A 25-Year Retrospective', *Journal of Biological Rhythms*, **13**, pp. 100–12.

Welsh, D. K., Takahashi, J. S. and Kay, S. A. [2010]: 'Suprachiasmatic Nucleus: Cell Autonomy and Network Properties', *Annual Review of Physiology*, **72**, pp. 551–77.

Wittmann, M. and van Wassenhove, V. [2009]: 'The Experience of Time: Neural Mechanisms and the Interplay of Emotion, Cognition, and Embodiment', *Philosophical Transactions of the Royal Society of London B*, **364**, pp. 1809–13.

Yamaguchi, S., Isejima, H., Matsuo, T., Okura, R., Yagita, K., Kobayashi, M. and Okamura, H. [2003]: 'Synchronization of Cellular Clocks in the Suprachiasmatic Nucleus', *Science*, **302**, pp. 1408–12.

Ye, R., Selby, C. P., Chiou, Y.-Y., Ozkan-Dagliyan, I., Gaddameedhi, S. and Sancar, A. [2014]: 'Dual Modes of CLOCK: BMAL1 Inhibition Mediated by Cryptochrome and Period Proteins in the Mammalian Circadian Clock', *Genes and Development*, **28**, pp. 1989–98.

Zakay, D. and Block, R. A. [1997]: 'Temporal Cognition', *Current Directions in Psychological Science*, **6**, pp. 12–16.