

Tense and Relativity

Andrew Bacon
University of Southern California*

Abstract

Those inclined to positions in the philosophy of time that take tense seriously have typically assumed that not all regions of space-time are equal: one special region of space-time corresponds to what is presently happening. When combined with assumptions from modern physics this has the unsettling consequence that the shape of this favored region distinguishes people in certain places or people traveling at certain velocities. In this paper I shall attempt to avoid this result by developing a tensed picture of reality that is nonetheless consistent with ‘hypersurface egalitarianism’ – the view that all hypersurfaces are equal.

1 Temporalism

Call ‘modalism’ the thesis that there is genuine contingency: for some P , it’s contingent whether P . Although there were initially some Quinean reservations about this thesis, especially during the early days of modal logic, modalism is now widely accepted and the practice of employing sentential modal operators to formalise modal claims rarely raises eyebrows any more. Since there are many parallels between modality and time – the adverbs ‘always’ and ‘sometimes’ seem to behave logically much like the modal counterparts ‘necessarily’ and ‘possibly’ – it is natural to wonder about the status of the counterpart claim that some things aren’t always true or always false. Call this ‘temporalism’. In higher orderese:¹

TEMPORALISM: For some P , it’s not always the case that P and not always the case that not P .

*Thanks to Cian Dorr and Jeremy Goodman for discussion and for sending me detailed comments on an earlier draft, and to an anonymous referee for several helpful suggestions.

¹Note that sometimes eternalism and temporalism are formulated in terms of first-order quantification over propositions: there are propositions which aren’t always true and aren’t always untrue. This formulation does not assume the coherence of this bit of higher orderese, however it raises distracting questions about the status of the monadic propositional truth predicate required to state it, which I would rather not address. Temporalism, as I have formulated it, circumvents this entirely since it does not favour a particular alethic ideology.

Temporalism is to be contrasted with its negation, eternalism: for all P , either it's always the case that P or it's always the case that not P . Although eternalists appear to be denying the existence of change, they tend to make sense of tense in English in other ways. While tense operators are everywhere redundant on this view², most actual eternalists maintain that English expressions like 'always' and 'sometimes' can be treated in terms of quantification over times instead.³

In recent debates about modality a lot of attention has been directed at the thesis of necessitism – that necessarily, everything is necessarily something (see Williamson [35]). There is a straightforwardly parallel debate about tense, which Williamson labels 'permanentism':⁴

PERMANENTISM: Always everything is always something.

Note that eternalism entails permanentism since, given that tense operators are redundant, it amounts to the vacuous claim that everything is something. To find an interesting thesis in permanentism one already has to accept temporalism. Permanentism can be contrasted with its negation, the thesis that sometimes something is sometimes nothing.⁵ A more specific example of a non-permanentist view is Prior's presentism.

Necessitism naturally brings to mind David Lewis's modal realism. Williamson's version of necessitism is distinctive, however: according to Lewis if there could have been a talking donkey, then there *is* a talking donkey, provided we take care to quantify unrestrictedly. For Williamson on the other hand, the possibility of a talking donkey entails a weaker thing: the actual existence of something that could have been a talking donkey. The possible talking donkey won't in general be a talking donkey, according to Williamson; indeed, it typically won't even be a concrete object.

There are temporal analogues to the Lewisian picture. For example, Ted Sider describes the 'block universe' account of time as one in which 'dinosaurs, computers and future human outposts on Mars are all equally real' (Sider [29], p11). The Williamsonian version of permanentism, by contrast, would deny the existence of dinosaurs, although would accept the existence of things that used to be dinosaurs.

Both Lewis and Sider appear to be describing alternative forms of necessitism and permanentism. Indeed, I think it is most natural to interpret both Lewis and Sider as rejecting modalism and temporalism respectively, thereby ac-

²Given the assumption that 'always' is factive, one can easily prove from eternalism that tense operators are redundant in the sense that 'sometimes P ', 'always P ' and P are equivalent.

³Note that such eternalists will still accept the rendition of eternalism in English – 'for all P , either it's always the case that P or it's always the case that not P ' – because either the first or second disjunct is a vacuously true quantified statement (for in either case one is quantifying into a formula that has no free variables).

⁴This debate is sometimes cast in terms of the Barcan formula and its converse: always everything ϕ s if and only if everything always ϕ s. See Barcan [2].

⁵Williamson calls this 'temporaryism'. In what follows, however, I shall attempt to avoid this terminology due to its potential to be confused with the thesis of temporalism outlined above.

cepting necessitism and permanentism in their vacuous forms (in stark contrast to Williamson).⁶

Cian Dorr has recently offered a more transparent way of cashing out the Lewisian picture in a way that accepts modalism (Dorr [7]). Consider an ordinary, non-talking, donkey, Muriel. According to the Dorrian version of modal realism, Muriel could have been a talking donkey because there in fact exists a talking donkey, and Muriel could have occupied the qualitative role that talking donkey in fact plays. All possibilities, on the Dorrian picture, are qualitatively indistinguishable multiverses. The differences between multiverses consist solely in which individuals occupy which qualitative roles. For example if one multiverse contains a farmer and a ballet dancer satisfying a certain qualitative description, then they all do. However they might still differ – I might be the farmer in one possibility and the ballet dancer in the other. This is clearly a necessitist picture, but one dramatically different from Williamson’s.

The same way of thinking can be applied to time.⁷ If there’s some object, x , which used to be a dinosaur it’s because there really is a dinosaur distinct from x , but whose qualitative role x used to play. The resulting picture is one in which no change is qualitative: there’s a four dimensional manifold with dinosaurs, computers and human outposts on mars dotted about all over it. The qualitative distribution of properties will remain constant, as time progresses, but which things occupy which qualitative roles changes. To rule out this form of permanentism we could adopt the following strengthening of temporalism:

QUALITATIVE TEMPORALISM: For some P such that it’s qualitative whether P , it’s not always the case that P and it’s not always the case that not P .

There is of course an issue here as to whether one can make sense of the operator locution ‘it’s qualitative whether’; I think one can, but it is not something I shall defend here.⁸ One can contrast this with the thesis of qualitative eternalism: For any P , if it’s qualitative whether P , then either it’s always the case that P or it’s always the case that not P .

I am inclined to think that qualitative eternalism is not particularly plausible. According to it, for example, some Roman general is crossing a river right now (provided we are quantifying unrestrictedly). Since this is a qualitative proposition that was once the case, it is always the case. Of course, this fact is not true because Julius Caesar is crossing the Rubicon, it is true rather because someone whose qualitative role Caesar used to play is crossing some river whose qualitative role the Rubicon used to play. But to my ear this explanation does nothing to make the initial claim sound any less surprising. You would not even find a propositional eternalist uttering this sentence, for they typically treat

⁶Of course, this is not to say that these views cannot capture the truth of English sentences using modal and temporal words such as ‘can’, ‘sometimes’, and so on. Such words will be given an analysis in which they quantify over worlds or times instead, and in which most expressions possess a covert world and time variable.

⁷This idea is explored in Dorr [6].

⁸This formulation of qualitative temporalism is my own: Dorr’s theory can be formulated using counterpart theory and does not depend on the coherence of the qualitiveness operator.

tenses very differently.⁹

A final question we might ask ourselves is how to think about the persistence of objects in this picture. According to a perdurantist picture, for example, ordinary objects are located at four-dimensional ‘worm-like’ regions of space-time, although have instantaneous temporal parts that are located at each temporal cross section of their locations. Another extremely natural position is that ordinary objects are instead located at constantly changing three-dimensional regions of space time. This latter position seems more natural to me, so I shall adopt it when an answer to this question is needed. There are many other possibilities as well and little I say in the following will turn on this dispute, so I will set it aside for the most part.

It is my view that the combination of permanentism and qualitative temporalism is the best version of temporalism on the market and deserves our considered attention.¹⁰ However the strongest objections to it, in my opinion, are that it does not appear to be consistent with modern physics,¹¹ prompting the recent suggestion by Dorr that the best way to reconcile temporalism with space-time physics is to reject qualitative temporalism.

To my knowledge all existing versions of qualitative temporalism are committed to the view that some regions of space-time are (in some sense) more special than others. That is, these temporalists reject what I’ll call:

HYPERSURFACE EGALITARIANISM: All regions of space-time are metaphysically on a par (at the fundamental level or otherwise).

Assuming classical physics the denial of hypersurface egalitarianism usually means that some particular simultaneity slice of the space-time is distinguished from the rest. This difference may manifest itself as a fundamental property possessed by only one hypersurface (as in Deasy [5]), but I will also be interested in views in which the difference consists in a distinguishing non-fundamental property.¹² On either picture, objects entirely in our future and past are not on a par with present objects – the ancient Greeks for example, have a lesser status, metaphysically speaking, than we do.

When we move to relativistic physics, matters change considerably: the denial of HYPERSURFACE EGALITARIANISM forces one to make the surprising claim that not only are not all times on a par, but that not all *places* are on a par

⁹In addition to the above, the view has other surprising consequences. For example, there must be continuum many people in order for there to be a person whose qualitative role I will occupy in x minutes time for each x in some interval. There is also a puzzle about whose qualitative role I’ll occupy after I die (presumably I won’t even occupy the role of a person after I’m dead); a similar puzzle also applies to Williamson.

¹⁰The view is explored, for example, by Williamson [35] §8.4, Sullivan [34], and Deasy [5]. It is sometimes associated with the ‘moving spotlight’ theory, although as Williamson notes, this term has other baggage that I won’t be as concerned with here.

¹¹See Sider [29] §2.4, for example. See also Rietdijk [27], and Putnum [26].

¹²Perhaps one is a temporalist but maintains that while tense operators are conceptually basic and cannot be reduced to tenseless notions, they are not metaphysically fundamental. On this picture there may be ways to single out a special hypersurface using non-fundamental tense-theoretic language without distinguishing any hypersurface from a fundamental perspective.

either, or that people travelling in certain *directions* and at certain *speeds* have a more accurate view of reality than others. (On these views the special region of space-time is usually a three-dimensional hypersurface of some sort, although it could be other shapes as well. On some theories it is a space-time point, in others it is ‘bowtie’ shaped, for example.)¹³ There are some striking analogies with the pre-Copernican ‘geocentric’ conception of the universe, but in this case it is potentially more radical: the Ptolemaic system has no implications about the *metaphysical* status of people located differently from us (for example). Since people in the U.S. are travelling at different velocities due to the rotation of the Earth, and are at different places from people in Australia, the hypersurface prioritarian must demote the metaphysical status of at least one of these people to that of whatever lesser status the Ancient Greeks possess in the classical variant.

