

Tensed Truth, Temporal Particularity, and the Fixity of the Past

Abstract

Our ordinary conception of time has it that there are temporal particulars: not only do people do things, but there are particular doings by people; not only are we born, but the birth of each one of us was a particular event, and each of us will have our own particular death. Temporal particulars in this sense are individuated, fundamentally, by their temporal locations or relations, rather than by their intrinsic or qualitative characteristics. In this respect they are unrepeatable, not just *de facto*, but as a matter of their very nature.

However, there is a tradition in philosophy that seriously downplays this aspect of our thinking about time. According to this tradition, the fundamental unit of temporal representation is *tensed truth*; the notion of an unrepeatable particular, individuated by its temporal location, is at best an abstraction from a complex of tensed truths. The aims of this paper are, first, to argue that the representation of temporal particulars is deeply implicated in our ordinary conception of the past as fixed and unalterable; and, secondly, to argue that the theorist of tensed truth is able to provide only a pale imitation of this aspect of our thinking. I will then reflect on the consequences of this for debates about the metaphysics of tense.

1 EVENTS, COUNTERFACTUALITY, AND THE SENSE OF THE PAST

Our ordinary conception of time has it that not only do people do things, but there are particular doings by people; not only are we born, but the birth of each one of us was a particular event, and each of us will have our own particular death. Events in this sense are *temporal particulars*, individuated by their temporal locations or

relations, rather than by their intrinsic or qualitative characteristics. In this respect individual events are unrepeatable, not just *de facto*, but as a matter of their very nature.

A sensitivity to the identities of such particular, unrepeatable events is, among other things, an important aspect of ethical life. In feeling regret or guilt for something one did, one's feelings are typically informed by a sense of how one might have acted otherwise—not just how one might act otherwise in other, similar circumstances, but rather how things might have gone, but actually did not, on that very occasion, which is now an unalterable part of one's life history.

More generally, I propose that a sensitivity to the identity of particular events matters to our ordinary understanding of the past as fixed, and of the contrast between actual and merely counterfactual histories.

One way of bringing this out is to consider some ways of representing time that are not sensitive to the identities of particular events. Here, for instance, John Campbell describes a creature who can represent periods of the circadian day, performing computations over these representations in order to time its behaviour:

The agent knows what phase the day is currently at: whether it is early, late or mid-morning, for example. And the agent may have discovered and stored information about what typically happens at various particular phases of the day—that breakfast is served at 10.00am on Forel's balcony, for example. And the agent may put this stored information to use in guiding action, as honeybees used to gather at Forel's balcony at breakfast time...So the agent arrives for food at the right time of day, and leaves shortly afterwards. Notice, though, that the agent so far has only the conception of time as (repeatable) phase. The agent does not draw, and makes no use of, the distinction between the morning of one day and the morning of another. (Campbell 2006, p. 2)

The creature of Campbell's example has quite a repetitive life. Most of the things it cares about, like getting food, do not differ much from one day to the next. In this respect the creature has, as Campbell puts it, 'no use' for, and indeed does not represent, the difference in numerical identity of one day and the next.

We can of course envisage a creature with a more complex life: for instance, an animal that hibernates might have reason to keep track not only of the time of day but also the time of year. This creature might therefore entertain a series of overlapping representations of its current position in different temporal phases,

and perhaps be able to integrate this information in various ways.

Intuitively, however, human time is different from phase time, and not only in the number and complexity of the time phases we keep track of. In remembering an event of one's past, one does not only think of it as occupying a position in various repeating phases, like time of day or season of the year; one thinks of it as having a unique temporal location that, as a matter of principle, one will not encounter again.

One expression of this conception is of course the use of a special symbolic code—for instance, a system of dates—to refer to specific events or their locations. Standard Western chronology exploits properties of the integer line in way that allows for any arbitrary terrestrial event to be assigned a unique temporal address. But, arguably, reference to events as uniquely located has a more fundamental role in our thinking about time, one that is independent of the use of any chronological system with this particular property. A more basic expression of a sensitivity to particular events, I suggest, is just the recognition that past events, having happened, will never come again, and thus are forever inaccessible to intervention—in contrast to future events, at least some of which are potentially within the reach of our agency and therefore open to deliberation.¹

To see this, reflect that a creature that represents only phases would not be able to grasp the fixity of past events. This is because it fails to represent the numerical identities of distinct events, and so fails to represent the distinctness of past and future instances of a phase. The creature in Campbell's example might represent breakfast as being over; but soon breakfast can be coming up again. These are, in reality, distinct events of breakfast, but the creature does not represent them as such. Insofar as the bee is capable of affecting when breakfast occurs, then, it simply figures in its cognitive life as something it can affect by acting appropriately at the

¹The neatness of this contrast would be disputed by defenders of the possibility of backwards causation, as well as those who hold that the proper objects of deliberation are just those events whose (subjective) probability is linked to our actions, even if they are causally prior to those actions—as in the case of Newcomb-like problems and retrospective petitionary prayer. A classic discussion of such cases is (Dummett 1964). I cannot get into these delicate matters here; I will just note that both of these possibilities are to some extent revisionary of the common thought that the past cannot be altered.

right point in the cycle. The idea that once breakfast is over, that particular breakfast is now gone for good and now cannot be affected, lies beyond the creature's grasp.²

This in turn limits the kind of understanding of its own actions available to the phase thinker, and the associated repertoire of emotional appraisals. The creature might, on a given occasion, note that its acting in such-and-such a way in a particular situation has led to a certain undesirable consequence, and thereby refrain from acting that way in similar situations in the future. To that extent, we might be tempted to say that the creature regrets acting as it did, and to attribute to it the counterfactual thought, 'Had I not done such-and-such, that outcome would not have occurred.' But its grasp of this counterfactual may not amount to anything more than simply the association of a generically bad outcome with the generic pairing of a type of action in a type of situation.

