Temporal location of events in language and (non) persistence of the past

FABIO DEL PRETE

Abstract

The article reviews some analyses of temporal language in logical approaches to natural language semantics. It considers some asymmetries between past and future, manifested in language, which motivate the "standard view" of the non-reversibility of time and the persistence of the past. It concludes with a puzzle about the changing past which challenges the standard view.

Keywords: time, events, necessity of the past

1. Introduction

Languages like Italian have tensed forms of the verb which, in interaction with temporal adverbs, express the *location of an event in time*. For example, the italicized verb forms in the Italian sentences (1a)-(1c), in interaction with the temporal adverbs in parentheses, locate a particular event in the past, the present and the future, respectively:

   (1)   a.   (Ieri) il presidente dell'UCI *era* a Bruxelles.

               'The president of UCI *was* in Brussels (yesterday)'

          b.   (Questa mattina) il presidente dell'UCI *è* a Bruxelles.

               'The president of UCI *is* in Brussels (this morning)'

          c.   (Domani) il presidente dell'UCI *sarà* a Bruxelles.

               'The president of UCI *will be* in Brussels (tomorrow)'

The verb forms in question do this with respect to a *reference point*—typically, the moment at which the sentence is uttered. The past, present and future of (1a)-(1c) are, respectively, past, present and future with respect to *that moment*, which we may call "the *hic et nunc* of the utterance". This can be seen as the "zero point" from which every event is apprehended.

The property just mentioned makes tense, by and large, an *indexical* category: to determine its contribution to meaning, we must first identify the context of utterance. For example, suppose that you run into an utterance of (1b) – maybe you see an inscription of this message in a letter, or you foresee it through the powers of your imagination as a future utterance – but you are not aware of the *hic et nunc* of it, so that you don't know whether the sentence was uttered one year ago or will be uttered in a century from now. Until you don't identify the *hic et nunc* of this utterance, you won't be able to determine its full meaning (being now September 1st 2020, does the utterance mean that the UCI president was in Brussels in the morning of September 1st 2019? or does it mean that she will be in Brussels in the morning of September 1st 2120?).

Linguists have viewed the function referred to above as fundamental to tense. Accordingly, they have defined tense as the *grammaticalised expression of location in time*. Bernard Comrie expresses this view in the following passage:

> [For grammatical categories, including tense, w]hat one finds most typically is the choice of the speech situation as the reference point […]. As far as tense is concerned, then, the reference point is typically the present moment, and tenses locate situations either at the same time as the present

moment, or prior to the present moment, or subsequent to the present moment (Comrie 1985).

The location of events in time is fundamental to human experience. Our lives happen in time, as successions of events. An important function of language is to represent the events making up our personal lives and our public history by locating those events within an established temporal frame and by determining temporal relations (*before*, *after*, etc.) between them on the background of such a frame. In this function, language does an invaluable service to memory – both to "private" memory about events from our personal lives and to "public" memory about events from our common history. In this article I consider the subject of the temporal location of events from the perspective of natural language semantics. The more general question accompanying my inquiry is what we can learn about fundamental properties of time from a consideration of linguistic phenomena. The specific questions arising from my inquiry bear in particular on the (non-)reversibility of time and the (non-)persistence of the past.

The article is structured as follows. I begin with a recollection of analyses of temporal language in logical approaches to natural language semantics (section 2). Then I consider some asymmetries between past and future that become manifest in language, which raise the questions of reversibility and persistence, and I present the "standard view" of the non-reversibility of time and the persistence of the past (section 3). Then I consider a puzzle about the changing past that challenges the standard view (section 4). I conclude suggesting an open-minded view on the fundamental questions.

2. Logical approaches to the semantics of temporal expressions
Arthur Prior is the philosopher who laid the ground to logical approaches to the semantics of tense, by founding modern tense logic

(Øhrstrøm and Hasle 1995; Copeland 2020; Goranko and Rumberg 2020). Preliminarily, we are interested in Prior's philosophical ideas about time, which are particularly well-suited to think about history. It is not an accident that the first modern treatment of so-called "historical modalities" (modalities such as "In 1932 *it was possible* for Great Britain to avoid war with Germany; but in 1937 *it was impossible*"; Thomason 1984), which were formalized by Richmond Thomason in the '70s and '80s, sprang from Prior's tense logic and from his views about indeterminism and what was later called "branching time" (the view that, at any moment in history, the world has *only one past* but *several possible futures* branching off from that moment; see section 3.2.).

Prior was a *temporalist*: he thought that human thoughts are fundamentally tensed and that the temporal notions of *pastness*, *presentness* and *futurity* are not expressible in the objectivist, atemporalist terms of the notion '*time instant t_1 precedes / follows time instant t_2 in the series of times*'. He was also a *presentist*: he thought of the present as the fundamental perspective of human thought on reality, while he conceived of the past and the future, derivatively, as "displaced presents"—the *past is what has been present* and the *future is what will be present* (we'll see shortly that this derivative conception of the past and the future as displaced presents lies at the heart of Prior's treatment of the past and the future *qua* grammatical categories, i.e. grammatical *tenses*).

In these respects, Prior's ideas differed dramatically from those of Willard Van Orman Quine. In his *Word and Object*, Quine writes:

> Our ordinary language shows a tiresome bias in its treatment of time. Relations of date are exalted grammatically as relations of position, weight, and color are not. This bias is of itself an inelegance, or breach of theoretical simplicity.

Moreover, the form that it takes – that of requiring that every verb form show a tense – is peculiarly productive of needless complications, since it demands lip service to time even when time is farthest from our thoughts. Hence in fashioning canonical notations it is usual to drop tense distinctions (Quine 1960).

For Quine, tensed sentences such as (1a)-(1c) are a logically imperfect way to convey certain meanings and should be improved by regimenting the tensed expressions of everyday language in the objectivist terms provided by the mathematical ordering of times (according to the relation of precedence), accompanied by tenseless predications about such objectively ordered times. For example, (1a) should be regimented as (1a')[1]:

(1)   a′.   The president is in Brussels at time t_1 and t_1 precedes the time t_0 at which event e_0 happens (and t_1 is included by the day preceding t_0).