One version of hypersurface egalitarianism is Dorr’s version of temporalism discussed above. On that view there is no change concerning which hypersurface is present — the only change consists in which physical objects occupy which qualitative roles. However, as of writing, there is no official answer to the question of which objects will occupy given qualitative roles n minutes from now in the relativistic setting.¹⁴ In this respect, at least, the relativistic version of the theory is incomplete. In this paper, I describe a version of *qualitative* temporalism, in both classical and relativistic settings, that accepts HYPERSURFACE EGALITARIANISM and thus, like the eternalist, treats all regions of space-time as on a par fundamentally (or otherwise). It is possible that some of the ideas of sections 3 and 4, involving time-shifted properties, could also be of help in fleshing out a non-qualitative temporalist picture, although it is beyond the scope of this paper to explore that idea here.

Before we move on a point of terminology is in order. It will be convenient to use the adverbial phrase ‘sometimes’ in place of the slightly more awkward ‘not always not’. One must bear in mind that despite its typographical form, this is no more a quantificational phrase than ‘always’ is (it is not semantically decomposable into ‘some’ and ‘times’). ‘Sometimes P ’ also has an unfortunate presupposition, which I do not always want enforced, that it’s been the case that P on two or more occasions. Despite the occasional awkwardness, I shall simply use it synonymously with ‘not always not’ (which has no such presupposition).

¹³See Zimmerman [36] for a recent example of the three-dimensional hypersurface view. The view that the present is just a space-time point is sometimes called ‘Here-nowism’ (see Hinchliff [14], Savitt [28], Sklar [31], and Stein [33], Skow [32]). See Sider [29] §2.4 for critical discussion of these possibilities.

¹⁴Many natural candidate answers give odd results. For example, one candidate answer is that I will take the role of the slice of my space-time worm that lies n minutes of proper time in the future of me (where that length of time is calculated by integrating proper time along the world-line I trace out in space-time). But this has the unfortunate result that people will become ‘out of sync’: if Alice and Bob meet today, and Alice does a lot of travelling, then when their space-time worms reconverge at a later time, Alice will be in the future light cone of Bob and Bob will only meet a counterpart of Alice whose qualitative role Alice used to occupy. Thanks to Jeremy Goodman for discussion here.

2 Temporalism in Classical and Relativistic Physics

Temporalism need not be something that is motivated by metaphysical considerations. It is something one might be driven to on the grounds that there are things we are ignorant about that cannot be traced to ignorance about the realm of eternal truths. Or perhaps there are things one can care about or be relieved about, albeit things that are essentially tensed, about the future or past (Prior [23]). Temporalism looks as though it offers us a way of accounting for these attitudes by offering us more contents for these attitudes to have. It is for this reason that the dispute over temporalism makes its appearance in the philosophy of language more often than it does in metaphysics (Kaplan [16], Montague [21], Brogaard [3]).

We might contrast temporalism with an apparently stronger position – the A-theory – which some take to be more metaphysically substantive version of temporalism. According to A-theorists the present is metaphysically distinguished. This might mean it is distinguished ontologically, as the presentist maintains, or perhaps it is special in some other way (as, for example, with the ‘moving spotlight theory’). The A-theory is usually characterised by the claim that there are tensed facts. Note, however, if all one means by a ‘fact’ is a true proposition, then temporalism already entails that there are tensed facts simply because there are true propositions that aren’t always true.¹⁵

There are stronger forms of the A-theory that do go significantly beyond temporalism. These views often conjoin to temporalism either (a) the thesis that tense operators are fundamental¹⁶, or (b) that there are fundamental tensed facts – fundamental truths that aren’t always true. These stronger versions of the A-theory do seem to make a fairly strong metaphysical distinction between times: of the things that are going on now and at other times, the things going on now are more fundamental. I prefer the weaker version of the A-theory. Regarding (a), perhaps tense operators are explainable in other terms – or, more likely in my opinion, they are conceptually basic. But either way they are not part of the fundamental furniture of the universe. Similarly, regarding (b) I’m inclined to think that all fundamental truths are eternal truths.¹⁷ Truths about the field values at particular space-time points, for example, are truths

¹⁵Temporalism says that there are propositions that are sometimes but not always true: if p is such a proposition, then either it or its negation is a temporary truth and thereby a tensed fact.

¹⁶Perhaps the tense operators are ‘structural’ in the sense of Sider [30]. See, for example, the characterisation of the A-theory in Sullivan [34] p150-152 (see also Deasy [5] for more discussion).

¹⁷This is one way in which I differ from Dorr [6]. Dorr (pc.) has expressed sympathy to the idea that not all fundamental truths are eternal. Note, moreover, that this plausibly follows from the principle that whatever is sometimes true is possibly true [8], a principle which Dorr accepts. Let F be the conjunction all fundamental truths. Suppose, for contradiction, that F is always true. Since it’s sometimes Tuesday and sometimes Wednesday, it follows that: sometimes (F and it’s Tuesday) and sometimes (F and it’s Wednesday). Given that whatever is sometimes true is possible, it follows that it’s possible that (F and it’s Tuesday) and possible that (F and it’s Wednesday), which contradicts the plausible thought that everything (including whether it’s Tuesday) supervenes on the fundamental. So F is not always true.

that are always true if ever true, whereas truths about who is presently sitting, and other non-fundamental truths, can change. According to this moderate version of the A-theory there is nothing distinctive about any particular slice of space-time, at least from the perspective of the fundamental facts.¹⁸

It is often argued that the A-theory has a special problem accommodating relativistic physics. One might have thought that by adopting a less metaphysically substantive version – temporalism – the issues in this vicinity are lessened. In my view it matters very little whether we adopt the stronger thesis about the fundamentality of tense or not: the idea that no time is distinguished from the others by a fundamental property is consistent with the idea that the present time is still distinguished in a non-fundamental way. Indeed the problem from space-time physics does not turn on the difference between the two ways the present could be distinguished: temporalism is therefore in just as much trouble as the stronger A-theory is.

The problem from space-time physics, in rough outline, is that it seems as though temporalism singles out one hypersurface as distinguished (in either a fundamental or non-fundamental way) which makes for trouble when combined with certain facts about the structure of these hypersurfaces that we get from physics (see, for example, Sider [29] Ch2§4.) The common way to run this problem is to note that if you're a temporalist there's a special time that is accurate:

ACCURACY: A time t is accurate iff, for any P , P if and only if at t , P .

Given some natural assumptions one can prove that no more than one time can be accurate.¹⁹ Now of course, it may turn out that according to some versions of temporalism being accurate isn't as metaphysically weighty as it would be according to a stronger version of the A-theory. However, as I already mentioned, the problems do not turn on any kind of judgment about the metaphysical weightiness of accuracy.

That said, there is a more serious problem with the above formulation of the objection, even at this embryonic stage. It is absolutely essential, in what

¹⁸There is a putative puzzle for moderate temporalism that for reasons of space I can only touch on briefly here. It is extremely natural to think that facts about who is presently sitting supervene on fundamental facts, like the values of fields at particular space-time points. However given what we have just said, the former kind of facts change, whereas the latter do not. The way out, I think, is not to relinquish the supervenience principle, but to allow the supervenience facts to change: whether I am sitting presently supervenes on the field values at a certain region of space-time, but tomorrow will supervene on the field values at a different region of space-time temporally separated from the original. The resulting view is one in which what's necessary is a temporary matter, and necessary truths can be sometimes false. Indeed, the existence of temporary necessities is a prediction of the standard approach to the logic of modality and tense, developed by people like Kit Fine [10] and David Kaplan [15]. Although Dorr and Goodman [8] have recently drawn attention to this somewhat surprising consequence of the standard view, and have recommended an alternative modal-tense logic, the orthodox view strikes me as the most natural way to be a moderate temporalist of the sort I have described, Dorr and Goodman's arguments notwithstanding.

¹⁹This follows from the assumption that no two distinct times make exactly the same propositions true. This assumption drops out of the view, discussed later, that times are just certain kinds of propositions, for example.

follows, that we distinguish two related theoretical entities often referred to as ‘times’. One of these is a sort of physical entity which arises out of modern physics: a certain kind of three dimensional region of space-time. The other is a certain kind of theoretical object that is fundamentally defined by its role in relation to the tense operators: for example, part of that role is that always P iff for any time, t , at t , P . It could turn out that these entities and the physical entities end up being the same thing. However it is always possible to reconstruct talk involving the tense-logical entities in terms of certain kinds of mutually exclusive propositions, and of the present time as the proposition that is presently true (see, e.g., chapter 10 in Prior [25], and Fine [10]).²⁰ I shall reserve the word ‘hypersurface’ – the higher dimensional analogue of a surface – for the physical entity, and ‘time’ for the tense-logical entity. The relation between the two is fraught.

Now, although it is built into the definition of a time, as we are using the word, that we can ask whether a proposition is true at it, there is no obvious sense to be made of truth ‘at a hypersurface’. Of course, one can trick oneself into thinking one has a feel for the notion, but it quickly becomes harder to make sense of when we take into account the abundance of tensed propositions that one would have to think of as true at a hypersurface (discussed in section 3) and is particularly unfamiliar when applied to the kinds of hypersurfaces you find in the relativistic setting.

The following is therefore an attempt to formulate the objection in more detail, without conflating times with hypersurfaces. The problem arises, in its most acute form, when we move from a non-relativistic theory, such as Newtonian physics, to a relativistic one. I shall focus here on special relativity because it exhibits most of the fundamental problems with reconciling temporalism with relativistic theories; similar issues arise in the context of the general theory of relativity.

How, then, do tensed propositions, such as the proposition that Ellen is sitting, relate to the eternal facts about space-time, hypersurfaces, and their properties? To get a clearer handle on this question it will be handy to introduce the notion of a *space-time event* as it is used by physicists: these are effectively space-time points – things that are both point-like and exist instantaneously. These things are not quite like events in the ordinary sense – ordinary events usually have non-zero duration and extension, whereas space-time events stop occurring the moment they start (they are instantaneous) and don’t take place in an extended region of space-time. More strikingly, we can talk of space-time events in the physicists’ sense occurring in empty regions of space-time. However despite these superficial differences they seem like a natural generalisation of the ordinary notion of an event – a space-time event is an extremely brief, extremely small event, that can occur even in a region of space-in which nothing seems to be happening. It also seems that we can talk, as I just did, of space-time

²⁰Fine defines a time as a proposition that is maximally strong among the sometimes necessary propositions. This definition is sensible only under some assumptions about the interaction of the modal and tense operators (see the comment in footnote 18 relating to Dorr and Goodman’s [8].)

events occurring. After all, we can talk of events in the ordinary sense, which can have arbitrarily small duration and extension, occurring: our notion of occurrence for durationless point-like events is, again, just a straightforward generalisation of that notion.²¹ Note that insofar as space-time events are very short and small kinds of events in the ordinary sense, which space-time events are occurring is a constantly changing matter. Clearly any space-time event will occur at sometime or other, but unlike ordinary events, cannot occur more than once since they are instantaneous.²² Letting E be the proposition that some particular event, e , is occurring this amounts to the following in the tense logical language:

Sometimes E .

For any P such that sometimes P , if it's always the case that if P then E , then it's always the case that if E then P .