By contrast, an important species of human counterfactual thought, and in particular the counterfactual assessment of our own past actions, is temporally specific. The thought that if one had done X in past circumstances C, Y would have come about, is not equivalent to the thought that if one were to do X in relevantly similar future circumstances C*, Y would come about—although they may be inferentially linked, and the same sorts of evidence may bear on both judgments. Their non-equivalence does not, moreover, turn essentially on any kind of causal indeterminism, or on the empirical impossibility of reproducing precisely the same circumstances in the future, but simply on the basic logical point that they concern numerically distinct occasions, one in the past and one in the future. (Or, more carefully, the past-directed conditional concerns some particular actual past occa-

²A reviewer at this point wondered how the phase thinker can really distinguish past and future at all if, from its point of view, every event is always both past and future—in which case it is hard to see how the thoughts 'X is past' and 'X is future' could have distinct psychological functional roles. One solution would be to associate distinct functional roles not with the bare labels 'past' and 'future', but with quasi-metrical representations of the degree of pastness or futurity of an event, perhaps as an analogue magnitude, so that (the analogue equivalent of) the thought 'breakfast was one hour ago' has a distinct functional profile from the thought 'breakfast is in an hour'. While this would solve the immediate problem of distinguishing the functional roles of past and future, though, the fact that the same event can be represented as to some degree both past and future shows that the creature in a deeper sense systematically conflates the past and the future, and to this extent cannot be credited with a grip on the idea that past events are fixed and unalterable.

sion, whereas the future-directed one concerns any future situation meeting the conditions specified in the antecedent.) Our grasp of this difference means we can regret our past actions in a particular way—one that is essentially distinct from any resolution to act differently in relevantly similar future circumstances, but rather is informed by the recognition that the time for acting otherwise is now gone for good.

Similarly, our grasp of the numerical difference of past and future events has implications for the range of attitudes we can take towards our future actions. In deliberating over the future one may simply be trying to bring about some generically desirable outcome, like finding food, or avoid some generically bad one, like injury or discomfort. This kind of deliberation is plausibly something a phase thinker would also be capable of. But one can also deliberate over such things as what kind of person to be, or what kind of life to have. And this kind of deliberation, arguably, involves a recognition of the future as a distinct domain from the past, and of the idea that future events differ from past events, with respect to the orientation of our agency, in being still open and up for grabs.³

2 SPACE AND THE INDIVIDUATION OF PARTICULARS

I have indicated one significant way in which the individuation of particular events matters to us in ordinary life. The question I want to consider in the rest of this paper is what the underlying structural features of our thinking about time must be in order for the notion of a temporal particular to play this role. This will eventually lead into questions of the underlying metaphysical commitments of ordinary temporal thought and experience.

Events, as I have been using the term, are particular occurrences individuated by their temporal locations, or by their temporal relations to one another.⁴ (This

³Thanks to a reviewer for encouraging me to consider our future-directed attitudes at this point.

⁴Which are more fundamental, locations or relations? As will emerge from the discussion of the Russell below, really it is the individuating role of relations that is my focus here. However, this is not necessarily to commit to the claim that all facts about temporal and spatial locations, or about the identities of particular located items, can be reduced or translated to facts about relations between them. For discussion of the difficulties for this claim, see (Hawthorne and Sider 2002). The point

is not the only way of understanding the term: for instance, in probability theory, an 'event' is something that can happen more than once, like a coin landing heads.) A natural way of understanding this is by way of analogy with the idea of spatial particulars as individuated fundamentally by their spatial locations and relations; and this raises a rich set of questions concerning the extent to which a grasp of temporal relations might be capable of mirroring the foundational role of spatial relations in the individuation of particulars.

An instructive place to start here is a discussion by Russell of the notion of a particular. Russell poses the question how spatiality is involved in the individuation of distinct instances of the same sensible quality. He considers the case of visually perceiving two patches of white, at different places in the visual field:

Let us suppose...that within one field of vision we perceive two separated patches of white on a ground of black. It may then be taken as quite certain that the two patches are two and not one. The question is: Can we maintain that there are two if what exists in each is the universal whiteness? (1956, p. 116)

Saying that they are distinguished by the respective spatial position is not very satisfactory, Russell thinks, since this just pushes the problem of individuation back onto some more obscure entities, spatial positions. Even if absolute spatial positions in some sense exist, they cannot be perceived directly, and so the distinctness of their spatial position cannot be what grounds the judgment of the distinctness of the two white patches.

In some cases, perhaps, the two patches might be distinguished with reference different qualities they are 'co-present' with. For instance, perhaps one patch is round and the other is square; since roundness and squareness are incompatible, they must be distinct. But,

It is obvious, however, that this method of distinguishing the two patches is altogether inadequate. The two patches are just as easily distinguished if both are square or both are round. So long as we can see both, no degree of likeness between them causes the slightest difficulty in perceiving that there are two of them. The difference of shape, whether it exists or not, is not what makes the patches two entities instead of one. (pp. 116-117)

I take from Russell in what follows, that detecting relations plays a key role in the cognitive project of individuating particulars, can be separated from this reductionist claim. Thanks to a referee for pressing me to clarify this.

He then considers the different possibility that such particulars are individuated by the different relations they stand in. He writes:

It may be said that the two patches are distinguished by the difference in their relations to other things. Suppose a surface of black with a small white space in the middle...Suppose we have another white patch, of exactly the same size and shape, entirely surrounded by red. Then, it may be said, the two patches of white are distinguished by differences of relation, since one is surrounded by black and the other by red. But if this ground of distinction is to be valid, we must know that it is impossible for one entity to be both wholly and immediately surrounded by black and wholly and immediately surrounded by red. I do not mean to deny that we do know this. But two things deserve notice—first, that it is not an analytic proposition; second, that it presupposes the numerical diversity of our two patches of white. (p. 117)

Russell's point here is that there is no logical proof from the premise that the universal whiteness stands in the wholly-surrounded-by relation to two different colours to the conclusion that there are two numerically distinct patches of white. Nevertheless, we do take ourselves to know that one single thing cannot bear this same relation to two distinct things—and, Russell claims, we can know this on the basis of visual inspection. More generally, he says, we know of the spatial relations of perceived objects that,

They, or some of them, must be asymmetrical, *i.e.*, such that they are incompatible with their converses: for example, supposing “inside” to be one of them, a thing which is inside another must not also be outside it. They, or some of them, must also be transitive, *i.e.*, such that, for example, if x is inside y and y is inside z , then x is inside z —supposing, for the sake of illustration, “inside” to be among fundamental spatial relations...It follows that some at least of the fundamental spatial relations must be such as no entity can have to itself. It is indeed self-evident that spatial relations fulfil these conditions. But these conditions are not demonstrable by purely logical considerations: they are synthetic properties of perceived spatial relations. (pp. 117–118)

Russell thereby connects the individuation of particulars as numerically distinct with the construction of a system of spatial relations which satisfy the relevant ordering properties.⁵ Understanding the question of the individuation of spatial par-