Prior's advice is radically different: we should take our tensed talk seriously for what it is, not trying to reduce it to something else. The use of tensed forms, far from being logically imperfect, is a legitimate and meaningful way of expression. Instead of regimenting our language to avoid such forms, we should rather devise a new logic capable of formalizing their meaning.

---

[1] In Quine's regimentations, the verb form "is" is to be taken as the *tenseless* present of mathematics, as it shows up in eternally true statements such as "2 + 2 *is* equal to 4", while the specification "at time t_1" makes explicit – once and for all – the temporal reference which is left implicit by the logically imperfect sentences of our ordinary language, thus making the propositions expressed by those sentences *eternally true* or *eternally false*. The event e_0 introduced in the regimentations represents the event of their utterance.

In the following sections I consider some logical approaches to temporal language, starting from Prior's approach, and then moving to refinements of the latter which were motivated by certain limits inherent to Prior's tense logic.

2.1. Priorian tense operators

In this presentation I use the graphical device of a line oriented from left to right to represent the temporal flow of events. The left part is associated with the past, the right part with the future. A designated point "t_0" is associated with the present – the time at which the utterance happens.
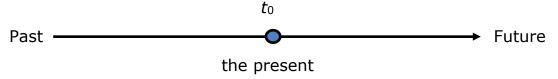
t_0

Past ——————————●——————————▶ Future

the present

Fig. 1. The time line

Prior (1957) proposed a logical language containing two *tense operators*[2]: one for the past, one for the future. Syntactically, these operators, PAST ("it has been the case that …") and FUT ("it will be the case that …"), are prefixed to a *proposition p*; this yields another proposition – a *past tense proposition* PAST(p), formalizing a sentence like (1a), or a *future tense proposition* FUT(p), formalizing a sentence like (1c). Prior's idea, which he revived from Medieval Scholastics, was that a proposition takes a truth value, True or False, *at* a time. The value of a proposition at a time t_1 may differ from its value at a different time t_2. For instance, proposition (1b) may be True at September 20th, 2020 but False at September 21st, 2020 (the reason being that the president

---

[2] Prior (1957) actually introduced other temporal operators. For our purposes it will suffice to consider the operators described in the main text.

may be in Brussels in the morning of the former but not of the latter). Prior's operators are defined below:

1.   Prior's tense operator PAST:

- PAST(p) is true at t *if and only if* p is true at a time t' in the past of t.
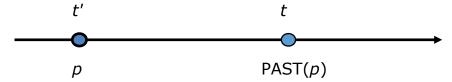

Fig. 2. Representation of the meaning of "PAST(p)"

2.   Prior's tense operator FUT:

- FUT(p) is true at t *if and only if* p is true at a time t' in the future of t.
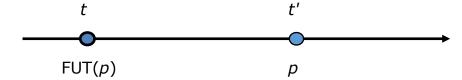

Fig. 3. Representation of the meaning of "FUT(p)"

To exemplify, suppose you uttered (1a) on the 20th of August, 2020:

(1)   a. (Ieri) il presidente dell'UCI era a Bruxelles

The Priorian logical structure of (1a) is (2):

(2)   PAST (the president of UCI *is* in Brussels)

and (2) is evaluated as in (3):

(3) is true at the time of August 20th, 2020 *if and only if* 'the president of UCI is in Brussels' is true at some time t' in the past of August 20th, 2020

So far, we have been implicitly assuming that the time at which a proposition is evaluated is the present time t_0. Though this is often the case, it is not *always* so. The proposition to which a tense operator is prefixed may be itself of the form PAST(p') or FUT(p'), so that the final proposition may have one of the following forms:

(4)   a.   PAST(PAST(p'))

        b.   PAST(FUT(p'))

        c.   FUT(PAST(p'))

        d.   FUT(FUT(p'))

Each of the formulas (4a)-(4d) may be associated to a particular tensed construction from natural language. Relevant tensed constructions are exemplified by the italicized parts of (4a')-(4d') (we read these sentences as making up a historical narrative, uttered at 2 PM on the same day as the reported events):

(4)   a'.   At 9AM, the president *had already arrived* in Brussels.

        b'.   He *would give* his speech at 10AM.

        c'.   By 6PM he *will already have returned* to Linate Airport.

        d'.   He *will be about to embark* for Rome then.

(4a') is the *past of a past*: the president's arrival in Brussels is presented as past not just with respect to the present but to a past 9AM time. (4b') is the *future of a past*: the president's speech is presented as future not with respect to the present but to that past 9AM time. (4c') is the *past of a future*: the president's return to Linate Airport is presented as past not with respect to the present but to a future 6PM time. Finally, (4d') is the *future of a future*: the president's embarkment for Rome is presented as future not just with respect to the present but to that future 6PM time.

The discussion of (4a')-(4d') should clarify that the times at which the innermost tense operators in (4a)-(4d) are evaluated are distinct from the time of utterance: they are displaced either in the past or the future of it by the outermost tense operators. While the broad temporal aspects of the meanings of (4a')-(4d') highlight a nice feature of the iteration of tense operators allowed in Prior's tense logic, there are finer temporal aspects of those meanings that lie beyond the reach of this logic, thus showing its limits. One such crucial aspect is the *reference to particular time points*, so common in natural language, which is typically achieved via time adverbials (see [4a']-[4d']).

2.2. Problems for the Priorian analysis of tense
Consider the following scenario (due to Barbara Partee). While driving on the highway to go to the sea, you utter (5) as you remind that, at the point in which you left home, you forgot to turn off the stove:

(5) I didn't turn off the stove!

The sentence clearly *refers to a particular time […] whose identity is generally clear from the extra-linguistic context,* just as the identity of the *he* in [the sentence *He shouldn't be in here*] is clear from the context (Partee 1973).