The latter claim encodes the idea that e is instantaneous: informally, if e had occurred over more than one occasion we could let P be the proposition that e was occurring on a particular one of these occasions. Then it would always be the case that if P then E , but sometimes the case that E and not P , contradicting the principle in question.²³

From these assumptions we can partition space-time into non-overlapping regions under the (eternal) equivalence relation between events: sometimes both e_1 and e_2 occur.²⁴ Note furthermore that our partition over space-time has a natural temporal order. It is a theorem of standard tense logic that:

If sometimes P and sometimes Q then either sometimes [P and Q], or sometimes [P and it was the case that Q] or sometimes [Q and it was the case that P].

²¹We can get a handle on the notion of an unoccupied space-time event occurring by asking ourselves whether the physical event e would have been occurring, had that space-time event been occupied by e .

²²One might think that I am here running together duration (a tense logical notion) and spatiotemporal extension (a physical notion) in exactly the way I cautioned against above. However I am making this judgment on the empirical basis that the duration of an ordinary event and its spatiotemporal extension are always proportional to one another; similar parallels can thus be drawn in the limiting case of space-time events.

²³Although the idea that space-time events occur at most once seems extremely natural, it's important to note that there are temporalist views that deny this and maintain that space-time events can occur for extended durations. One version of the Prior-Goldblatt view can be understood in this way (see Prior [24] p203 and Golblatt [12]: an event occurs for as long as it is space-like separated from us. (Although one could describe another version of the Prior-Goldblatt view in which must predicates, like 'occurs', are context sensitive.)

²⁴This relation is obviously reflexive and symmetric (reflexivity requires that every event occurs sometimes). To show that it's transitive suppose that sometimes e_1 and e_2 occur, and that sometimes e_2 and e_3 occur. Since it's always the case that if e_1 and e_2 occur e_1 occurs, it follows by the second principle that, always, if e_1 occurs, both e_1 and e_2 occur. By symmetrical reasoning it's always the case that if e_2 occurs, both e_2 and e_3 occur. Putting these together we know that always if e_1 occurs, e_3 occurs, and since e_1 sometimes occurs both e_1 and e_3 sometimes occur. The equivalence relation is moreover eternal provided we assume that 'always' satisfies the characteristic principle of S5.

We can say that an equivalence class comes before another if there is an event in each, e_1 and e_2 , such that sometimes [e_2 occurs while it was the case that e_1 occurred]. Assuming the above principle this linearly orders our partition of space-time.²⁵ Using tense logic alone, and the notion of occurrence for space-time events we can tell a considerable amount about the temporal structure of our space-time manifold.

Note that these observations appear to give one region of space-time in particular a privileged status: the fusion of all events that are occurring. What shape is this distinguished region? If one is assuming a classical Newtonian picture, for example, there's a unique way of carving up space-time into a so-called 'foliation' of three-dimensional hypersurfaces, with events located within the same hypersurface to be thought of as occurring simultaneously. It's natural to think that in a Newtonian universe the partition of space-time alluded to earlier will be a foliation, and that the distinguished region will be a single three dimensional hypersurface from this foliation. Indeed, one can prove this from some empirical observations about the connection between event-occurrence and space-time structure in a Newtonian universe:²⁶

Always, if e_1 is occurring then: e_2 is occurring iff e_2 lies on the same surface of simultaneity as e_1 .

The notion of simultaneity here is an eternal relation between space-time events that falls directly out of the geometry of classical space-time and is not a tense-logical notion.

It is tempting to find the fact that a unique simultaneity slice is privileged somewhat surprising since the fundamental facts, according to the Newtonian picture, do not single out any particular hypersurface as special, and it is unclear how a non-fundamental distinction between hypersurfaces could emerge in this story either. One picture, which would need some more fleshing out, would be that the question of which hypersurface is the special occurrent hypersurface is one whose answer supervenes on the fundamental facts. But since the occurrent hypersurface changes – whichever hypersurface is occurring now won't be tomorrow, for example – the supervenience must be a temporary one, singling out different hypersurfaces as time progresses.²⁷

In special relativity the matter is complicated further, for in that case not only does the fundamental geometry fail to single out a special hypersurface, but

²⁵Of course it is consistent with what I said that this partition (which by permanentism exists eternally) is always linearly ordered by this relation, but that the specific way in which the partition is linearly ordered changes over time. If we also assume the S5 principle for 'always', a completely standard assumption, we can argue that this ordering will never change: if sometimes e_1 occurs while it was the case that e_2 occurred then it's always the case that sometimes e_1 occurs while it was the case that e_2 occurred.

²⁶Note: since we do not in fact live in a Newtonian universe, we do not make these sorts of empirical observations. They only seem to be approximately true when low relative velocities are involved.

²⁷Of course, this view posits the existence of temporary necessities, and so is not entirely uncontroversial. For more discussion see the treatment in Dorr and Goodman [8] mentioned also in footnote 18.

it doesn't even give us a unique way of foliating space-time. This turns on the fact that in special relativity the fundamental geometry of space-time doesn't give us any way to make sense of the temporal separation between the occurrence of two events – the notion of the amount of time elapsed between one event and another isn't something one can straightforwardly read off the geometry of space-time (compare this with the Newtonian case). The geometrical facts don't even tell us whether the temporal separation between two events is 0, i.e., whether the events are simultaneous with one another. The geometry only tells us what the *spatio-temporal* separation between two events is.

But surely, one might reason, ordinary talk of events spatially separated from us occurring still makes sense even if it is not something that falls directly out of the fundamental geometry of space-time. There are, after all, many ordinary concepts that cannot simply be reduced to the geometrical properties of space-time – why should occurring be any different? Perhaps, then, we can still talk about the totality of events that are occurring right now, even if only in a non-fundamental way.

There are actually lots of ways one could attempt to reconcile these kinds of tensed concepts with special relativity.²⁸ Let me briefly describe the simplest, and the one that most people have in mind when they talk about the A-theory in the context of special relativity. Recall that if we assume a standard tense logic one can divide space-time into a linearly ordered partition of non-overlapping regions. Natural partitions of this sort fall out of the geometry of space-time as well in the form of certain kinds of hypersurfaces, called planes of simultaneity. These things have two notable properties, either of which define them: (i) they are 'geodesically generated' surfaces – three dimensional surfaces that consist of all the straight-lines passing orthogonally through some time-like vector and (ii) they are orthogonal to some timelike Killing field.²⁹ A Killing field is an assignment of vectors to space-time points with the property that any two entities with trajectories 'following' the vectors will always be the same space-time distance apart from each other. Each one of these surfaces can be given the flat Euclidean geometry we're familiar with from non-relativistic physics. Foliations are simply partitions of space-time into these kinds of hypersurfaces.

Could the partition of space-time generated by the tensed logical equivalence relation describe above be one of these foliations? Imagine an inertially moving

²⁸See for example Zimmerman [36], Müller [22], Prior [24], Goldblatt [12], Skow [32].

²⁹Both of these definitions of a plane of simultaneity are chosen so that they generalise straightforwardly to the setting of general relativity. However although these two definitions are equivalent in a flat Minkowski space-time, when we move to the context of general relativity they can come apart (one can find manifolds that do not admit a timelike Killing field, and not every time-like Killing field is orthogonal to a spacelike hypersurface). Note that in the context of general relativity people often invoke the notion of a Cauchy surface: a space-like surface that intersects every time-like or null curve exactly once. However Cauchy surfaces don't behave much like planes of simultaneity: even in the context of special relativity, with a flat Minkowski space-time, there are Cauchy surfaces that aren't planes of simultaneity. (It also turns out that there are some deviant space-times – space-times that fail to be 'globally hyperbolic' – that don't have any Cauchy surfaces, although these are much less common and effectively allow for time travel (see, for example, the discussion of Gödel universes in Malament [20], Gödel [11]).)

measuring device that emits a beam of light which is then reflected back from some distant mirror. The spatio temporal location where the beam first hits the mirror is an event (or near enough), e . We can tell whether the event e is occurring in terms of observable (temporal) truths about the measuring device, which is in spatial proximity to us: if for some ϵ it was the case ϵ seconds ago that the device emitted the beam and will be the case ϵ seconds from now that it will receive the beam then e is occurring. One can similarly relate the proposition that e will occur and has occurred to facts about what will and has happened to the device; if for example, there exist $\epsilon > \delta$ such that it was the case ϵ seconds ago that the device emitted the beam and will be the case, in δ seconds, that it will receive the beam, then e has occurred. Assume that our clocks allow us to accurately measure facts about what has and will happen to the measuring device in ϵ seconds, and that the observations made by our measuring device indicate which events are occurring. Then, as a matter of empirical fact, the observations we make with such devices are consistent only with the thesis that the fusion of events occurring right now is a plane of simultaneity, and that the partition generated by the equivalence relation ‘sometimes both e_1 is occurring and e_2 is occurring’ will be a foliation of space-time into planes of simultaneity.³⁰

Although this result might at first seem to vindicate the compatibility of temporalism with relativity, things aren’t quite that simple. Firstly, unlike in a classical space-time, the geometry of special relativity permits many distinct ways of partitioning space-time into a foliation of planes of simultaneity. Moreover, it turns out that the process described above for empirically determining which events are simultaneous will deliver one of these alternative foliations if performed by another measuring device that is moving inertially relative to the original. The result is that the foliation of space-time into classes of seemingly simultaneous events picked out by our observations depends essentially on the observers trajectory. Indeed, it turns out that any foliation of space-time could be observed to be ‘correct’ by some inertially moving observer: so in some sense all of the foliations of space-time are on a metaphysical par. But we also know that there is a unique way of partitioning space-time by a tense-logical relation defined in terms of event occurrence, so there is a unique foliation corresponding to our tensed talk (if this partition corresponds to a foliation at all).

So far we have been focussing only on questions concerning the occurrence of events, and we saw that the events presently occurring determine the hypersurface that is ‘present’. However presumably everyday properties, like whether someone is presently sitting, also depend on the local properties of the present hypersurface. It is natural to think, for example, that whether a person’s temporal worm – the fusion of space-time points they at sometime occupy – intersects the present hypersurface in a sitting shape determines whether or not they are presently sitting. What goes for sitting presumably goes for many other things that matter to us; whether people are happy or sad, whether they are popular or rich, alive or dead, and so on. The question of which hypersurface is presently

³⁰See Malamant [19] §3.1 for more discussion of this procedure (Handbook of the Philosophy of Physics). Note that this procedure breaks down in the context of general relativity.

occurring is therefore of greater significance than documenting which space-time events are occurring.

Observe also that if there is a special foliation, associated with a special inertial trajectory, it is extremely unlikely that it is a trajectory we are presently on. Remember that we are constantly moving – the earth is rotating, and it is moving around the sun. Even if by chance we at some time did happen to be moving along the special inertial trajectory, it would only be momentary. The chances are, in fact, that the special trajectory, whatever it is, is moving relative to us at quite a speed. If this is so, then even things close to us that appear to be sitting might not, in fact, be sitting!³¹ This, I take it, is unacceptable; it would entail widespread error.