⁵There is an unclarity in Russell's presentation. The idea that particulars are individuated by their standing in mutual ordering relations is different from, and not obviously connected to, the idea canvassed in the previous quotation that particulars might be distinguished by their having incompatible relational properties, like ‘being wholly and immediately surrounded by red’ and ‘...by black’. From Russell's surrounding discussion it is clear that he takes the former idea, of a system of transitive, irreflexive and asymmetric relations, to be the more fundamental.

particulars in this way, he claims, is an improvement on framing it in more intuitive terms:

The essential characteristic of particulars, as they appear in perceived space, is that they cannot be in two places at once. But this is an unsatisfactory way of stating the matter, owing to the doubt as to what a “place” is. The more correct statement is that certain perceptible spatial relations imply diversity of their terms; for example, if x is above y , x and y must be different entities. (p. 121)

There are a few doubts that might be raised here about the generality of Russell’s conclusion. Arguably, we are not warranted in taking the spatial relations in question unrestrictedly to have the ordering properties in question. For instance, the relation of *being inside* is in fact not necessarily either asymmetric or irreflexive: neither property holds on the surface of a torus. Similarly, neither are relations of *left* or *above* on the surface of a sphere. We can imagine discovering that our world in fact has a non-standard topology, so that the ordering properties of spatial relations are different from what we had supposed.⁶

Nevertheless, Russell can be read here as making a more limited point about the role of spatial awareness in the individuation of particulars. In the case of the two white patches, recognising the spatial relations they stand in is clearly sufficient for recognising their distinctness. More generally, the way we individuate spatial particulars exploits the fact that we operate with an array of relational spatial notions—‘to the left of’, ‘inside’, ‘above’, and so on—that can be applied simply on the basis of visual observation, and, under normal conditions, ‘imply diversity of their terms.’⁷ These relations, moreover, can be understood to generate an order: a system of relations within which the particular relata of those relations are each located.

⁶See, for instance, (Reichenbach 1958)’s vivid descriptions of adventures in non-Euclidean space.

⁷A reviewer suggested that individuating spatial particulars also requires a grasp of the relevant logical properties, such as asymmetry, transitivity, etc. While I am sympathetic to this suggestion, some might regard it as too demanding. Arguably, the visual system creates distinct object files partly on the basis information about objects’ (relative) spatial locations, thereby exploiting the logical properties of spatial relations to individuate objects, but without having to represent those logical properties. It might perhaps be suggested that conceptual (as opposed to perceptual) individuation of particulars requires a grasp of these logical properties. However, again, I think it is arguable that a thinker could simply use concepts such as ‘to the left of’ to individuate objects that stand in those relations, without grasping in full generality the higher-order concepts of asymmetry, transitivity, and so on. But this is a delicate matter that warrants further discussion.

Now, in line with Russell's conclusion, a natural suggestion would be that the individuation of events in time is parallel to the individuation of particulars in space in just this respect: that recognising a temporal relation, such as temporal precedence, between two events is normally sufficient for recognising their numerical distinctness. For example, in recalling an event in episodic memory, one thereby recognises it as prior to, and thereby distinct from, any similar such events in the future. Connecting this with the discussion of the previous section, the thought would be that the understanding of the past as fixed, and an associated grasp of historical counterfactuality, is coeval with a grasp of temporal relations as generating an order among events.

But someone might insist that our understanding of time is fundamentally not like this. It might be thought, rather, that the notion of something's being past is a primitive and irreducible matter. The difference between, say, its raining and its having rained is a difference in kind of fact, rather than a difference in the location of a particular within a framework. This primitive tense theorist would likewise object to the idea that we do discern relations between events that imply the relations are distinct individuals. The idea that spatial experience presents us with relations among particulars is fairly intuitive. By contrast, the primitive tense theorist may hold that experience does not really present us with temporal relations at all, but only with what is happening now.⁸

The challenge for this primitive tense theorist is to say what it is about our thought, and in particular our thought about the past, that distinguishes it from that of the phase thinker introduced in the last section. The phase thinker can sort events according to their status as past, present, or future: for example, Campbell's creature is able to represent breakfast as coming up soon, as happening now, or as having just happened. But, as we have seen, these are not categorisations of a

⁸This claim might of course be disputed: after all, there is a perceptible difference between a sequence consisting of red flash followed by a green flash, and a sequence consisting of a green flash followed by a red flash, which on the face of it is a difference in temporal relation. The tense theorist will have to construe this difference as one of tensed contents, rather than temporal relations. Adjudicating this disagreement on phenomenological grounds is a tricky matter—for some discussions see, for instance, contributions to (Phillips 2017). In what follows I will focus on the cognitive role of temporal relations for individuating particulars, setting aside questions of their presence in experience.

particular event: the creature can represent breakfast now over, now as coming up. Thus, the change shift from present to past does not, from the creature's point of view, represent the fact that that very event can now never be altered or revisited, but at most that now is not the time to do something about breakfast.

Similarly, the primitive tense theorist's ideology allows that the same event—or rather, the same tensed proposition—can be correctly represented as both past and future. So, like the phase thinker, this theorist does not understand the past and future as distinct domains comprising different sets of individuals. In the next couple of sections, I will develop this challenge for the primitive tense theorist, and explore the resources they have for meeting it.

3 EVENTS AND TENSED TRUTH

We can give the primitive tense theorist a name. The idea that temporal structure is articulated fundamentally in terms of the tenses, rather than relations between particulars, finds its clearest expression in the work of Arthur Prior. At the heart of Prior's approach to time is the idea that tense is deeply analogous, both formally and substantively, to modality.

The basic elements of temporal thought or discourse are *tensed propositions*—propositions which may change in truth-value over time, and do not need to specify a particular time in order to receive a complete evaluation as simply true or false. For instance, 'The nectarine is ripe' expresses a complete proposition, one that goes from being simply false to simply true as the nectarine ripens. Temporal information is expressed by means of sentence-level operators that shift the evaluation of a tensed proposition to another time, in something like the way that modal operators shift the evaluation of a sentence to merely possible worlds. Prior explains:

If [an] expression constructs a sentence out of one other sentence it is an adverb or an adverbial phrase, like 'not' or 'It is not the case that' or 'allegedly' or 'It is alleged that', or 'possibly' or 'It is possible that'...I want to suggest that putting a verb into past or future tense is exactly the same sort of thing as adding an adverb to the sentence. 'I *was* having my breakfast' is related to 'I *am* having my breakfast' in exactly the same way as 'I *am allegedly* having my breakfast' is related to it, and it is only an historical accident that we generally form the past tense by modifying the present tense, e.g. by changing 'am' to 'was' rather than by tacking on an adverb. (2003a, pp. 12–13)

One way of understanding the metaphysical territory staked out by Prior's approach to time is in terms of the idea that time is fundamentally unlike space because time, unlike space, is not a dimension of particularity. Rather, tense, like modality, expresses a way, or mode, for something to be the case.