Partee notes that there are two possible Priorian analyses of (5), according to whether PAST takes scope above negation, i.e. (5'), or below it, i.e. (5"):

(5')   PAST(NOT(I turn off the stove))
(5")   NOT(PAST(I turn off the stove))

Consider the evaluation of (5') first. Assuming that you utter (5) at 5PM, the truth conditions of your utterance under analysis (5') are as follows:

(6) (5') is true at 5PM *if and only if* 'I turn off the stove' is false at some time t' in the past of 5PM

Notice that these truth conditions come down to requiring that *sometimes in the past you did not turn off the stove*. Turning now to (5"), the truth conditions of your utterance under analysis (5") are as follows:

(7) (5") is true at 5PM *if and only if* 'I turn off the stove' is true at no time t' in the past of 5PM

These truth conditions come down to requiring that *never in the past did you turn off the stove*.

The truth conditions predicted for your utterance under either one of (5'), (5") are both empirically inadequate. On the one hand, those predicted under (5') are *too weak*: there are lots of moments in the past in which you did not turn off the stove and might so satisfy the truth conditions in question; your utterance, however, is not about just *any* such moment, but about *a particular one* among them. On the

other hand, the truth conditions predicted under (5") are *too strong*: there are lots of moments in the past in which you turned off the stove; however, your utterance is not meant to deny that there is any such moment, but only to assert that you did not turn off the stove at a *particular moment* in the past.

The consideration of Partee's case clarifies, even without considering a sentence containing a time adverb, that Prior's tense logical approach is hopeless in dealing with the phenomenon of reference to specific times. We had anticipated the problem discussing the referential aspects of the narrative (4a')-(4d'): while Prior's tense operators allow one to move backward and forward in time (to evaluate propositions with respect to times displaced from the *hic et nunc* of the utterance), they lack the power of targeting specific moments in the way that time adverbials do. Therefore, if we want a satisfactory analysis of temporal location in language, while keeping to Prior's philosophical attitude, we need to refine Prior's logic so as to give his tense operators "referential power".

2.3. Hybrid Tense Logic

The following passage is from the entry *Hybrid Logic* of the *Stanford Encyclopedia of Philosophy*:

> Hybrid [tense] logics are logics that result by adding further expressive power to ordinary [tense] logic. The most basic hybrid logic is obtained by adding so-called *nominals* which are *propositional symbols of a new sort, each being true at exactly one [moment]* (Braüner 2017; emphasis added by FDP).

Let's illustrate what this means by applying it to Partee's case. Let's suppose that "*t**" (a *nominal*) is a very specific proposition, True at

exactly one moment, i.e. that particular moment in the past of your utterance in which you should have turned off the stove (in other words, "*t**" is a proposition that univocally describes the state of the world at the moment in question). Given this property of "*t**", it follows that the conjunction "*t** & NOT(I turn off the stove)" is a proposition True at *at most* one moment: *if* this conjunction is True at *any* moment, it cannot but be True at the unique moment at which "*t**" is True. In Partee's case, the meaning of your utterance of (5) can then be analysed along the lines of the logical paraphrase "it has been the case that the state of the world is as described by *t** and I do not turn off the stove," where "it has been the case that" is just Prior's past tense operator, "the state of the world is as described by *t**" is the very specific proposition that is True *only* at the crucial moment in the past referred to in your utterance, and "I do not turn off the stove" is the proposition (probably True at many moments in the past) that your utterance claims to be True at that same moment.

The hybrid tense logical analysis of (5) is thus as in (5‴), and the resulting truth conditions of your utterance are as in (8):

(5)   I didn't turn off the stove!

(5‴)   PAST(t^* & NOT(I turn off the stove))

(8)   (5‴) is true at 5PM *if and only if* 'I turn off the stove' is false at some time t^* in the past of 5PM coinciding with the moment the speaker is referring to (univocally described by the nominal "t^*").

These truth conditions can be depicted as follows (the green spot on the time line represents the presupposed antecedent event of you turning on the stove and is added to have a complete pictorial representation of the situation, the red spot represents the event of you not turning off the stove at the crucial moment identified by the nominal "t^*"):
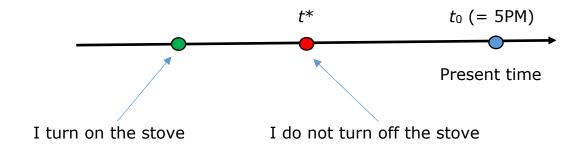
Fig. 4. "I didn't turn off the stove"

In conclusion, the introduction of *nominals* – propositions that are True at exactly one moment – endow Prior's tense logic with the referential power that this logic lacked.

2.4. Hybrid Tense Logic and Reichenbachian semantics of tense

An account of tense that is usually presented as an antagonist to Prior's is due to Hans Reichenbach, who proposed an apparently more complex analysis of natural language tenses. A strong empirical motivation for Reichenbach's analysis comes from *compound tenses*, e.g. the *Past Perfect* in (9), as opposed to the *Simple Past* in (10):

(9)   [CONTEXT SENTENCE: Mary arrived at 3PM.] Peter had left.
        (Inference: Peter's departure happens earlier than Mary's arrival)

(10)   [CONTEXT SENTENCE: Mary arrived at 3PM.] Peter left.
        (Inference: Mary's arrival happens as early as Peter's departure)

In his *Elements of Symbolic Logic*, Reichenbach writes:

"A particularly important form of token-reflexive symbol is found in the tenses of verbs. The tenses determine time with reference to the time point of the act of speech, i.e., of the token uttered. *A closer analysis reveals that the time indication given by the tenses is of a rather complex structure*. […] From a sentence like [(9)] we see that the time order expressed in the tense does not concern one event, but two events, whose positions are determined with respect to the point of speech. We shall call these time points the 'point of the event' and the 'point of reference'. In the example the point of the event is the time when Peter [left]; the point of reference is a time between this point and the point of speech. In an individual sentence like the one given it is not clear which time point is used as the point of reference. This determination is rather given by the context of speech" (Reichenbach 1947; I added the emphasis).