There is another version of TEMPORALISM that has become much more popular recently that one might hope to appeal to at this juncture: relativism (see MacFarlane [18] chapter 11, Pinillos [1]). The version of relativism subscribing to the thesis of temporalism must accept all the consequences we have drawn from temporalism so far. However, if I were a relativist I would typically go on to make two further claims. The first is that, by some miraculous coincidence, the preferred foliation is always perpendicular to *my* trajectory *right now* (that's right: *I*, A. Bacon, play a special role in determining how time flows. This conclusion is arrived at just by applying the relativist semantics). The second is that the norms of assertion are different: one doesn't aim to assert truths, one only aims to assert things that are true relative to your own perspective, as determined by your trajectory. The second point is there to try and make the first sound less absurd: although everyone other than me is potentially going about asserting false things, it turns out there's nothing wrong with asserting false things if you're moving relative to me. These maneuvers notwithstanding, it is very hard to take the original assertion as anything less than absurd: even if everyone else can be pardoned for asserting falsehoods, it's the thought that I, in particular, am the only person asserting truths that is so hard to swallow.³²

Before we move on, it is worth noting why the problem of widespread false belief does not arise in the Newtonian case. It's certainly true that Socrates exists, by permanentism, and that relative to his time he thought that he was alive. But of course, Socrates is not alive. But we cannot infer from this that there are loads of people, like Socrates, who think that they are alive when they aren't. Socrates, for example, doesn't think he is alive, he *thought* (correctly) that he *was* alive. Socrates isn't a person or a thinker – these are tensed notions

³¹Note that in the setting of special relativity there's a fairly natural measure we can put on the class of possible foliations of space-time into planes of simultaneity (one that derives from the geometry of the space-time, and is independent of the choice of coordinate system). Given an point p , and any finite spatially separated time-like trajectory t that does not intersect the lightcone through p (representing, e.g., the period of my sitting) the measure of foliations that have a plane of simultaneity that intersects both p and t is 0.

³²Relativists will typically point out that I'm not special in an interesting sense, because everyone is special relative to their own perspective. This strikes me as a distraction: in addition to all the relativised properties – of being an asserter of truths relative to p , for each perspective p – there's also the property of being an asserter of truths simpliciter which only I have. See Cappelen and Hawthorne [4] for more discussion of these kinds of worries.

– he *was* a person and *was* a thinker. Things existing entirely before or after the present hypersurface simply aren't the kinds of things that can have thoughts (they are just things that used to have thoughts, or will have thoughts).

Might the analogous move be made in the relativistic case? This would involve saying that things spatially separated from us and moving relative to us aren't the kinds of things that can have thoughts or beliefs. This is patently absurd: most people are moving relative to me, yet it clearly doesn't follow that only a handful of people have thoughts or beliefs. It seems clear that this type of response does not extend to solve our problems with relativity.

3 Timeshifted Properties in Classical Physics

In the last section I gave an argument that temporalism appears to do two things: (i) single out as special a preferred foliation of space-time (under the equivalence relation: sometimes e_1 and e_2 occur) and (ii) single out a preferred hypersurface from that foliation as the hypersurface corresponding to the present time (the fusion of the events that are presently occurring). In the next sections I shall argue that both of these results are slightly misleading: I will attempt to flesh out the temporalist picture in way that makes it clear that the tensed facts puts each foliation and hypersurface on a par with one another. I'll start, in this section, with the simpler of the two points: that the tensed facts, in a clear and natural sense, do not really single out a hypersurface as special. This point holds even in the non-relativistic case so it makes sense to see what happens to our argument in this context. Indeed, it is more straightforward, pedagogically, to start with that case before moving on to the more complicated issue of foliations in the relativistic setting.

The key concept here will that of a time-shifted property. Let me start with some examples. The property of sitting is a temporal property – Ellen, for example, might have it sometimes but not always. Perhaps she was sitting five minutes ago, but not now. However, although she does not have the property of sitting, she does have another temporal property, which we'd express with the predicate 'was sitting five minutes ago'. This is also a property she has sometimes but not always: she has it now, but she won't in five minutes. The length and order in which one switches between having this property and not having it is exactly parallel to that of the property of sitting, but for the difference of a five minute delay.³³

Let us provisionally adopt the notation '–5-sitting' to denote the latter property and '0-sitting' to denote the former. It is clear how to generalise this; we can talk about people x -yawning, x -talking, x -jumping, and so on, for any positive or negative number x , and this just means the person was or will be yawning, talking, et cetera in x minutes time (I am counting in minutes, but

³³Formally one needs a slightly richer tense logical framework to define these properties – one needs metric tense operators in addition to the ordinary Priorian operators – but it is clear which properties I'm talking about, whether we understand them in terms of special tense operators or not.

nothing important will turn on this).

It is important not to confuse time-shifted properties with time-indexed properties. A time indexed-property is a property like the property of sitting at 5:00pm GMT on February the 22nd, 2015. Time-indexed properties are eternal – if you ever have the property of sitting at that particular date and time you always have it. Time-shifted properties, on the other hand, are temporary – Ellen has -5 -sitting now, but won't in a few minutes.

The second thing to note is that I express the property of 0-sitting in English with the expression 'is sitting'. However to express the property of -5 -sitting in English I have to use a more complicated construction. One might take this observation to be indicative of a more fundamental asymmetry between the two properties. Perhaps it suggests that the property of 0-sitting is somehow more basic than that of -5 -sitting – that one gets the latter property from the former by some tense-logical operation, but not vice versa. This could of course be true, but if it is, it is not entailed by the thesis TEMPORALISM that I am exploring or any of the other principles discussed in section 1. One could, for example, adopt the alternative hypothesis that -5 -sitting is more basic, and that one only obtains the property of 0-sitting from it by applying some tense-logical operator. Or one could adopt the hypothesis that all of these temporal properties are on a par with one another, and that tense-logical operations are simply ways of converting between these different properties.³⁴

Let us explore this last thought some more. It is obvious that our present notation subtly steers us toward the asymmetric interpretation, according to which the property of 0-sitting is more fundamental. However the only reason we indexed this property with the number 0 was because it happened to be expressed by the simple expression 'is sitting' and not the more complicated ones needed to express the other properties; we could just as well have centered our numbering system around the property of -5 -sitting instead: all the relevant structural features would remain intact. An even better strategy would be not to use real numbers at all, but to label the properties by a set of abstract indices \mathcal{I} that have some of properties of the real numbers, but not all. Most importantly it should have the metric structure of the reals, the directional structure, but not the necessarily the additive and multiplicative structure that allows us to distinguish 0 from other numbers.

We can represent this structure by a function $d : \mathcal{I} \times \mathcal{I} \rightarrow \mathbb{R}$, which induces

³⁴I have not here taken sides on what it means for a class of properties to be metaphysically 'more special' than another class. One might think that this notion ought to be spelt out in terms of the easiness by which we can refer to that property, in accordance with Lewis's claim that the natural properties are reference magnets (see Dorr and Hawthorne [9] for some discussion of this role). This might be taken to indicate that 0-sitting is more natural than -5 -sitting. Note, however, that being easy to refer to does not always correspond with the pretheoretic idea of fundamentality: an alien species physiologically very different from us might take a minute to process all of their sensory experiences, so that there is a constant lag between events in the world and their experience of them. Such a species might find that simple easily expressed predicates for the properties of -1 -sitting particularly convenient and easy to use, and may even adopt more complex constructions for expressing the property of 0-sitting.

an ordering over the indices been given by: $\alpha \leq \beta$ iff $d(\alpha, \beta) \geq 0$. The structure we want can be given as follows:

$$d(\alpha, \beta) = 0 \text{ iff } \alpha = \beta.$$

$$d(\alpha, \beta) = -d(\beta, \alpha).$$

$$d(\alpha, \gamma) = d(\alpha, \beta) + d(\beta, \gamma) \text{ for every } \beta$$

$d(\alpha, \cdot)$ is surjective for each α . (For any $\alpha \in \mathcal{I}$ and $x \in \mathbb{R}$ there's a $\beta \in \mathcal{I}$ such that $d(\alpha, \beta) = x$.)

Our abstract indices are a bit like the reals, \mathbb{R} , in some respects, but we have thrown some information away, such as the location of the origin. (Note that the particular structure we are looking at above are specific to the non-relativistic setting, in which the temporal aspects of space-time is structured like the reals; we will later consider modifying some of these constraints.) For any index α , and real number x , there is a unique index which we shall write $\alpha + x$ which has the property that $d(\alpha, \alpha + x) = x$.³⁵

On the view we are exploring, time-shifted properties are all on a metaphysical par. In order to reflect this idea perspicuously in our notation, talk of the properties of 0-sitting and -5-sitting will now be replaced with the property of α -sitting and β -sitting for a suitable choice of α and β in \mathcal{I} . There are lots of different ways of assigning indices to replace 0 and -5 that would be equally acceptable – the only constraint is that $d(\alpha, \beta) = -5$. Finally, for each real number, x , we can introduce the timeshifting operation on properties. Let P be the set of all properties, understood as functions from times (or world time pairs) to functions from individuals to truth values. We write $[x]_P$ to denote the timeshift function of x minutes: the function that maps P onto itself which, for example, maps the property of α -sitting to the property of $\alpha + x$ -sitting.

It is important to emphasize at this juncture that the abstract indices themselves are no more than formal devices for labelling classes of related properties; they serve merely as a notational convenience, and the class of properties picked out by a given index are not ‘about’ that index in any interesting sense. The property we express in English with the word ‘spherical’, for example, is surely a paradigm example of a qualitative property, even though it’s picked out in our notation using an expression that *looks* like it contains something occupying the place of a singular term (it does not really: we should think of ‘ α ’ as merely a typographical marking, like a prime symbol or subscript).³⁶ One might argue that unlike the property of being spherical, the property of being spherical five minutes ago isn’t qualitative since it appears to be about the number five, or a

³⁵This index is unique: if $d(\alpha, \beta) = d(\alpha, \gamma) = x$ then $d(\beta, \gamma) = d(\beta, \alpha) + d(\alpha, \gamma) = -x + x = 0$. So by the first axiom $\beta = \gamma$.

³⁶It should be noted that according to some versions of endurantism, predicates like ‘spherical’ are context sensitive and covertly about particular times. On this view properties like *being spherical* are not qualitative either. As Haslanger [13] notes, however, there is no incompatibility between endurantism and the view that *being spherical* is intrinsic and qualitative once temporalism is assumed.

particular length of time. However remember that even though the property of being spherical five minutes ago on the present view is expressed by linguistically applying a certain tense operation to the predicate ‘is spherical’, this doesn’t reflect a fundamental asymmetry between the *property* of being spherical and being spherical five minutes ago: on this picture the latter property is just as qualitative as the former. Thus, I think, we really are describing a version of qualitative temporalism – the property of being spherical, for instance, is qualitative, and moreover things can have such properties temporarily.