Prior often expressed his view of time in the formal language of tense logic, which he developed with colleagues throughout his career. Tense logic is one of a family of intensional logics which work by enriching standard extensional logics with a family of non-truth-functional sentence-level operators. The basic elements are:

- A stock of propositional variables $p, q, r...$
- The truth-functional connectives $\sim, \rightarrow, \&, \vee$.
- The four tense operators P, F, H, G .⁹

The standard semantic approach to intensional languages of this sort uses a notion of truth with respect to an index of evaluation (for instance, a possible world), and treat the operators as shifting the point of evaluation of the embedded sentence. In this respect the only substantive difference between tense logic and ordinary modal logic is that there are two sets of operators, corresponding to two different 'directions' in which the evaluation might be shifted. Intuitively, the points of evaluation are to be interpreted as times, and the two 'directions' are interpreted as respectively earlier and later.

However, Prior stands out among intensional logicians for the particular metaphysical importance he attached to the object language itself, as a perspicuous representation of the logical structure of reality, and the rigour with which he insisted on this point. Accordingly, he was typically ambivalent about the standard semantic approach as applied to tense logic. Prior tended to insist that, while the semantic approach might be useful device for exploring the language's formal properties, it should not be mistaken for a genuine interpretation; only the object language of tense logic itself is a truly perspicuous representation of the reality of time.

⁹In what follows, I consider just the propositional language of tense logic, and completely ignore any of the complications arising from combining tense logic with the logic of quantification.

These points make Prior's philosophical use of tense logic a particularly apt case study for a view of time on which there are, fundamentally, no temporal particulars. The crucial feature of tense logic is that its expressions are all *general with respect to time*. A tense-logical proposition can be now true, now false, now true again; hence, it expresses something repeatable. Suppose, for instance, that it is currently raining, and it also was raining at some time in the past. That is,

(1) It is raining, and it was the case that it is raining.

Or, in tense logic,

(2) Rain & $P(\text{Rain})$

Are these the same raining event or are they distinct? From the point of view of the tense primitivist, the question has no answer: all we can say is that raining is in the past of raining. But being in the past is not a relation that implies the diversity of its terms, but rather a way for something to be the case.

This way of putting the point might seem counterintuitive. A common way of understanding tense logic is in terms of the idea that it represents temporal reality from a particular temporal point of view—that of the present. And this might seem to be enough to provide a way of distinguishing events in terms of their temporal relations to the present. That is, although the propositions of tense logic are in themselves general, they are used in a way that expresses information relative to the time of utterance or evaluation, and thus instances of the same tense-logical proposition can be distinguished relative to that time. For instance, (1) above has it that rain is occurring in the present, that is, at the time of evaluation; and also at a non-present time. Similarly, the statement

(3) It was sunny, and it will be sunny

distinguishes two instances of it being sunny as, respectively, past and future. So although both (1) and (3) involve a repetition of the same object language expression, nevertheless their truth would seem to require two distinguishable events happening at different times.

However, this interpretation rests on a subtle confusion, one that is invited by the persistent temptation to understand the notion of tensed truth in terms of that of truth at a time. (3), for instance, says that it is sunny both in the past and the future. But the judgment that there are two distinct events here requires the additional assumption that the same time cannot be both past and future. And not only is this not asserted by (3); the very idea of particular times at which these instances obtain is an additional bit of ideology, one that finds no direct expression in the object language.

Prior was very clear on this point. He emphasised that the idea that tense logic is centred on the present is at best an expedient way of speaking and at worst seriously misleading. Tense logic does not express a view on which the present is one privileged perspective among a multitude of other perspectives; rather, it expresses a view on which the notion of the present is *redundant*. Elsewhere he explains:

...the reality of the present consists in what the reality of anything else consists in, namely the absence of a qualifying prefix. To say that Whitrow's lecture is past is to say that *it has been the case* that Whitrow is lecturing. To say that Scott's lecture is future is to say that *it will be the case* that Scott is lecturing. But to say that my lecture is present is just to say that *I am lecturing*—flat, no prefixes. The pastness of an event, that is to say its having taken place, is not the same thing as the event itself; nor is its futurity; but the presentness of an event *is* just the event. (Prior 1972, p. 322)

In other words, it is a mistake to think of the 'present' of the tense primitivist as indexically picking out one time among others; and it is similarly a mistake to think of the tenses as referring to other, non-present times. Thus, although whatever is asserted outside the scope of any tense operator is thereby, in a certain trivial sense, asserted to be present, it is not the case that whatever is asserted within the scope of *P* or *F* is to be understood as referring to something which is *not* present.

The analogy with modal thinking is illuminating here. Suppose it is raining, and one thinks, "It might not have rained; but even then, rain still would have been possible." In other words, it is raining, but it might have been that it wasn't raining although it might have rained; or, in modal logic,

$$(4) \text{ Rain } \& \diamond (\sim \text{Rain} \& \diamond \text{Rain})$$

In this envisaged possibility, is the merely possible rain the same rain as the actual rain? Intuitively, the question does not admit of a determinate answer. The mere possibility of rain is not a particular raining, and on the face of it there is no clear sense to questions about the identity of this possibility.

Now, of course we can and do make distinctions among possibilities easily enough. Light drizzle is a different possibility from a thunderstorm; an hour of rain is a distinct possibility from just a minute of rain; and so on. We can go on and on making further distinctions between possible scenarios—in principle, right up to the point of a maximum, fully determinate description of the whole universe.