In this passage, Reichenbach draws his famous three-way distinction:

- Speech point (S, time point of the token)
- Reference point (R, the time of reference with respect to which the event is temporally located)
- Event point (E, the running time of the event)

This seems indeed to introduce some complexities with respect to Prior's analysis. In particular, the notion of *reference point* does not seem to correspond to any construct in Prior's tense logic. However, recall that Prior's tense operators can be iterated (e.g. [4a]-[4d]); if we add to this the referential power provided by nominals, the *prima facie* difference between the two approaches might vanish.

It has indeed been shown that Hybrid Tense Logic allows for a unification of Prior's and Reichenbach's insights about tense. A contribution in this direction comes from the work of Patrick Blackburn (Blackburn 1993, 1994). More recently, Blackburn and Jørgensen (2016) have argued "that Prior and Reichenbach are best viewed as allies, not antagonists" and they've done so "by combining the central insights of Prior and Reichenbach in the framework of hybrid tense logic". We illustrate how such a unification works via a couple of examples.

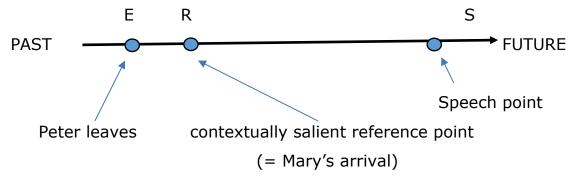For sentence (9), we have the following Reichenbachian analysis:

Fig. 5. Reichenbachian analysis of "[Mary arrived at 3PM.] Peter had left"

Using "r^*" as a nominal describing the reference point R, we obtain the Hybrid Tense Logic analysis (9') for (9):

(9')   PAST(r^* & PAST(Peter leaves))

That is: it has been the case in the past that r^* is True – and this takes us back to the time of Mary's arrival – and it has been the case in the past of this other time that Peter leaves. This gives us the correct temporal order by which Peter's departure precedes Mary's arrival.

For (10), we have the following Reichenbachian analysis (temporal coincidence between the event point and the reference point is depicted as spatial contiguity of the respective spots on the time line):
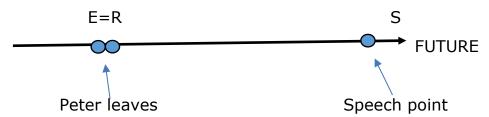
Fig. 6. Reichenbachian analysis of "[Mary arrived at 3PM.] Peter left"

Again, using "r^*" as a nominal describing the reference point R, we obtain the Hybrid Tense Logic analysis (10') for (10):

(10')   PAST(r^* & Peter leaves)

That is: it has been the case in the past that r^* is True – again, this takes us back to the time of Mary's arrival – and that Peter leaves. This correctly gives us the coincidence between Mary's arrival and Peter's departure.

2.5. Interaction between tense and temporal adverbs

Time adverbs are typically used to constrain the temporal location of events (Dowty 1979). Consider (1a) again, repeated as (11):

(11)   (Ieri) il presidente era a Bruxelles
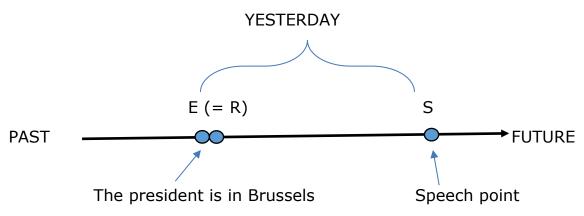
YESTERDAY

Fig. 7. Semantic contribution of "yesterday"

Semantically, adding the temporal adverb to the tensed verb in (11) comes down to constraining the temporal location of the event within the boundary of a particular time in the past: that past interval which corresponds to the day before the day of the utterance.

2.5.1. Operator-based analysis of temporal adverbs
David Dowty presents a problem for a *prima facie* plausible analysis of temporal adverbs in terms of Prior-style tense operators. He considers the following tentative analysis for the adverb "yesterday" (Dowty 1979):

- YESTERDAY(p) is true at t *iff* p is true at some time t' in the past of t which is within the day before t
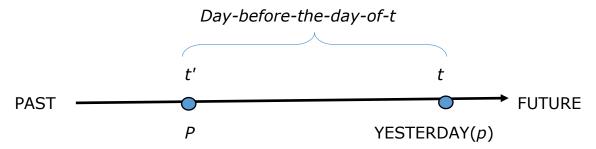
Day-before-the-day-of-t

Fig. 8. Operator-based analysis of "yesterday"

Let's assume now that (11) has the logical structure in (12) below. This formula then receives the truth conditions in (13):

(12) YESTERDAY(PAST(The president is in Brussells))

(13) "YESTERDAY(PAST(The president is in Brussells))" is true at *t iff* "PAST(The president is in Brussells)" is true at some time *t'* in the past of *t* which is within the day before *t iff* "The president is in Brussells" is true at some time *t''* such that, for some time *t'* in the past of *t* which is within the day before *t*, *t''* is in the past of *t*

These truth conditions can be depicted as in Figure 9:

YESTERDAY

t''   *t'*   *t*

The P. is in Br.   PAST(the P. is in Br.)   YESTERDAY(PAST(the P. is in Br.))
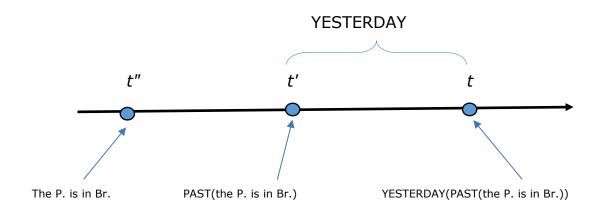
Fig. 9. "YESTERDAY(PAST(The president is in Brussels))"

Clearly, this analysis does not work: it predicts that (11) is true now if the president was in Brussels one year ago. In other words, the analysis in question predicts that (11) means the same as "At some point within yesterday it was true that the president had been in Brussels in the past of that point," which clearly is not the case. The same problem arises if we reverse the order of the operators PAST and YESTERDAY in (12), as shown in (14) (Dowty 1979: 323).