What, then, happens to our argument that there is a uniquely special hypersurface? In the last section we singled this out as the fusion of events which are occurring. Note that the word ‘occurs’ as I just used it picks out some particular timeshift property: α -occurring. But as we noted above, there is nothing metaphysically special about the property of α -occurring. The property of β -occurring is no less distinguished, and there is also a unique hypersurface that is presently β -occurring. Which surface is presently β -occurring is constantly changing, just as which surface is currently α -occurring is constantly changing – however, the surfaces which are α and β occurring, respectively, are and always will be *different surfaces*. Indeed, assuming that space-time has a classical Newtonian geometry, then for every hypersurface in our (unique) foliation of space-time there is an index β in \mathcal{I} such that the fusion of events presently β -occurring just is that hypersurface.³⁷

Thus, one might argue, the initial appearance of there being a uniquely special hypersurface is really a byproduct of the fact that we have a linguistic preference for a certain property – the property of α -occurring, in this case – over other equally natural properties. But why, exactly, do we have a linguistic preference for the property of α -occurring, and not, say, the property of β -occurring? Why is it that we just use the present tense words like ‘occurs’, ‘sitting’, ‘talking’ and so on to pick out the property of α -occurring, α -sitting, and so on, but we must use a more complicated past tense construction to pick out the β variants of these properties? A parallel question can be asked about thoughts: why do our simple present tense thoughts result in one believing that Ellen is α -sitting, and not β -sitting?

The answer, in rough outline, turns on the natural idea that both thinking and semantically expressing are also tensed notions: there’s a notion of α -expressing and α -thinking which the simple present tensed constructions bear to the α indexed properties, but these constructions also bear the β -expressing and thinking relations to the β -properties. To properly understand this response, however, we need to extend the operation of timeshifting, which we have so far only applied to predicates, to other semantical categories. To state the operation in full generality requires us to adopt a slightly more precise framework, such as the higher order intensional logic of Montague or Gallin (see Williamson

³⁷Given all this it is, moreover, natural to think that the tense-logical distances between two such hypersurfaces, as determined by our function d , and the geometrical distances, as determined by the geometry of the space-time, should be the same (or proportional depending on how we choose units). In other words, the geometrical distance between the fusion of the α -occurring events and the fusion of the β -occurring events should be proportional to $d(\alpha, \beta)$.

[35] section 5.4 and 5.5.) In either system one constructs a system of domains, D_τ , for each type τ out of some basic domains and domain constructors. The three basic domains consist of a set of indices, T , representing the set of times in our case (but could also be a set of worlds, or world time pairs), a set of individuals I , and the set of truth values, $\{0, 1\}$, all assumed to be disjoint from one another. The domain constructors include: the function space constructor $D_1 \rightarrow D_2$, the set of functions from D_1 to D_2 , and the n -tuple space constructor, $D_1 \times \dots \times D_n$. Different intensional logics behave in different ways, but any domain in the two systems of Montague and Gallin just mentioned can be obtained by applying, in some order, these two domain constructors to the basic domains.³⁸ For any domain, D , we can introduce the identity function, id_D , on that domain in the obvious way. We can then introduce for each domain, D , constructed in this way, and permutation σ of the times, T , the shift $[\sigma]_D$, a function from D to D (a member of $D \rightarrow D$), as follows:

$$\begin{aligned} [\sigma]_I &= id_I, [\sigma]_{\{0,1\}} = id_{\{0,1\}}, [\sigma]_T = \sigma. \\ [\sigma]_{D_1 \rightarrow D_2}(F) &= [\sigma^{-1}]_{D_2} \circ F \circ [\sigma]_{D_1} \\ [\sigma]_{D_1 \times \dots \times D_n}(F_1, \dots, F_n) &= \langle [\sigma]_{D_1}(F_1), \dots, [\sigma]_{D_n}(F_n) \rangle \end{aligned}$$

If the indices of our model are times, then the permutations we are particularly interested in are permutations of moving each time back or forth by some fixed amount, which we will simply write $[x]_\rho$ for real number x . It is straightforward to check that our definition above agrees with our notion of timeshifting for n -ary predicates, which are from the domain $T \rightarrow (I^n \rightarrow \{0, 1\})$. Singular terms are left fixed by timeshifting operations; I shall accordingly suppress the abstract indices when using singular terms.

Rather than scrutinize this definition in great detail, let us simply look at how this operation behaves on the domain $(T \rightarrow 2) \rightarrow (T \rightarrow 2)$ of operators. This will allow us to apply it to the operator expression ‘sentence S means that ...’, or ‘Ellen thinks that ...’ which are of central concern to us here. Suppose that I α -believe that Ellen is α sitting. If $d(\alpha, \beta) = -5$ then the operator ‘I β -believe that’ is the result of applying the $[-5]_{(T \rightarrow 2) \rightarrow (T \rightarrow 2)}$ to the operator ‘I α -believe that’, and this just amounts to ‘5 minutes ago I α -believed that in 5 minutes time ...’. Thus β -believing is just the complex attitude I bear to a proposition, P , when five minutes ago I believed that in five minutes time, P .

(Small digression: Notice that this result arises because in our definition of type-shifting on function spaces we ‘conjugate’ the function: we compose it on either side with the timeshift operator and its inverse respectively. This definition is not unnatural.³⁹ In fact, the need for non-sentential tense operations behaving like those defined above arises in natural language. In some languages

³⁸Domains constructed this way will result in our generating surplus domains but all the types of interest in either Montague’s or Gallin’s system will be among the domains of our definitions.

³⁹The word ‘conjugation’ comes from group theory, where this kind of operation is applied all over the place.

the simple Priorian theory of tense is adequate: to talk about the past or future, on the Priorian theory, one simply applies a past or future tense operator to a present tensed sentence. In Chinese, for example, past and future events are expressed entirely by the means of adverbs – there is no verb inflection, or anything resembling tense – and is thus amenable to this Priorian style of theory. In many other languages, including English, however, we talk about the past and future by inflecting verbs, which therefore requires us to be able to apply tense operations to multiple semantical categories. For example, ‘Alice said that Bob ran’ is true if Bob’s alleged running was simultaneous with the time of Alice’s speech. Its truth does not require that Bob’s alleged running be prior to Alice’s speech; although this would be incorrectly predicted by a Priorian analysis which captured the tenses both inside and outside the embedded context (‘said’ and ‘ran’) by applying sentential tense operators.⁴⁰ On the other hand the sentence ‘Alice said that Bob is running’ is true only if Bob’s alleged running is present.⁴¹ Here the Priorian theory would put a single tense operator up front, to capture the past tense of ‘says’, and would incorrectly require only that Bob’s running coincide with the time of Alice’s speech. A benefit of our theory, on the other hand, is that in addition to sentential tense operators, we have other tense operations that can be applied to a wide range of semantical categories. In particular we can apply tense operations to operators themselves (the kinds of things associated with verbs like ‘says’ and other propositional attitudes); this allows us, for example, to change the tense of an attitude verb without thereby also modifying the tense of the complement clause. The definition of timeshifting involving conjugation is thus consistent with this data in a way that the Priorian theory isn’t. For more discussion of these problems see King [17], and Brogaard [3].⁴²)

Let us look now at the properties of meaning. ‘Ellen is sitting’ means that Ellen is sitting; or being precise about the indices, α -means that Ellen is α -sitting.⁴³ What does this present tense sentence β -mean, assuming still that $d(\alpha, \beta) = -5$? It β -means that P iff 5 minutes ago it α -meant that in 5 minutes time, P . Now English words are not constantly changing their meanings: in particular, ‘Ellen is sitting’ α -meant the same thing five minutes ago that it α -means now. Thus 5-minutes ago ‘Ellen is sitting’ α -meant that Ellen is α -sitting, which is logically equivalent to the claim that in five minutes time, Ellen is β -sitting. So ‘Ellen is sitting’ β -means that Ellen is β -sitting.

So, in short, it is not entirely true that we linguistically favour the property of α -sitting over β -sitting: we α -express the former property with the present

⁴⁰More precisely, it captures the past tense of ‘Alice says...’ with ‘WAS[Alice says ...]’ and the past tense of ‘Bob runs’ with ‘WAS[Bob runs]’, and thus predicts ‘WAS[Alice says WAS[Bob runs]]’ as the truth conditions for this sentence.

⁴¹Presumably Bob’s alleged running must also overlap the time of Alice’s speech. Note also that the direct speech report ‘Alice said ‘Bob is running’’ is true, even if Bob’s alleged running is past; it is sometimes possible to get a true reading of a false indirect speech report by reinterpreting it in terms of the corresponding direct one.

⁴²King uses these kinds of arguments to argue against temporalism.

⁴³Note that given our definition of timeshifting for individuals, α -Ellen is the same individual, Ellen, for each α .

tensed word ‘sitting’, but not the latter, yet we β -express the latter with this word and not the former.

Let us now try to fit all these thoughts together into the wider picture. Let us start with an analogy that will be heuristically valuable. Consider the following version of presentism. The entirety of the universe is contained within a persisting three-dimensional manifold in which all objects have spatial extent only – there is no physical temporal dimension, and nothing is temporally extended. Entities located within the three dimensional manifold are continually changing their properties and their locations and so on. However unlike the standard forms of presentism there are a few important differences. The first is that although the entities are continually changing their properties, no entity ever goes out of existence or comes into existence, although they might stop being located somewhere. The second is that there is something special about one of the spatial dimensions: when an object changes its location, it always ends up to the right of its earlier locations, although there are no constraints about its location in the other two dimensions. All objects are continuously moving to the right in some path or other, as though there were some force constantly pushing them in that direction. More generally, some of the properties an entity has, for instance geometrical properties like being round, supervene on the eternal properties of certain regions of space. Which region determines these properties changes as time goes on, but the region that determines these properties is continually moving to the right. Lastly, although I spoke of each objects unique ‘location’, it turns out that there are infinitely many other relations that behave just like location. Some of these relations put each object to the right of where they’re located, and others put them to the left, but all of these location relations are changing in unison, each pushing the object through space in the same overall direction. Likewise, each object typically has an abundance of connected properties – properties a bit like sitting, talking, and so on – which are determined by the goings on in certain regions of space, with the region that determines these properties constantly moving. There is finally an interesting connection between the purely spatial separation along the distinguished dimension of an objects location at one time, and then at another: the distance between the objects location now and in x minutes time will be directly proportional to x (i.e. the component of the objects velocity along the distinguished dimension is constant).

The analogy between our view and this version of presentism, I hope, should be clear. To perfect the analogy we must simply increase the dimensionality of the manifold by one – the temporal dimension is distinguished from the others simply by the special role it plays in physics: things are moving at a constant rate along one of the dimensions as a matter of physical law. The connection this physical dimension has to time proper (as understood in terms of tense operators) is mediated by this physical role. It is also a picture in which there is change. The objects in our system are continually changing their properties, but because there is an abundance of these properties this does not result in any particular region of the space-time being picked out as special.