However, even at this point, it is not clear we are really entitled to speak of distinct *particular* possibilities, with the associated distinction of numerical and qualitative identity. Rather, we distinguish possibilities just by further specifying what they are like. Ordinary modal thinking does not have much use for the idea of descriptively identical, but numerically distinct, possibilities.¹⁰ We can, if we like, picture possibilities as locations in a modal ‘space’, ordered according to some notion of relative similarity or dissimilarity. Arguably, ordinary counterfactual reasoning involves making some such judgments of relative similarity among possibilities. But the idea of a ‘space’ here adds nothing beyond a way of organising possibilities according to their intrinsic features; it is not a framework for the individuation of possibilities in the way that physical space is a framework for the individuation of

¹⁰A reviewer invited me to consider here David Lewis’s argument against linguistic ersatzism about possible worlds on the grounds that linguistic or propositional specifications of possible worlds are unable to allow for distinct but indiscernible possible individuals (Lewis 1986, pp. 157–158) I do not know exactly what to make of this argument of Lewis’s, but I have a couple of points to make about its relation to the present dialectic: i) Lewis is principally concerned with the problem of distinguishing between indiscernible individuals *within* a single possible world; by contrast, he considers the problem of indiscernible worlds to be ‘harmless’, precisely because there is no particular reason why a theory of modality should rule for or against the possibility of indiscernible worlds. And ii) Lewis is of course arguing that our ordinary modal judgments commit us to a certain conception of the modal domain, namely as a framework of particular ‘big concrete objects’, in contrast to a traditional conception of modality as a primitive modification or mode of a state of affairs. But the tense primitivist holds that tense is fundamentally analogous to modality *as traditionally conceived*, rather than as conceived by Lewis. So to say that modal thought according to Lewis has a certain feature is not to say that temporal thought according to the tense primitivist has a corresponding feature. And indeed, as I emphasise below, it remains a signal limitation of the tense primitivist’s conception of time that times—maximal constellations of tensed truths—conform to a version of the Identity of Indiscernibles.

particular objects. By contrast, Russell's point about space was that our judgments of the distinctness of particulars, like the two white patches, are often not based on any intrinsic difference, but rather on recognising an 'external' relation between them, namely a spatial relation.¹¹ It is this role for spatial relations that does not seem to have any analogue in ordinary modal thought.

Similarly, the tense primitivist might distinguish events or times after a fashion in terms of what else is the case at those events. For example, the statement,

- (5) It was the case that it is raining and sunny, and it was the case that it is raining and not sunny

distinguishes two 'instances' of rain by specifying a further bit of information, whether or not it was sunny. So in a sense we can say that these are two distinct events. Taken to the limit, this approach can in principle be extended to uniquely individuate particular temporal locations, by identifying a time with a maximal consistent totality of tensed truths. Quantifying over tensed propositions, an instant-proposition i may be defined as any proposition satisfying the following:

$$(6) i \vee \exists p i \vee F i$$

$$(7) \forall p ((i \rightarrow p) \vee (i \rightarrow \sim p))^{12}$$

It should be clear from our discussion of Russell what the limitations of this approach are. Although we do sometimes distinguish events in terms of what else is happening contemporaneously with them, this is not essential to the idea of them being distinct individuals. Rather, as with the spatial case, our ordinary conception has it that events in time are often distinguished from one another just on the basis of temporal order relations between them, regardless of any intrinsic difference between them. If, for instance, one sees two identical flashes of light one after the other, simply perceiving that one precedes the other is, under normal conditions,

¹¹By an external relation I mean one that does not supervene on the intrinsic properties of the relata. For discussion see (MacBride 2016).

¹²This approach, which resembles the strategy of identifying possible worlds with maximal propositions, is developed by Prior in his later papers, collected in (Prior 2003b); see also (Meyer 2009) for a related approach.

sufficient grounds for recognising their distinctness, irrespective of whatever else may be the case contemporaneously with each flash. But we have seen that the tense primitivist's basic notions do not allow for events to be distinguished in this way. In making tense fundamentally analogous to modality, the tense primitivist thereby downplays the central place for relations of order, over and above qualitative variation, in our thinking about time.

This so far is just to say that the basic elements of tense logic, as interpreted by Prior, do not provide materials for talking about particular events or their temporal locations. Nevertheless, it might be that there are ways of approximating talk about particular times and events that emerge when we look at the logic as a whole. The next section will investigate this possibility.

4 TENSE LOGIC AND TEMPORAL STRUCTURE

In the last section I suggested that considerations of order play only a quite peripheral role in ordinary modal reasoning. Nevertheless, as is well known, the axioms of modal logic can be thought of as corresponding to conditions on a relation of relative possibility. In the possible worlds model theory, each modal axiom characterises a class of models, specifiable in terms of an ordering property of an accessibility relation on possible worlds, for which the relevant set of axioms is sound and complete. This encourages the idea that, just as the axioms of modal logic characterise conditions on the relation of relative possibility, the axioms of tense logic can be thought of as expressing postulates about the topology of time, by characterising ordering properties of the earlier-later relation.¹³

This looks like precisely what the tense primitivist needs to appeal to in order to recover talk of particular events. The axioms of tense logic can be used to construct an order of temporal relations, and events can then be distinguished with reference to their position in the temporal order.

One immediate point to note is that there is no axiom schema corresponding to the property of irreflexivity, and hence no general principle of tense logic which

¹³See, for example, (Meyer 2013; Newton-Smith 1980), as well as ch. 2 of (Prior 1967).

is interpretable as expressing that temporal order is irreflexive.¹⁴ The same goes for the property of antisymmetry. In other words, there is no general way of stating in a tense-logical idiom the idea that temporal priority is a relation that, in Russel's terms, implies diversity of its terms. In the light of the above discussion, this should be entirely unsurprising: as we have seen, one of the distinctive marks of the tense primitivist's conception is that such ordering relations do not play a central role.

A symptom of this is that in tense logic there is no expressible difference between a circular model of time, in which every event is both earlier and later than itself, and an infinitely repetitive one, in which the same sequence of events happens over and over again. Any tense-logical formula which is true on the one model will be true on the other. Tense logic therefore cannot distinguish the cosmic possibilities of cyclical time and eternal linear recurrence.

The tense primitivist might not be very concerned by this. Intuitively it is quite hard to hold the possibilities of linear and cyclical recurrence apart anyway, and it does not seem that our ordinary conception of time is very strongly committed to distinguishing them. One might indeed argue that these are not in fact conceptually distinct scenarios, and it is a virtue of the tense primitivist's conception that it does not spuriously multiply possibilities. In other words, the tense primitivist's position will be that, insofar as we operate with some notion of temporal order, the logic of this notion does not strictly require it to express an irreflexive and antisymmetric relation.

One might instead look to the idea of transitivity to capture what is distinctive about our conception of the time in contrast to the phase thinker's. Transitivity is, arguably, essential to any conception of temporal succession as forming an order. Unlike irreflexivity, it has a corresponding pair of axioms, the S_4 -like:

(PS₄) $H\varphi \rightarrow HH\varphi$

(FS₄) $G\varphi \rightarrow GG\varphi$

¹⁴Strictly speaking what this means is that the modal system which is sound and complete for the class of irreflexive models is just the basic system K, which is also sound and complete for the class of all models. For explanation of this, see (Cresswell and Hughes 1984, ch. 2).