(14)   PAST(YESTERDAY(The president is in Brussells))

This alternative analysis does not work either since it predicts that (11) means the same as "At some point in the past it was true that the president had been in Brussels on the day before that point," which is also true now if the president was in Brussels one year ago.

In conclusion, we need to provide an account of the interaction between tense and temporal adverb which does not run into the empirical problem above – possibly keeping to Dowty's idea to treat "yesterday" as a Prior-like temporal operator.

2.5.2. Temporal adverbs in Hybrid Tense Logic
Hybrid tense logic helps us out of the problem of the interaction of tense and time adverbs. The crucial insight is twofold: (a) the past tense and the temporal adverb both take scope over the nominal $r*$ representing the reference point, and (b) they do so independently from one another. This scope and independence properties are shown by the hybrid tense logical formula in (15), which is evaluated as in (16):

(15)   PAST($r*$ & the president is in Brussels) & YESTERDAY($r*$)

(16) (15) is true at time t *iff*

(a)   'the president is in Brussels' is true at some t' in the past of t coinciding with the moment the speaker is referring to which is univocally described by $r*$;

(b) $r*$ is true at a time t' in the past of t which is within the day before t.

The truth conditions in (16) correctly capture the temporal fact (involved in the meaning of [1a]) that 'the president is in Brussels' is true at some time *t'* in the past of *t* which is within the day before *t*.

To summarize, the present section has been essentially devoted to presenting an outstanding formal approach to the semantics of tense, keeping an eye to its historical development. For the sake of this exposition, I have assumed the representation of time as a directed line, largely used by linguists, without problematizing it. In the next section I turn to some substantive questions regarding time as is represented in language.

3. Directionality of time in language and branching futures

In this section I present linguistic data supporting the view of (a) the directionality of time and (b) the asymmetry between the uniqueness of the past and the plurality of the (possible) future(s).

3.1. Temporal asymmetries: polarised scalar adverbs

Having regard to the properties of time, some telling asymmetries appear as soon as we consider data such as the contrasts below (two interrogative marks in front of a sentence indicate that this is semantically anomalous):

(17) a. È già tardi.
        'It is already late'
      b. ??È ancora tardi.
        'It is still late'

(18) a. È ancora presto.
        'It is still early'
      b. ??È già presto.
        'It is already early'

The so-called "phase adverbs" *già* 'already' and *ancora* 'still' display a puzzling behaviour in their interaction with the gradable adverbs *presto* 'early' and *tardi* 'late': each one of the former can meaningfully modify only one of the latter. Why so?

Let's characterize the contrasts above by drawing some pictures. First, we associate the adverbs *presto* and *tardi* with two regions of the time scale, separated by a transition point, i.e., a moment in time at which it is not early anymore and it is not late yet; the region associated with *presto* is located to the past side, while the region associated with *tardi* to the future side – obviously so, since what happens earlier in time becomes past at a point at which what happens later is future.
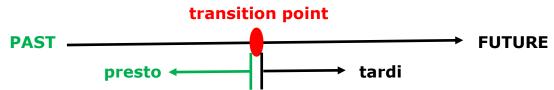

Fig. 10. Transition between the two regions presto – tardi on the time scale

The rationale behind the "pastward" orientation of the *presto*-arrow and the "futureward" orientation of the *tardi*-arrow is simple: *presto* has degrees that grow as you move from later times to earlier times, while *tardi* the other way round.

Next, we associate the adverbs *ancora* and *già* with two regions of *some* scale (not necessarily the time scale, for reasons clarified below), separated by a transition point, i.e., a point at which a certain property, P_1, does not hold anymore while the opposite property, P_2, does not hold yet; the region associated with *ancora* is located to the "low values" side of the scale while the region associated with *già* is located to the "high values" side.
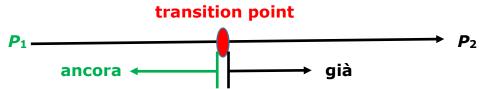
Fig. 11. Transition between the two regions ancora – già on a relevant scale (associated with gradable properties P_1, P_2)

It may not be obvious why *ancora* and *già* should be associated with regions of the scale in this way. Here is why:

- *ancora* indicates that some gradable property, P_1, continues to hold as one proceeds to values on the scale that are ever closer to the Point of Transition, *PT*; starting from *PT*, P_1 ceases to hold and a gradable property P_2, opposite to P_1, begins to hold;
- *già* indicates that P_2 has begun to hold and will moreover hold at all higher values on the scale.

Hereafter, I refer to *ancora* and *già* as "polarised scalar adverbs": their meanings are defined on some relevant *scale* (which may be the time scale or other) and they are "polarised" since they are associated to opposite *poles* on that scale.

3.1.1. Polarised scalar adverbs in non-temporal domains
Consider the following context:

Context NT1. We are looking for *tall* persons for a certain task (say, to pick fruits from apple trees). The minimal height required to do the task is 1m80cm (the *transition point*). Two candidates come to us, Leo and Teo, whose heights are 1m85 and 1m75, respectively.

We can represent the situation of Context NT1 via a scale of height measures, ordered from smaller measures (in a region of the scale that we associate with the property 'short') to greater measures (in a region of the scale that we associate with the property 'tall'):
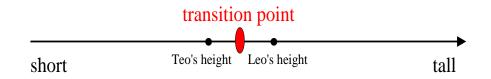

Fig. 12. The situation of Context NT1

Now, consider (19), as uttered in Context NT1:

(19)   Leo è 1m85, è *già* (abbastanza) alto. Teo è 1m75, è *ancora* (troppo) basso.
'Leo is 1m85, he's already (enough) tall. Teo is 1m75, he's still (too) short'.

This sentence would be true in the envisaged situation. However, sentence (20), as uttered in Context NT1, would be meaningless:

(20)   Leo è 1m85, è ancora alto. Teo è 1m75, è già basso.
'Leo is 1m85, he is still tall. Teo is 1m75, he is already short'.

Let's now consider the following alternative to Context NT1:

Context NT2. We are looking for *short* persons for another task (say, to work in a room with a low ceiling). The maximal height allowed for this task is 1m80cm. Two candidates come to us, Leo and Teo, whose heights are 1m85 and 1m75, respectively.