It is also worth comparing this picture with Williamson’s description of

the temporalist picture. The way that Williamson talks suggests that objects don't to have many interesting properties unless they are located somewhere (in Williamson's terminology, if they are 'chunky'). For example Williamson (p13-14) writes: 'A temporaryist may ask a permanentist: but what is the past or future coin *now*? Similarly, a contingentist may ask a necessitist: but what is the merely possible coin *actually*? Such questions are in effect a challenge to give more 'concrete' properties of the alleged temporarily or contingently non-concrete thing. However, it is contentious to presuppose that everything non-abstract has such more 'concrete' properties.' But if the word 'chunky' and 'located' express α -chunkiness and α -locatedness, rather than the β variants, then this looks like a language dependent distinction at best – one that does not have much metaphysical significance. Why is being β -chunky or β -located not any less of a 'concrete' property (unless one insists, unduly, that we focus on some penumbrally connected notion of α -concreteness over β -concreteness)?

Let us now attempt to give a toy model of a formalized temporal language. A particularly perspicuous language would be one that had primitive predicates for each of the time shifted properties, thus exhibiting no bias for some timeshifted properties over others. To do this we must require that we have, for each predicate of the less perspicuous language (containing predicates such as 'is sitting') an infinite set of predicates in the perspicuous language, 'is α -sitting' for each α in \mathcal{I} . If our language is higher order this process must be repeated for each grammatical category. In principle we should even index the the logical constants, however since these are extensional the indexed variants will all be logically equivalent to one another. Of course, most of the non-logical vocabulary will be redundant on this way of setting things up: the metric tense operators can be used to convert between the different timeshifted predicates, so one could in principle take the vocabulary set corresponding to one of our abstract indices and define the rest from that. However any choice of indexed properties, in this regard, would be just as good as any other.

Let us begin with a first order tense logical language L , containing un-indexed predicates corresponding to ordinary predicates of English like 'is sitting'.⁴⁴ In addition to its non-logical predicates, L contains a pair of unary tense operators \triangleright and \triangleleft and for each $x \in \mathbb{R}$ there is a unary operator $[x]$. \triangleright and \triangleleft are to be read 'it will be the case that' and 'it was the case that'. From these one can define operators 'it will always be the case that P ', written $\blacktriangleright P$, as $\neg \triangleright \neg P$; 'it was always the case that P ', written \blacktriangleleft , can be defined analogously. $[x]P$ on the other hand is to be read as 'in x minutes time P '. When x is negative $[x]P$ clearly implies $\triangleleft P$, and when positive $\triangleright P$. A model of L consists of (i) a set of linearly ordered times, T , with the metric structure of \mathbb{R} , which in the intended model will be identified with a set of hypersurfaces, (ii) a domain D , and (iii) a binary function $\llbracket \cdot \rrbracket$, that assigns to each time $t \in T$, and n -ary predicate P in L a subset of D^n , written $\llbracket P \rrbracket_t$ and a unary function $\llbracket \cdot \rrbracket$ that assigns to each constant c an element of D , written $\llbracket c \rrbracket$. For each time

⁴⁴All of the following ideas extend to higher order intensional languages; I focus on the first order case to keep the number of moving parts to a minimum.

$t \in T$ truth relative to a time and variable assignment, written $t, v \Vdash \phi$, can then be defined inductively. The clauses for the quantifiers and truth functional connectives are completely standard, the only clauses worthy of mention are those governing the tense operators:

$$t, v \Vdash [x]\phi \text{ iff } t + x, v \Vdash \phi$$

$$t, v \Vdash \triangleright\phi \text{ iff } t', v \Vdash \phi \text{ for some } t' > t$$

$$t, v \Vdash \triangleleft\phi \text{ iff } t', v \Vdash \phi \text{ for some } t' < t$$

In the first clause, $t + x$ refers to the unique time $t' \in T$ greater than t whose distance from t is x if x is positive, or the unique time less than t whose distance from t is $-x$ if x is negative.

From L we can construct a more peripicuous language L^+ , containing primitive predicates for each timeshift of the predicates of L . L^+ does not contain L 's predicates, but if P is an n -ary predicate in L , then L^+ contains an infinite collection of corresponding predicates P_α , one for each $\alpha \in \mathcal{I}$. Otherwise L^+ is just the same as L . It is straightforward to extend our model just described to L^+ . One must construct an extended interpretation function, $\llbracket \cdot \rrbracket^+$ that assigns extensions to the new predicates in L^+ . This can be achieved by selecting a fixed element, $\alpha \in \mathcal{I}$ (it does not matter which) and stipulating that $\llbracket P_{\alpha+x} \rrbracket_t^+ = \llbracket P \rrbracket_{t+x}$. Recall that every index $\beta \in \mathcal{I}$ is identical to $\alpha + x$ for some $x \in \mathbb{R}$ so this definition is complete.

However I think there is a better way to think of the above model, which will be much more useful when we try to generalise. Let us rather think of times in our new model as a distance preserving mapping from our abstract indices, \mathcal{I} , to times in our original model, T . If, for example, f assigns α some element of T , then due to the distance preserving nature of f , f 's action on the remain indices is fixed: $f(\alpha + x) = f(\alpha) + x$. Since every abstract index is of the form $\alpha + x$ for some real x , these facts completely determine f 's action on the indices. These functions have a natural order, where $f < g$ iff $f(\alpha) < g(\alpha)$ for some α (note that if $f(\alpha) < g(\alpha)$ for some α , $f(\alpha) < g(\alpha)$ for every α). The functions also have a metric structure, given by the metric defined over T : the distance between f and g is simply the distance in T between $f(\alpha)$ and $g(\alpha)$ (this distance is the same for any choice of α). We can thus talk about f being x minutes to the future of g . If f is one such mapping then we can determine the extension of a predicate, P_α , at f by looking at the extension assigned to the unindexed predicate P at $f(\alpha)$ in our original model of L : $\llbracket P_\alpha \rrbracket_f^+ := \llbracket P \rrbracket_{f(\alpha)}$.

To see how all of this formalism can actually be put to use, let us start by describing a fairly concrete example of an L model, and then an L^+ model. In the model of L we identify T with a foliation of a classical space-time into hypersurfaces (which is uniquely determined by the structure of classical Newtonian space-time). Our domain D shall consist of a single person, Ellen, who is associated with a certain four-dimensional space-time worm. We also have a predicate S whose intended interpretation is the property of sitting. The extension of the unary predicate S contains Ellen at a hypersurface t ($\text{Ellen} \in \llbracket S \rrbracket_t$) if

and only if the shape of ts intersection with Ellen's space-time worm is a sitting shape (I assume sitting is purely a geometrical property for now). Otherwise the extension of S is empty at that hypersurface. One can see by inspection that the formalization of the sentence 'Ellen is sitting' is true relative to some times, and false relative to others provided her space-time worm intersects some but not all hypersurfaces in a sitting shape.

To extend this model to L^+ , then, we must know whether Ellen is α -sitting at each time, for each choice of α . Here a time simply maps each index, α , to a hypersurface, and that indirectly determines whether Ellen is α -sitting. If it maps α to a hypersurface that intersects Ellen's space-time worm in a sitting shape then she's α -sitting relative to that time, and if it doesn't she isn't.⁴⁵ According to later times α will be mapped to temporally later hypersurfaces, which again will indirectly determine whether Ellen is α -sitting at these later times. This all motivates our crucial formal insight, which is summarised below:

TIMES: Times can be represented by mappings from abstract indices to hypersurfaces.

In section 2 I distinguished the tense-logical notion of a time from the physical one. Here that distinction bears its fruit, for the above tense-logical entities do not correspond uniquely to hypersurfaces (i.e. physical times): each tense-logical time bears some relation (via some index) to every hypersurface. It is important to stress, however, that this definition encodes no more than a convenient formal representation of times. The abstract indices are nothing more than a helpful technical device introduced as a way of labelling certain classes of properties that bear an important similarity to one another. Identifying times with mappings from these indices to hypersurfaces is merely a compact way of retrieving facts about what the extensions of these classes of properties are at each time from the eternal facts about the distribution of eternal properties over the space-time manifold.

These entities are moreover formally easy to manipulate because all of the operations that we need to be able to define on times can be fairly straightforwardly lifted from the corresponding operations on hypersurfaces. For example, to make sense of the $[x]$ operator we need to know what it means for one time to be x minutes in the future or past from another. This is achieved as follows: if f is a time, and x a real number, we can define the time x minutes to the future of f , written $f + x$, as the mapping which takes α to the hypersurface that is temporally separated from the hypersurface $f(\alpha)$ by x minutes. The ordering on times is induced in the same way – f comes before g iff the hypersurface f maps any particular index to comes before the hypersurface g maps that index to.

⁴⁵Note that one does not need to be a perdurantist in order to have the concept of an objects space-time worm. I simply mean the fusion of regions that object is sometimes located at. (I'm presently using 'located' in this definition to pick out a certain relation – α -location say – however all choices of α would produce the same fusion.)

4 Timeshifted Properties in Relativistic Physics

The extension of these ideas to the relativistic case presents a number of new problems. Let us suppose, as we argued earlier, that the fusion of events that are presently occurring forms a plane of simultaneity. Of course, by ‘occurring’ I just mean α -occurring for some α : as before there are other hypersurfaces representing the events that are β -occurring for other indices β , where what is now β -occurring are just those things that were once α -occurring, or will be α -occurring. Given our argument in section 2, these hypersurfaces form a particular foliation of space-time. Thus although we avoid commitment to the existence of a ‘special’ hypersurface we apparently are still committed to a special foliation of the space-time. We saw this to be particularly problematic when combined with certain features of special relativity.

A natural response would be to postulate a greater variety of tensed properties. In the classical case we had the following fact about timeshifted properties: for any α and β either the property of α -sitting just is the property of having been β -sitting x minutes ago (for some x), or the other way around. Perhaps there are pairs of tensed indexed properties that aren’t related in this way. Our earlier argument established, for example, that when α and β are comparable – a timeshift, $[x]$, by some real number x can take you from one kind of property to the other – the fusion of events α -occurring and β -occurring respectively are parallel planes of simultaneity. But perhaps there are pairs of indices α and β that determine classes of properties that are not comparable in this way – in these cases we have no reason to think the fusion of events α -occurring and β -occurring are parallel to one another.