Focusing just on the past-directed axiom **PS4**, what this intuitively says is that the past of the past is also the past.

Meanwhile, instead of the B axiom of symmetry, we can postulate the following ‘mixing’ axioms, corresponding to the requirement that earlier and later are converses:

(PB) $\varphi \rightarrow GP\varphi$

(FB): $\varphi \rightarrow HF\varphi$

The past-directed version of these intuitively says that what is happening now will always have happened. Moreover, PS4 and PB together imply the S5-like:

(PS5) $P\varphi \rightarrow GP\varphi$

i.e. that what is past will always remain past. Together, then, these appear to express a conception of the past as *cumulative*: as time passes, more and more past facts come into being, and none gets discarded. This seems like an apt enough expression of the idea that the past is fixed and unchangeable—what has been done will always remain done, purely as a matter of the logic of time.¹⁵

We might bring out the importance of this conception again by contrast with our phase thinker. On day 0, let us suppose, the creature could in some sense remember the sunset of day -1 and anticipate the sunrise of day 1, and hence represent them as, respectively, past and future. And on day -1 it could remember the sunset of day -2, and so represent that as past. Yet on day 0 it has not only forgotten the sunset of day -2; it cannot so much as make sense of the idea that there are two distinct past sunsets, one in the past of the other. The creature’s memories and anticipations are like a file of, respectively, recent and upcoming events, and

¹⁵This raises the tricky question whether the future-directed analogue $F\varphi \rightarrow HF$, that what will be the case was always going to be the case, should be thought of as expressing a form of fatalism about the future. Assuming it does generates a trilemma for the primitive tense theorist here: Either we accept both **PS5** and its mirror image **FS5**, thus embracing fatalism; or we reject both, thus denying the fixity of the past as a logical principle; or we accept **PS5** and reject **FS5**, thus complicating the logic of time. Many discussions of future contingents in the setting of tense logic are essentially explorations of the third horn of this trilemma. By contrast, in what follows I suggest we reject the trilemma by denying that **PS5** is an adequate expression of our intuitions about the past.

these files get *updated* with new information as time passes. As far as the creature is concerned, there is no more to the notion of the past than just those events that are at least potentially within the scope of its memory file. Because the creature's memory only reaches back to the events of the recent past, it thus lacks any conception of the past as cumulative. From its point of view, new occurrences replace the previous events of the past; the past gets dragged along in the wake of the present, like the tail of a comet.

Now, we noted above the tempting thought that our cumulative conception of the past can be captured in terms of an implicit acceptance of a logic of tensed truth as governed by the axioms **PS4** and **TB**, and amounting to recognition of the transitivity and duality of earlier-later. Yet, as I will argue in the remainder, this again rests on a subtly question-begging interpretation of the tense operators.

First, note that the phase-using creature's representations of events need not, at any one time, violate the principle **PS4**. For example, we supposed that on day 0 it represents as past the sunset of day -1. Suppose additionally that it represents the sunrise of day 0 as past, and represents that sunset of day -1 as in the past of—i.e., as prior to—that sunrise. So in this case the past of the past is represented as past, in accordance with **PS4**. On the other hand, in the case we considered above, on day 0 it represents the sunset of day -1 as past, and on day -1 it represents the sunset of day -2 as past, but on day 0 it does not—perhaps could not—represent as past the sunset of day -2. But, by the same token, it also fails to represent the sunset of two days ago as in the past of the past: it just does not represent it at all. In general, the failure of the creature to be able to keep track of how the past stays past does not mean that it ever actually represents something as in the past of the past but not as past. The unstable and changing character of its conception of the past is something that can be appreciated only from an external theorist's description of how the creature's tensed representations evolve over time.

The more general point here is that the principle embodied in **PS4**, when interpreted internally, from the time-bound perspective of the creature whose thought is being modelled, is not adequate to capture our intuitions about the constancy of the past across multiple temporal perspectives. A more thorough illustration of

this point can be provided by considering a semantics for tensed representations in terms of a branching structure of alternative histories.

Here is one way this can go.¹⁶ Instead of taking truth to be relative to a time t , we instead take truth to be relative to a pair $\langle t, h \rangle$ of a time and a *history*, such that $t \in h$. (A history for these purposes is just a linearly ordered set of times.)

The standard semantic clauses for the tense operators are then restricted to times within the history of evaluation. For instance:

P: Pp is true at $\langle t, h \rangle$ iff, for some $t' \in h$ such that $t' < t$, p is true at $\langle t', h \rangle$.

We can then introduce a further semantic notion of *overall* truth, or Truth (still relative to a time), such that a formula is True at $\langle t, h \rangle$ if it is true at all $\langle t, h' \rangle$ such that $t \in h'$ (the h -index for Truth is vacuous.) Determinate truth, rather than truth relative to a possible history, is what agents are fundamentally concerned with: in the context of linguistic communication, Truth is linked with assertability; and a belief about the past will count as correct at a given time just in case its content is True with respect to that time and any history containing that time.

It is important to be clear about what the structure and its branches represent. The structure is intended to model how a thinker's tensed beliefs evolve over time. But this is not to say that the thinker themselves represents time as having a branching structure. The structure as a whole does not directly represent the agent's beliefs about the structure of time. Rather, the fact that a given node in the model belongs to multiple different histories represents a conception of time on which the truth about non-present may be underdetermined.

This approach is most familiar as a way of modelling the idea that the future is open in such a way that future truth is indeterminate.¹⁷ Similarly, a structure with backward as well as forward branches models of tensed truth for a conception

¹⁶What follows is a sketch of the kind of supervaluational branching-time semantics developed by (Belnap and Green 1994; Dummett 1981; Thomason 1970) and others. There are other alternative implementations of a branching-time semantics, for instance the many-valued approach of (Łukasiewicz 1968). I am focusing on the supervaluational approach because it has the important feature of preserving the logic of a standard interpretation of the operators.