Again, we can represent the situation of Context NT2 via a scale of height measures, but notice that this time the height measures will be ordered from greater ones to smaller ones:
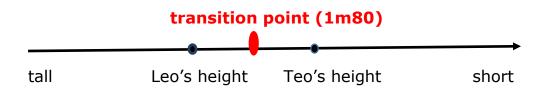

Fig. 13. The situation of Context NT2

If we now consider (20) again, but this time as uttered in Context NT2, we can see that, far from being meaningless, (20) is now true. On the other hand, (19), which was true as uttered in Context NT1, as uttered in Context NT2 is meaningless!

We interpret the observations above by noting the following properties of *già* and *ancora* (we are focusing here on their use in *non-temporal domains*):

- "Già" indicates *upward persistence of a property*.
  [In Context NT1: Leo's height is enough and the same is true of any height *following* Leo's on the relevant scale (Fig. 12); in Context NT2: Teo's height is small enough and the same is true of any height *following* Teo's on the relevant scale (Fig. 13)].

- "Ancora" indicates *downward persistence of a property*.
  [In Context NT1: Teo's height is not enough and the same is true of any height *preceding* Teo's on the relevant scale (Fig. 12); in Context NT2: Leo's height is not small enough and the same is true of any height *preceding* Leo's on the relevant scale (Fig. 13)].

(Compare these properties with what was captured in Figure 11).

3.1.2. Polarised scalar adverbs in the temporal domain
Take the following context:

> *Context T1.* We are waiting for our friends, Leo and Teo, at a birthday party. They are expected to arrive by 9PM. Leo arrives by the expected time, but at 9:10PM Teo still has to arrive.

We can represent this situation via the usual time scale (where times are ordered from earlier to later):
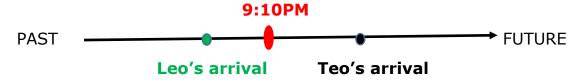

Fig. 14. The situation of Context T1

Suppose that at 9:10PM we utter (21):

> (21) Leo è *già* arrivato. Teo *ancora* non è arrivato.
>   'Leo has already arrived. Teo still has not arrived'

This sentence is true in the envisaged situation. On the other hand, sentence (22), as uttered in Context T1, would be meaningless:

> (22) ??Leo è ancora arrivato. Teo è già non arrivato.
>   'Leo has still arrived. Teo has already not arrived'

Let's focus on the left sentence of (21) (*Leo è già arrivato* 'Leo has already arrived'). Figure 15 depicts that part of its meaning –

contributed by *già* – that I characterize in terms of *forward persistence of a state (in time)*:

9:10PM

Past ⟶ Future

Leo's arrival   *e*   *e*   *e*

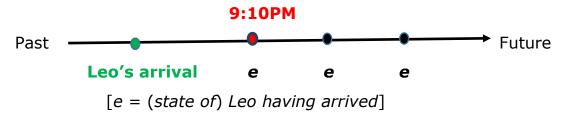[*e* = (*state of*) *Leo having arrived*]

Fig. 15. Forward persistence of Leo having arrived

The idea is simple: Leo arrived (at some point before 9PM), and from that point onward it will be forever true that Leo arrived.

Turning to the right sentence of (21) (*Teo ancora non è arrivato* 'Teo still has not arrived'), Figure 16 depicts that part of its meaning – contributed by *ancora* – that I characterize in terms of *backward persistence of a state (in time)*:

9:10PM

Past ⟶ Future

¬*e*   ¬*e*   ¬*e*   **Teo's arrival**

[¬*e* = (*state of*) *Teo not having arrived*]
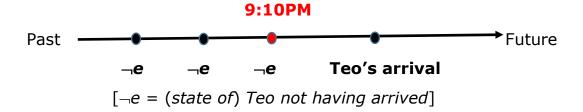
Fig. 16. Backward persistence of Teo not having arrived (yet)

The idea, again, is simple: Teo has not arrived (yet) by some point after 9PM, and moving to the past of that point it is always true that Teo has not arrived (yet).

We interpret the observations above by noting the following properties of *già* and *ancora* (this time we are focusing on their use in *temporal domains*):

- "Già" indicates *forward persistence of a state in time*; with a formula:

  $\exists t \in I \, [\varphi(t) \, \& \, \forall t' \in I \, [t < t' \rightarrow \varphi(t')]]$

  [Leo's having arrived holds at 9:10PM and it also holds at any time *following* 9:10PM on the time scale]

- "Ancora" indicates *backward persistence of a state in time*; with a formula:

  $\exists t \in I \, [\varphi(t) \, \& \, \forall t' \in I \, [t' < t \rightarrow \varphi(t')]]$

  [Teo's not having arrived holds at 9:10PM and it also holds at any time *preceding* 9:10PM on the time scale]

Let's now turn to the following alternative to Context T1 – this is in all respects like Context T1, except that we are now trying to *reverse the time scale* (via the "thought experiment" added in italics to the context description)[3]:

> *Context T2.* We are waiting for our friends, Leo and Teo, at a birth-day party. They are expected to arrive by 9PM. Leo arrives by the expected time, but at 9:10PM Teo still has to arrive. *We feel angry at Teo – he is always late! – and to defocus from our unpleasant feeling we imagine Teo's arrival as being already there, to the future, as much real as Leo's past arrival, and we imagine that we are moving through time from Teo's arrival back to Leo's arrival.*

This time the situation is represented as follows (notice the peculiar orientation of the time line from the future to the past, i.e. the ordering of times *from later to earlier*):

---

[3] We want to see if we can rescue (22) by inverting the "natural" (*earlier-to-later*) order of times. Recall that, in the passage from Context NT1 to Context NT2, we were able to rescue (20) by inverting the order of the degrees of height.