Here is one extremely natural reason to postulate the existence of such temporal properties. Imagine that you and I are moving inertially relative to one another, but we both begin our experiment from the same location. Suppose we are both going to observe Ellen, who is presently sitting, and measure how long she sits for before standing up. As we know, one of the surprising features of special relativity is that our measurements might end up differing: I might measure six minutes while you measure ten. Who, then, is right, and who is wrong? I maintain that both measurements are correct, but we are measuring different temporal properties of Ellen. You are perhaps measuring the length of time for which Ellen is α -sitting, and she may indeed be α -sitting for ten minutes, while I am measuring how long Ellen was β -sitting, which also may be six minutes. Neither of us have any false beliefs, and you may even correctly infer from your observations (by applying something called a ‘Lorentz transformation’) how long she was β -sitting, and I might be able to infer how long it was that she was α -sitting. If both of these properties exist, however, there will be no timeshift that takes us from one to the other: β -sitting will not simply amount to α -sitting x minutes ago, for instance. This idea therefore involves a slight deviation from the way in which we introduced these properties in the last section; these properties are not straightforwardly obtained by applying timeshift operations to the properties we express with familiar predicates like

‘is sitting’.⁴⁶

It turns out that there’s a natural way to extend the model we gave in the previous section to the relativistic setting that delivers exactly the verdicts suggested above. Recall that in that model we identified times with mappings from abstract indices to hypersurfaces taken from a classical space-time manifold. If f and g are two such mappings then f maps each index to a hypersurface that is separated by a fixed temporal distance from the hypersurface g maps the same index to. Thus it follows that for some real number x , $g = f + x$ (i.e. for all α , $g(\alpha)$ is the hypersurface separated by x minutes from $f(\alpha)$). So in general, once we’ve picked a single time, f , we know that all other times will be of the form $f + x$ for some real number x . This same idea can be applied to the relativistic case. Let f be some bijective mapping from abstract indices to the set of all possible planes of simultaneity. As in the classical case, we must, for each real number x , define the time $f + x$ which we can think of as the time that comes x minutes before or after f , depending on x ’s sign. The totality of times of the form $f + x$ will form a linearly ordered set whose structure is exactly parallel to the structure of our time mappings in the classical setting. In this regard, then, we do not deviate at all from the classical picture: the structure of *times* is completely linear, and the resulting tense logic will be standard linear tense logic. Moreover the structure of these times will not depend on the trajectory of the observer: two people moving inertially relative to one another can talk to one another using the same tense operators as the other.

So it all boils down to whether, given a bijective mapping f from abstract indices to planes of simultaneity, we can define the function $f + x$. Fortunately our original definition survives unscathed. In special relativity, given a plane of simultaneity there’s a well defined operation that moves that plane backwards or forwards by a certain time-like interval, x , whilst keeping it parallel with its original orientation. Thus the function $f + x$ is just the function that maps each index α to the hyperplane that is separated by a time-like interval of x into the future or past (depending on whether x is positive or negative) from the hyperplane $f(\alpha)$.

This definition also extends to the context of general relativity. If the space-time is ‘static’ – i.e. it admits a timelike Killing field orthogonal to a spacelike hypersurface – then we can identify a special class of hypersurfaces that generalise planes of simultaneity in a natural way: those hypersurfaces orthogonal to some timelike Killing field. A hypersurface like this can be moved backwards and forwards along its Killing field by an interval x , giving us the requisite structure. If the space-time does not admit a timelike Killing field of this sort, then a more general class of hypersurfaces are available: geodesically generated surfaces – smooth hypersurfaces consisting of all geodesics that pass orthogonally through a given time-like vector. Instead of having a time map each index to a plane of simultaneity we map each index to a pair of a normalised time-like vector and the geodesic surface it generates, S . These entities can similarly be

⁴⁶Although the contextualist view I develop shortly would allow them to be introduced as the properties expressed by these familiar predicates in various different contexts.

shifted backwards and forwards in the way required.⁴⁷

At any rate, we can see, at least *formally*, how to extend the model of the last section to various relativistic settings. But what exactly is the resulting picture? Let us set aside the special problems that arise in GR for the time being and restrict ourselves to the case of special relativity. At a time f , Ellen is α -sitting iff $f(\alpha)$ is a hypersurface that intersects Ellen's space-time worm in a 'sitting' shape. Whether Ellen is α -sitting will change as time evolves: maybe she's α -sitting but in ten minutes she won't be. This would happen in our model if $f(\alpha)$ intersects Ellen's space-time worm in a sitting shape, but the hypersurface parallel but whose time-like separation from $f(\alpha)$ is ten minutes to its future does not intersect Ellen's space-time worm in a sitting shape. What's interesting is that there might be another hyperplane, h , intersecting Ellen's worm at the same region of space-time that $f(\alpha)$ does that isn't parallel to $f(\alpha)$. Since f is a bijection, it follows that there is an index β that gets mapped to h .

One natural question to settle is whether Ellen is also β -sitting (at time f)? Perhaps the answer is yes in this case, although it's a feature of special relativity that not all geometric properties are preserved under Lorentz transformations: if my trajectory is orthogonal to the plane of simultaneity $f(\alpha)$ I might observe something to be square shaped, while you might observe it to be a parallelogram if your trajectory was orthogonal to a different plane of simultaneity such as $f(\beta)$. We in principle have the resources preserve certain shape properties one might have thought couldn't survive in special relativity by saying that one and the same object could be both an α -square and a β -parallelogram. However this is a choice point; one could develop the theory of shapes in a more objective way.

For our purposes, it is more interesting that this model delivers the kinds of temporally diverse properties alluded to above. Supposing that Ellen is β -sitting at time f , then she will remain β -sitting at all times $f + y$ such that fastforwarding the hyperplane $f(\beta)$ by y intersects her worm in a sitting shape. It follows, however, that the length of time for which Ellen will be β sitting need not be the same as the length of time for which she'll be α -sitting: perhaps she'll be α -sitting for six minutes and then she'll α -stand up, but it won't be ten minutes until she β -stands up. This might all sound quite puzzling, but, I think, it is no more puzzling than the phenomena that special relativity itself seems to force upon us. The underlying moral of relativity, on this picture, is that when we're moving along certain trajectories we can directly observe and measure certain temporal properties which cannot be directly measured by people moving along other trajectories (although we can indirectly calculate them). This is in fact consistent with the example we described earlier. If my trajectory is orthogonal to $f(\alpha)$, I might measure what I would describe as 'Ellen sitting' for ten minutes, before I observe what I would describe with

⁴⁷Given a geodesically generated surface, (S, ξ) , we can construct the surface separated by a space-time interval of x into the future or past of (S, ξ) as follows. Let $\gamma(t)$ be the unit speed geodesic passing tangentially through ξ ($\gamma'(0) = \xi$). $(S, \xi) + x$ may then be identified with the surface geodesically generated by $\gamma'(x)$ – the tangent to the curve γ at x .

the words ‘Ellen standing up’. What I have measured, completely accurately, is the length of time for which Ellen was α -sitting. Someone else moving in a trajectory orthogonal to $f(\beta)$ might instead observe what they would describe as ‘Ellen sitting’ for six minutes. This person would not have been incorrectly measuring how long Ellen was α -sitting – ten minutes – but correctly measuring the length of time that she was β -sitting, namely six minutes.

This, in brief, is how we account for the phenomena of time dialation associated with special relativity. It is extremely natural to pair this view with a kind of contextualism about temporal predicates. Roughly people who are at rest relative to one another tend to express the same kinds of properties with predicates like ‘sitting’, ‘standing’, and even with expressings like ‘saying’ and ‘uttering’. Let’s suppose that me and my co-movers express α -sitting, α -standing, α -saying, α -uttering and so on.⁴⁸ To ease readability I shall surpress the reference to the index α , and just use the expressions ‘sitting’, ‘uttering’, etc, since, after all, in our present frame these just mean the same as ‘ α -sitting’, ‘ α -uttering’, and so on. People moving inertially relative to us right now express different properties with these same predicates; their β variants. The view is thus a contextualist view. However note that this is not because the tense operators, ‘it will be the case (in x minutes)’ and ‘it was the case (in x minutes)’, are context sensitive – these express the same things relative to all inertial trajectories – it is because simple tensed predicates are context sensitive.

A somewhat striking aspect of the view I have been describing is that time has an objective rate, which every one can agree on no matter what their trajectory is: the tense operators, as I mentioned earlier, express the same thing in all contexts regardless of ones trajectory. The flipside is that objects possess a wider variety of temporal properties: there are, as we saw, many sitting-like properties which a person possesses for different amounts of time.

Since we are postulating the existence of an objective rate of time, it’s worth working through what happens to the infamous ‘twin paradox’ – a case which is usually taken to demonstrate the impossibility of objective time. These kinds of cases have the following structure: suppose that we’re initially moving together along the same inertial trajectory, and that you briefly accelerate away from me and travel inertially for a little while, before accelerating again, reversing your direction, so that you eventually meet up with me again. The paradox is that, according to special relativity, when we meet up again we will have measured different lengths of time passing. I might have measured a total of ten minutes passing, while you might have measured a total of six minutes. Given that the view I have been describing posits that time has an objective rate of time it is tempting to ask how much time *really* passed between our leaving each other and meeting again? Was it really ten minutes, six minutes, or some other number altogether?

⁴⁸One could argue that if a large linguistic community was roughly co-moving with one another, then even though they often move at small relativistic speeds relative to one another, there’s a certain amount of linguistic deference which prevents a constant tiny changes of meaning. (Perhaps they go by the linguistic communities average trajectory, or something like that.)

The truth of the matter, however, is that this question is ambiguous. Call the space-time event of you leaving me e_1 , the event of you reversing your direction e_2 and the event of our meeting back up e_3 . The length of time between the occurrence of e_1 and e_3 may be ten minutes, if I am measuring the length between their α -occurrence (suppose here that the present time maps α to a plane orthogonal to my trajectory). However, for different choices of β , the length of time between the β -occurrence of these events can vary. Suppose that when you start moving away from me until when you turn around, at e_2 , you measure the difference between the β -occurrence of e_1 and e_2 : three minutes. After turning around you measure the length of time between the γ -occurrence of e_2 and e_3 to be also three minutes. So we can deduce that the following is true at the time, f , at which we are initially together:

In three minutes time the following will be true: [e_2 will γ -occur and in three minutes time e_3 will γ -occur].

It follows simply by tense logic that in six minutes time e_3 will γ -occur.⁴⁹ We also know by the setup that it is ten minutes until e_3 will α -occur.

So before you are about to depart we can ask how long until we meet up again. In this case there are two salient questions in the vicinity which we must not confuse: it won't be 10 minutes until we α -meet up again, and it won't be another 6 minutes until we γ -meet up. This situation on its own is not that perverse. Even in the Newtonian setting we can find properties that behave like this: for example, if $\delta = \alpha + 4$ then it'll be 10 minutes before we α -meet up but only 6 minutes until we δ -meet up. What is puzzling about the twin paradox is that two initially synchronised people – people whose watch's measure the length of time between the α -occurrences of events – can become unsynchronised. After accelerating your watch will instead measure the length of time between the β -occurrence of events, and then the length of time between the γ -occurrence of events.⁵⁰ (The sense in which your watch measures the time between β -occurrence of events can be spelt out as follows: in five seconds time the second hand on your watch will have α -ticked thrice, but β -ticked five times.)