¹⁷For discussions of this approach to the open future, see the contributions to (Correia and Iacona 2013).

of time on which the past is also less than fully determinate. For example, this is one natural way of modelling the truth about the past according to a verificationist for whom past truth consists just in the existence of present evidence.¹⁸ On this interpretation, we can think of the branches of the model which overlap at any given time as corresponding to what is consistent with the evidence available at that time. As evidence can be degraded or destroyed over time, the verificationist holds that what are now truths about the past will no longer obtain in the future. The expression of this in the model is the appearance of new backward branches as we move forwards in the model, so that what was at one time true on all past histories, hence True, will at a later time be true on only some histories, and so not True. Since they allow that past truths can cease to obtain, there is a clear intuitive sense in which the verificationist whose conception of time is modelled by this structure lacks what I have called a cumulative conception of the past: they allow that facts about the past might disappear as time goes by.

Despite this, the semantics for the tense operators just sketched leaves their logic untouched. Classical logic and all of the standard tense-logical axioms hold on each history, and so their instances are all True. In particular, given that each history is a transitive, linear order, each instance of **PS4** and **PS5** comes out True. The verificationist thus accepts all the same tense-logical principles as a realist about past truth; yet, as we just saw, their conception of the past allows that truths about the past can be obliterated as evidence is destroyed. The problem is that this difference does not show up in the pattern of tensed truths that they accept at any one given point in time, but only in the *interplay* of tensed truth as they shift between temporal perspectives. This, however, is not capturable within the object language of tense logic itself, as it is fundamentally a calculus of deductive relations between tensed truths as apprehended from a single temporal perspective.¹⁹

¹⁸A classic exploration of verificationism about the past is (Dummett 1978); the next few paragraphs here recapitulate one strand of Dummett's discussion. An alternative interpretation is in terms of ontological presentism, and the view that past truth is grounded in what presently exists plus the (possibly indeterministic) laws of physics. For discussion of this last case, see (Dawson 2020; Markosian 1995, 2013).

¹⁹As we have seen, this is not a characterisation that Prior would have been entirely happy with, because he resists the idea that the present is one among many temporal perspectives. Nevertheless, it is still correct to say, even by Prior's lights, that tense logic does not model the interaction of tensed

To put the point more formally, as just noted, the formula $Pp \rightarrow GPp$ is valid on this semantics, hence True at any given time. What will fail to be the case, however, is the following: if Pp is True at t , then Pp is True at every t' such that $t < t'$. This is because, as we go to later times, these later times will have more backward branches; and p may fail to be true at one of these later-joining branches. The problem is to see how this can be recognised by the agent. At no time will the agent be prepared to assert, or in some other way display a commitment to, the statement ‘One day, *that* [i.e. some past occurrence] will not have happened any more’. From their point of view, time is always a fully linear order in which the past remains fixed.

We might try to bring this difference down into the object language by introducing into the logic a determinacy operator, which converts the Truth of a formula into its truth. That is, Dp is true at $\langle t, h \rangle$ iff p is True at $\langle t, h \rangle$. We can then trace entailments between truth and Truth in ways that reveal cross-temporal discontinuities. For instance, although $DPp \rightarrow DGPp$ is valid, $DPp \rightarrow GDPp$ is not. In general, prefixing each occurrence of a tense operator within a formula with D will not preserve validity.

The question is how we are to make sense of the meaning of the determinacy operator as anything more than a technical innovation. In a forward-branching model, with branches representing possible futures, an operator like D that generalises over the branches might be interpreted as expressing some kind of objective historical modality that evolves over time. We can make good enough sense of the existence of multiple forward branching points in the model as an expression of the ordinary idea that some things were once possible, but no longer are; and that in the future, some things which are now possible will no longer be possible any more. This kind of modality is plausibly connected with ideas about the direction of causality, and in particular with the idea that it is possible to interfere with or interrupt a causal process while it is going on, but not once it has concluded.

By contrast, it is much harder to make good sense of the idea of a possible past, or of a form of backward-looking modality which develops over time, possible pasts coming into being as time passes. We have no backward-looking analogue of truths from multiple different temporal perspectives.

ideas of prevention and interruption, by means of which different alternative pasts might newly become possible. The only viable intuitive interpretation of the different branches is rather the epistemic one on which the appearance of new branches reflect the destruction or degeneration of evidence about the past. Our D operator then will express the idea that what comes after it is implied or necessitated by present evidence, and its dual $\sim D \sim$ that what comes after it is compatible with present evidence.

This is at least a coherent way to interpret a temporal determinacy operator. Yet on this interpretation the possible failure of $DPp \rightarrow DGPp$ does not express any distinctive commitment of the verificationist, and in particular does not capture the idea that past truth can vanish. It rather simply expresses the truism that evidence can be destroyed over time, so that what is now implied by the present evidence may not be so implied in the future. What the verificationist is committed to, by contrast, is not simply this truism, but rather the more radical possibility that what is now true (or rather True) about the past may not be so in the future. And this commitment, I suggest, has no way of being expressed in the tense primitivist's ideology.

The immediate point here in reflecting on the verificationist's picture is that our realist convictions about the past cannot be captured solely in terms of any pattern of acceptance of tensed truth at a given time. The wider point is that these convictions concern not just how things appear from a single temporal perspective, but the way in which our thought spans multiple perspectives as time goes by. The verificationist, like Campbell's phase thinker, never commands a clear view of the temporal order in which their thinking is embedded—despite the fact that their reasoning about time may well respect the principles of tensed truth embodied in **PS4** and **PS5**. The tense-logical principles hold at each temporal perspective they occupy, yet fail to cohere to represent a single consistent temporal reality. There is no single actual timeline of events which each successive perspective is a perspective on, but rather a shifting succession of incommensurate timelines, the succession of which the thinker themselves is unable to represent. In an important sense, then, this thinker lacks a single stable conception of the past as a domain of events in the

actual world.²⁰

By contrast, the idea of a temporal particular at the centre of this discussion is that of something individuated by its location in an encompassing order of temporal relations. Attitudes that target temporal particulars, such as episodic memories or memory-based thoughts, locate their objects and their thinkers within that order.²¹ In having such attitudes, and maintaining them over time, we thereby orient ourselves in the order of events and their relations that is the temporal world. The suggestion that I have been developing in this paper is that our idea of the past as a domain of fixed and unalterable events is dependent fundamentally on our entertaining such particularistic attitudes, rather than the acceptance of general principles of tensed truth.

5 THE METAPHYSICS OF TIME AND TENSE

I have suggested that attitudes to particular events, individuated by their temporal locations, form a central component of our ordinary conception of time, specifically the past, and moreover play an important role in ethical life; and I have argued that this is fundamentally a different matter from entertaining attitudes to tensed truths, or to complexes of tensed truth. One question this raises is the extent to which our ordinary conception is committed to an ‘eternalist’ or ‘tenseless’ metaphysics, rather than a tensed metaphysics on which the tenses are fundamental.