9:10PM

Future —————————•————————⬤————————•————————▶Past

Teo's arrival   **Leo's arrival**

Fig. 17. The situation of Context T2

Crucially, (22) continues to be meaningless, even if considered as uttered in Context T2 – hence, against the reversed time scale of Fig. 17. On the one hand, the left sentence of (22), if it could ever mean anything, would mean the same as 'at 9:10PM, Leo is still (in the state of having) arrived;' this doesn't seem to be a meaningful thought to be ever expressed! On the other, the right sentence of (22), if it could ever mean anything, would mean the same as 'at 9:10PM, Teo is already (in the state of having) not arrived;' this, too, doesn't seem to possibly make any sense!

The moral of this section concerns a fundamental difference between time *versus* other types of scales. The different behaviour of the adverbs *ancora / già* in the non-temporal contexts NT1 - NT2 and in the temporal contexts T1 - T2 suggests that the natural ordering characterizing time cannot be reversed in the same way as the *smaller-to-greater* relation characterizing the scale of degrees of height (we chose the pair of predicates *alto - basso* but we could have made our point by picking a different pair of non-temporal antonyms).

3.2. Asymmetry between past and future

Section 3.1. has shown some evidence that time is represented as *directional* in language: unlike other kinds of scales that natural language meanings rely on, time is characterized by a "natural" ordering that is not reversible. We consider now whether there are other properties of time, revealed by linguistic data, that may have a relation to this directionality property. It turns out that there is another one, on which

an abundant literature on the asymmetry between the linear past and the branching future has flourished.

3.2.1. Many futures, one past

There are data suggesting that the traditional linear representation of time is too simplistic – this representation may be useful for the purpose of analysing certain temporal aspects of natural language meaning, but it is limited when one comes to the analysis of more subtle data at the crossroads of temporality and modality.

Consider the contrast between (23a) and (23b):

(23)   a. The meeting will take place in Room C, unless the management makes the Aula Magna available;

b. The meeting took place in Room C, unless the management made the Aula Magna available.

On the one hand, (23a) expresses a condition for the future occurrence of the meeting in Room C, namely that the management will not choose the Aula Magna instead. The "unless" clause in this sentence introduces this condition without making it so that the main clause in the future tense ("The meeting will take place in Room C") is not asserted – the latter *is* asserted in spite of the "unless" clause, which only has the effect of qualifying the assertion (Condoravdi 2003). In (23b), on the other hand, the "unless" clause has a different effect: it makes it so that the main clause in the past tense ("The meeting took place in Room C") can no longer be seen as asserted, once the "unless" is processed.

Another relevant contrast is that between (24a) and (24b):

(24)   a. The management made the Aula Magna available *before* the meeting took place in Room C;

> b. The management made the Aula Magna available *after* the meeting took place in Room C.

On the one hand, in (24a) the meeting taking place in Room C is still *future* with respect to the management making the Aula Magna available and the very fact that the management made the Aula Magna available may well have had the effect that the meeting didn't take place in Room C after all (because it probably took place in Aula Magna instead). On the other hand, in (24b) the meeting taking place in Room C is already *past* with respect to the management making the Aula Magna available, and the fact that the management made the Aula Magna available cannot possibly have had the effect that the meeting didn't take place in Room C – what is past cannot be canceled.

Let's consider one final contrast, the one between the Question-Answer pairs in (25a) and (25b):

> (25) a.   Q: Where will the meeting take pace?
> A: The ministers have two places available. It *may take place* in Room C or in Aula Magna';
> b.   Q: Where did the meeting take pace?
> A: The ministers had two places available. It *may have taken place* in Room C or in Aula Magna.

Only in (25b) does the modal mean that the speaker *does not know* which room, while the issue is objectively settled. In (25a) the modal does not (at least, not necessarily) mean that the speaker does not know which room and it means instead (at least, under a very plausible reading of the sentence) that the issue is objectively open.

The general idea brought to salience by such linguistic contrasts is an old and venerable idea: the past is fixed and determinate, the

future is open and indeterminate (Aristotle 1984, Prior 1967, Thomason 1984, Belnap et al. 2001, Condoravdi 2003).

3.2.2. "Branching time" approaches to the past-future asymmetry

At this point, the Priorian attitude toward temporal language shows up again. A little bit more formally, the idea is that our world, considered from the perspective of any moment m_0 in its history, has a unique actual past and present – the actual past and the actual present at m_0 – but many metaphysically (not just *epistemically*) possible futures – the possible futures at m_0. A very natural way to represent this view diagrammatically, and to translate it next into a formal model, is the following:

a) represent the unicity of the past at m_0 as a single oriented line (usual left-to-right orientation) that ends at m_0 (see the green line in Figure 18);

b) represent the plurality of the future at m_0 as a manifold of oriented lines branching off from m_0 – all such lines begin at m_0 and proceed rightward, possibly breaking into different lines at subsequent moments (see the red lines in Figure 18);

c) represent the world itself, as considered from the perspective of moment m_0, as the cluster of all histories passing through m_0 – where by "history" it is meant a complete linear path made up of moments temporally connected with each other (Figure 18 shows four such histories, h_1-h_4).
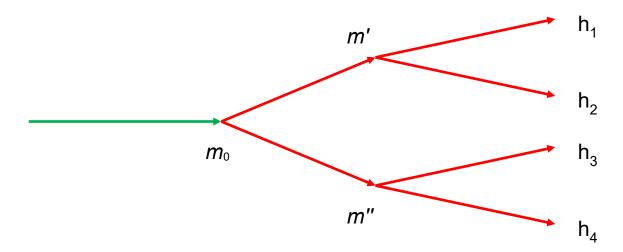
Fig. 18. "Branching Time" representation of our world from the perspective of moment m_0

The resulting model has been known as "Branching Time" (Prior 1967, Thomason 1984, Belnap et al. 2001).