Our earlier discussion of context sensitivity helps dispell some other puzzles that arise in connection with the twin paradox. Suppose we are currently together, and you are about to start your journey. In five minutes time, however, you will be at the point in your journey at which you are about to α -turn around (I could simply report this using the phrase 'the point in your journey at which you turn around' – I am simply adding the indices for clarity). Suppose at this point you utter (i.e. α -utter) the sentence 'it's been three minutes since we parted ways', since this is indeed the amount of time you have experienced. It is extremely natural to think that in doing so you will α -express a truth. To say otherwise would be to invite the 'false belief' problem we raised in section 2: that anyone in motion relative to me is potentially going around uttering

⁴⁹I am appealing to the inference from $[x](A \wedge [y]B)$ to $[x + y]B$.

⁵⁰These will be the same if the speed of your outward journey matches the speed of your return journey relative to you're trajectory.

falsehoods. If you have α -expressed a truth, however, it is not the proposition that it's been three minutes since we α -parted ways, because this would be a falsehood in five minutes time. The theory we suggested earlier, however, predicts that you α -express the true claim that it's been three minutes since we $(\beta - h(5))$ -parted ways; this is because at the time of your α -utterance you are α -moving in a trajectory orthogonal to the hyperplane of $(\beta - h(5))$ -occurring space-time events. Thus in different contexts, the same sentence can be used to α -express different propositions.

4.1 The Twins Paradox

In this section I shall describe in more detail the way in which the twins paradox gets resolved in the contextualist framework outlined. In particular, it brings to light a possible confusion that can arise about the contextualism I've described which takes a little bit of disentangling. Those who do not wish to delve into the details can skip this section.

Let's begin by highlighting a feature of the version of contextualism proposed here. We shall show shortly that according to this kind of contextualism, anyone moving at rest relative to me will be resolving the contextual parameters in the same way as I am and will continue to do so as long as they are at rest relative to me. For example, I will express the same property with the word 'sitting' as those at rest relative to me and the property expressed will be constant. By contrast, it is also possible to show that people moving inertially relative to me will not only be speaking in a different context, but the context they're in will be constantly shifting, depending on how long they have been moving away from us. For such people the context they are in at the beginning of the day will be different from the one they are in at the end of the day, unless some special effort to speak in my context is made. Thus while the word 'sitting' expresses a constant property for me, the person who is moving relative to me will be continually changing which property they express with their uses of the word 'sitting'.

We can make this a bit more precise as follows. Let us suppose that Alice, moving inertially relative to me, says that Ellen is β -sitting when she utters the sentence 'Ellen is sitting'.⁵¹ Let us suppose, as above, that although Ellen will sit for the next ten minutes, in six minutes time Ellen will stop β -sitting. Thus if Alice meant the same thing by 'sitting' throughout this period, namely β -sitting, then she ought to be willing to utter 'Ellen is sitting' for the next six minutes, before switching to 'Ellen is standing'. But this is not what we in fact observe: we observe time dialiation regarding periods of uttering to the same extent that we do with regard to periods of sitting. Although from Alice's perspective she appears to be "accepting" 'Ellen is sitting' for a period of six minutes that is because she is reporting the length for which she β -accepts this sentence; in fact she accepts (i.e. α -accepts) 'Ellen is sitting' for the ten minutes we observe from our perspective. It follows from the general kinds principles

⁵¹Recall that I am talking about what Alice α -says and α -utters here.

we have been expressing so far that Alice will accept ‘Ellen is sitting’ for the same length of time for which Alice is sitting, namely the next ten minutes. Thus Alice will be accepting ‘Ellen is sitting’ for several minutes after Ellen has stopped β -sitting. Since, by hypothesis, Ellen is speaking truthfully now and in the future it follows that she will no longer express the property of β -sitting with the word ‘sitting’ in x minutes time, but rather she will express the property of $(\beta - h(x))$ -sitting where h is some linear function of the time elapsed that is determined by her relative velocity to me.⁵²

This result looks as though it reinstates a problem we have directed at other views: it looks as though one frame of reference is being distinguished from the others. In some frames of reference people speak from a constant context; for people moving in that kind of frame the word ‘sitting’ will express the same property – say, α -sitting – throughout the day. In other frames, however, the context will change over time – for people moving along the latter kind of frame words like ‘sitting’ will express different properties at the beginning of the day than at the end.

As indicated, this may at first appear to indicate an asymmetry between me and Alice – she is constantly changing what she means by ‘sitting’ whereas I’m not. But it is not really an asymmetry: Alice is constantly changing what she α -means by ‘sitting’, but there’s another relation her utterances bear to propositions – β -meaning. Alice’s utterances constantly β -express the same proposition, and it is my utterances whose β -meanings are continually being resolved in different ways.

This confusion arises roughly because we are forgetting that according to the view in question theoretical terms like ‘context sensitive’ are themselves context sensitive. It can be used in some contexts to mean a sentence which α -expresses different things in different contexts, and in other contexts it can be used to mean a sentence which β -expresses different things in different contexts. While people moving relative to me may be continually α -expressing different things with their words, the things they β -express remain constant, and it is us who are continually shifting which things we β -express with our tensed vocabulary.

The situation I’ve just described is predicted by the following general method for calculating, at a time f , what a sentence, S , α -expresses in a context, c . Firstly find where the hypersurface $f(\alpha)$ intersects the trajectory of the speaker of c . Then we find the hypersurface, h , orthogonal to that speaker’s trajectory. Finally we obtain an index, $\beta = f^{-1}(h)$, which we use to interpret the speakers utterances: this just means that her uses of ‘sitting’ will α -express β -sitting, ‘standing’ will α -express β -standing, and so on. According to this model, people moving orthogonally to $f(\alpha)$ will always α -express the same thing with the word ‘sitting’. People not moving orthogonally to $f(\alpha)$ will presently α -express a different thing with the word ‘sitting’ than they will α -express in five minutes time, but their utterances involving ‘sitting’ will instead β -express a constant meaning, where β is such that $f(\beta)$ is orthogonal to their trajectory.

⁵² h can be calculated explicitly but the details are not particularly illuminating for our present purposes.

5 Conclusion

The view I have described here has some affinities with eternalism: it is contextualist about simple present tensed predicates like ‘is sitting’, and does not distinguish any hypersurface as more special than any other. The resulting view is thus very much at odds with the conventional permanentist version of the A-theory – the moving spotlight theory – in which a particular region of space-time is metaphysically favoured. However unlike eternalism the properties tensed predicates express in different contexts are not time indexed properties, like sitting at a particular date, or sitting relative to an inertial trajectory; they are temporary properties whose extensions change over time. The result is a version of qualitative temporalism that does not engender widespread error when combined with modern physics.

References

- [1] N. Ángel Pinillos. Time dilation, context, and relative truth*. *Philosophy and Phenomenological Research*, 82(1):65–92, 2011.
- [2] Ruth C. Barcan. A functional calculus of first order based on strict implication. *Journal of Symbolic Logic*, 11(1):1–16, 1946.
- [3] Berit Brogaard. *Transient truths: an essay in the metaphysics of propositions*. Oxford University Press, 2012.
- [4] Herman Cappelen and John Hawthorne. *Relativism and monadic truth*. Oxford University Press, 2009.
- [5] Daniel Deasy. The moving spotlight theory. *Philosophical Studies*, 172(8):2073–2089, 2015.
- [6] C. Dorr. The expressivist and the relativist.
- [7] Cian Dorr. How to be a modal realist. 2010.
- [8] Cian Dorr and Jeremy Goodman. Diamonds are forever. *Nous*, Forthcoming.
- [9] Cian Dorr and John Hawthorne. Naturalness. 2013.
- [10] Kit Fine. Prior on the construction of possible worlds and instants. *World, times, and selves*, 1977.
- [11] Kurt Gödel. Rotating universes in general relativity theory. In *Proceedings of the International Congress of Mathematicians Edited by LM Graves et al., Cambridge, Mass. 1952, vol. 1, p. 175.*, volume 1, page 175, 1952.
- [12] Robert Goldblatt. Diodorean modality in minkowski spacetime. *Studia Logica*, 39(2-3):219–236, 1980.

- [13] Sally Haslanger. Endurance and temporary intrinsics. *Analysis*, pages 119–125, 1989.
- [14] Mark Hinchliff. The puzzle of change. *Nous*, pages 119–136, 1996.
- [15] David Kaplan. *Demonstratives*. 1977.
- [16] David Kaplan. *Themes from Kaplan*, chapter Demonstratives. Oxford University Press, 1989.
- [17] Jeffrey C King. Tense, modality, and semantic values. *Philosophical perspectives*, 17(1):195–246, 2003.
- [18] John MacFarlane. *Assessment sensitivity: Relative truth and its applications*. Oxford University Press, 2014.
- [19] David B Malament. Classical relativity theory. *Handbook of the philosophy of science. Philosophy of physics. Part A*. Amsterdam: Elsevier, pages 229–275, 2006.
- [20] David B Malament. *Topics in the foundations of general relativity and Newtonian gravitation theory*. University of Chicago Press, 2012.
- [21] Richard Montague. The proper treatment of quantification in ordinary english. 1973.
- [22] Thomas Müller. *The language of special relativity*. Bibliothek der Universität Konstanz, 2004.
- [23] Arthur N Prior. Thank goodness that’s over. *Philosophy*, 34(128):12–17, 1959.
- [24] Arthur N Prior. *Past, present and future*, volume 154. Clarendon Press Oxford, 1967.
- [25] Arthur N Prior and Per FV Hasle. *Papers on time and tense*. Oxford University Press, 2003.
- [26] Hilary Putnam. Time and physical geometry. *The Journal of Philosophy*, pages 240–247, 1967.
- [27] Wim Rietdijk. A rigorous proof of determinism derived from the special theory of relativity. *Philosophy of Science*, pages 341–344, 1966.
- [28] Steven F Savitt. There’s no time like the present (in minkowski spacetime). *Philosophy of Science*, pages S563–S574, 2000.
- [29] Theodore Sider. Four dimensionalism: An ontology of persistence and time. 2001.
- [30] Theodore Sider. *Writing the Book of the World*. Oxford University Press, 2011.

- [31] Lawrence Sklar. Time, reality, and relativity. 1981.
- [32] Bradford Skow. Relativity and the moving spotlight. *The Journal of Philosophy*, pages 666–678, 2009.
- [33] Howard Stein. On Einstein–minkowski space–time. *The Journal of Philosophy*, pages 5–23, 1968.
- [34] Meghan Sullivan. The minimal a-theory. *Philosophical studies*, 158(2):149–174, 2012.
- [35] Timothy Williamson. *Modal logic as metaphysics*. Oxford University Press, 2013.
- [36] Dean Zimmerman et al. Presentism and the space-time manifold. *The Oxford handbook of philosophy of time*, pages 163–244, 2011.