Primitivists about tense, including Prior, have often presented tense logic both as a metaphysically perspicuous representation of temporal reality, and as an accurate formal regimentation of the structures underlying temporal discourse in natural language. The first of these makes it an apt expression of metaphysical views about the reality of tense. On the second of Prior’s aspirations, it is generally thought that Prior was wrong, and that the expression of tense in natural language

²⁰This conclusion mirrors the ‘thin red line’ problem from branching models of the open future, that there is no single actual course of events common to each point in the model, and so it is no longer coherent to talk of the ‘actual future’.

²¹This is not to suggest that singling out a particular event requires having some other way of picking out its location, like a date and time. It is rather that singling out an event makes it possible to frame questions about its temporal location in relation to one’s own, and to other events—even if these cannot be given an informative answer.

is not intensional in the way he thought. For the last few decades it has been usual for parties to this debate to avoid getting seriously involved with tricky empirical issues around the semantics of tense in natural language—issues which have come to centre around the explanation of phenomena such as the interaction of tense with indirect speech reports, anaphoric binding within tensed constructions, intersentential relations within longer stretches of discourse, the so-called 'double-access' sentence, and so on. These issues are recognised to be extremely complex and delicate, and generally best left to linguists and semanticists. Defenders of a tenseless metaphysics thus tend to emphasise that they are concerned with the nature of time itself, and not with the language of time.

Nevertheless, proponents of tensed metaphysics of time typically do not take themselves to be propounding an error theory. It would thus be a sensible move for a defender of a tensed metaphysics to maintain that whatever conclusions linguists draw about the semantics of tense—in particular, whether an adequate semantics for tensed language should be extensional or intensional—this in some way fails to be revealing of the deeper metaphysical commitments of ordinary thought about time.

If what I have argued is along the right lines, then this latter move is also problematic: the tensed picture is more radically revisionary of ordinary thought than is often assumed. This is not to argue directly against it, but it does undermine one of the principal motivations for adopting a tensed metaphysics of time.

Equally, this is not to argue that we are unequivocally committed to an eternalist picture of time, or that this commitment would be unproblematic. In particular, it may remain obscure how to reconcile our convictions about the openness of the future with the idea that past and future are just subdomains of a single overarching domain of events. This thought might lead one to suspect that there is a fundamental tension at the heart of our ordinary conception of time, between the idea of time as a system of located particulars and the idea of the difference between past and future as a modal difference between fixity and openness. I hope to have contributed to the exploration of these issues by clarifying the character of our commitment to the former of these.

REFERENCES

- Belnap, Nuel and Mitchell Green (1994). 'Indeterminism and the Thin Red Line'. In: *Philosophical Perspectives* 8, pp. 365–388.
- Campbell, John (2006). 'Ordinary Thinking About Time'. In: *Time and History: Proceedings of the 28. International Ludwig Wittgenstein Symposium, Kirchberg Am Wechsel, Austria 2005*. Ed. by Michael Stöltzner and Friedrich Stadler. De Gruyter, pp. 1–12.
- Correia, Fabrice and Andrea Iacona, eds. (2013). *Around the Tree: Semantic and Metaphysical Issues Concerning Branching and the Open Future*. Dordrecht, Netherland: Springer.
- Cresswell, M. J. and G. E. Hughes (1984). *A Companion to Modal Logic*. New York: Methuen.
- Dawson, Patrick (2020). 'Hard Presentism'. In: *Synthese* 198.9, pp. 8433–8461. DOI: 10.1007/s11229-020-02580-9.
- Dummett, Michael (1964). 'Bringing About the Past'. In: *Philosophical Review* 73.3, pp. 338–359.
- (1978). 'The Reality of the Past'. In: *Truth and Other Enigmas*. Duckworth, pp. 358–374.
- (1981). *Frege: Philosophy of Language*. 2nd ed. orig. 1973. London: Duckworth.
- Hawthorne, John and Theodore Sider (2002). 'Locations'. In: *Philosophical Topics* 30.1, pp. 53–76.
- Lewis, David K. (1986). *On the Plurality of Worlds*. Malden, Mass.: Wiley-Blackwell.
- Łukasiewicz, Jan (1968). 'On Determinism'. In: *The Polish Review* 13.3, pp. 47–61.
- MacBride, Fraser (2016). 'Relations'. In: *Stanford Encyclopedia of Philosophy*.
- Markosian, Ned (1995). 'The Open Past'. In: *Philosophical Studies* 79.1, pp. 95–105. DOI: 10.1007/bf00989786.
- (2013). 'The Truth About the Past and the Future'. In: *Around the Tree: Semantic and Metaphysical Issues Concerning Branching Time and the Open Future*. Ed. by Fabrice Correia and Andrea Iacona. Springer, pp. 127–141.

- Meyer, Ulrich (2009). 'Times in Tense Logic'. In: *Notre Dame Journal of Formal Logic* 50.2, pp. 201–19. DOI: 10.1215/00294527-2009-007.
- (2013). *The Nature of Time*. Oxford: Clarendon Press.
- Newton-Smith, W. (1980). *The Structure of Time*. Routledge & Kegan Paul.
- Phillips, Ian, ed. (2017). *The Routledge Handbook of Philosophy of Temporal Experience: Routledge Handbooks in Philosophy*. New York: Routledge.
- Prior, A. N. (1967). *Past, Present and Future*. Clarendon Press.
- (1972). 'The Notion of the Present'. In: *The Study of Time*. Ed. by J. T. Fraser, F. Haber and G. Muller. Springer Verlag, pp. 320–323.
- (2003a). 'Changes in Events and Changes in Things'. In: *Papers on Time and Tense*. 2nd ed. orig. 1962. Oxford University Press.
- (2003b). *Papers on Time and Tense*. 2nd ed. Oxford University Press.
- Reichenbach, Hans (1958). *The Philosophy of Space and Time*. orig. 1927; trans. Maria Reichenbach and John Freud, 1957. New York: Dover.
- Russell, Bertrand (1956). 'On the Relations of Universals and Particulars'. In: *Logic and Knowledge: Essays, 1901-1950*. orig. 1911. Macmillan.
- Thomason, Richmond H. (1970). 'Indeterminist Time and Truth-Value Gaps'. In: *Theoria* 36.3, pp. 264–281.