Let me illustrate the basic concepts of Branching Time by discussing the structure of Figure 18. This has four histories: h_1, h_2, h_3 and h_4. Moment m_0 belongs to all four. It can thus be said that there are four possible futures at m_0: the four paths that our world can take after m_0, following one or the other of h_1-h_4. Moment m' only belongs to h_1 and h_2. It can thus be said that m' only has two possible futures, corresponding to h_1 and h_2, while h_3 and h_4, which were possible futures at m_0, are no longer possible futures at m'. Similarly, moment m'' only belongs to h_3 and h_4 and can then be said to only have two possible futures (in general, futures shrink as we move forward in time; we come back to this shortly). Regarding the ordering of moments via the *earlier-to-later* relation, not all moments are temporally connected with each other: while m_0 is in the past of both m' and m'' (and is so connected with both), m' is not in the past of m'' nor is m'' in the past of m' (thus, m' and m'' are not connected with each other). The fact is

that m' and m'' are moments lying on *alternative futures*; each of these futures may end up becoming actual, but, crucially, they could not become *both* actual (though many alternative courses of events may be possible at a given point, history always ends up realizing only one of them).

Notice that the Branching Time model lends itself naturally to a "growing block" view of the past: the past of our world "grows" as we move forward in time, for instance from m_0 to m' (the past at m' is a "larger block" than the past at m_0). This is a sense – the *only* sense – in which the past can be said to "change". In all other respects, it is immutable. In particular, what is past at a moment m_i continues to be past at any later moment m_j; in this sense, there may be *no loss of past* as time goes by. One thus finds a sharp contrast with the future: what is a possible future at a moment m_i may no longer be a possible future at a later moment m_j; in this sense, there may be loss of possible futures as time goes by.

The Branching Time model allows us to account for the asymmetries observed in the interpretation of (23a,b), (24a,b) and (25a,b) above. To prove this, I should provide formal analyses in Branching Time of the relevant linguistic constructions featuring in those sentences. For reasons of space, I'll confine myself to discussing the *before* / *after* contrast in (24a,b), against the temporal structure of Figure 18. Suppose that (i) the management makes the Aula Magna available at m_0, (ii) the meeting takes place in Room C at m' but not at m'' (notice that supposition (ii) is coherent since m' and m'' are on *alternative* futures for m_0, and what happens in one future may not happen in the other). Suppose further that, (iii) right after m_0, the world takes the path of history h_4, thus going through the moment m'' in which the meeting does not take place in Room C. Given (i)-(iii), an utterance of (24a) made right after m_0 would be True: the Aula Magna is made available at a moment, m_0, at which the meeting taking place

in Room C has not occurred yet and the latter event occurs at a moment, m', in a possible future of m_0; however, the event in question does not occur in a different future of m_0, i.e. h_4, which ends up becoming the actual future.

Suppose now that (a) the meeting takes place in Room C at m_0, (b) the management makes the Aula Magna available at m', and (c) right after m_0, the world takes the path of history h_1, thus going through the moment m' in which the management makes the Aula Magna available. Given (a)-(c), an utterance of (24b) made right after m' would be True: the Aula Magna is made available at a moment, m', at which the meeting taking place in Room C has already occurred, since the latter event occurs at a moment, m_0, in the past of m'; since m' has only one past, it is not possible that the meeting taking place in Room C does not occur in a different past of m' (once this event has occurred, there is nothing you can do which could make it not have occurred!).

The two reasonings above are just a sketch of an account of how *A before B* can be true in a history without *B* ever being true in that history and how *A after B* can be true in a history only when *B* is also true in that history[4].

3.2.3. Formal analyses of the past and the future tenses in Branching Time

As we consider the future tense, the Branching Time model offers us different options to formally analyse it. The availability of different options is expected, given the complexity of the future dimension that is captured by this model.

One important option is the so-called "Peircean future" (Prior 1967):

---

[4] These properties are known as the "non-veridicality" of *before* (the truth of *A before B* does not entail that *B* is true) and the "veridicality" of *after* (the truth of *A after B* entails the truth of *B*). See Beaver and Condoravdi 2003.

- "FUT(p)" is true at m_0 *if and only if* "*p*" is true at a moment following m_0 *in every history passing through m_0*
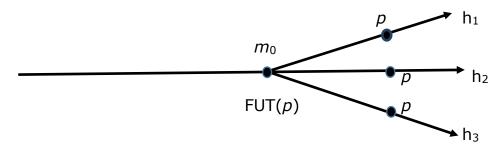

Fig. 19. "Peircean" future

On the Peircean solution, future tensed sentences are essentially *necessity modal sentences*. This solution thus predicts that the future contingent statement (FC) is equivalent to the necessity modal statement (NF):

(FC)   There will be a sea battle tomorrow.

(NF)   It is now necessary that a sea battle will occur tomorrow.

This can be criticized as an undesirable outcome of the solution in question.

Another notable solution is known as the "Ockhamist future" (Prior 1967):

- "FUT(p)" is true at m_0 *if and only if* "p" is true at a moment following m_0 *in the history h* that will be actual* (Prior 1967)
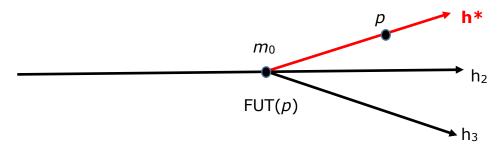

Fig. 20. "Ockhamist" future

The Ockhamist solution is based on the idea that, although there are many possible futures at m_0, only one of these will be actual (the one corresponding to h* in Figure 20) and a future tensed statement made at m_0 targets *this only future*. This solution has the following consequence: if I assert at m_0 "there will be a sea battle tomorrow" and on the day following the day of m_0 (in the history that becomes actual) there happens (not) to be a sea battle, then my assertion is (not) True, *regardless of what happens in the other histories that were possible futures at m_0*. This consequence has some intuitive appeal – we *do* look at what actually happened when we have to assess past utterances of future tensed sentences for truth (MacFarlane 2003). Besides this, the Ockhamist solution (unlike the Peircean) preserves the contingent character of (FC) above, as distinct from the necessity modal statement (NF). I'll leave it open what is the empirically most adequate analysis of the future tense in Branching Time.

Let's turn to the past tense. It should be obvious that Branching Time does not provide us with several options in this case – for a "backward looking" tense operator (like PAST), it does not make any difference whether the temporal structure is the traditional linear structure, or whether it is the tree-shaped structure of Branching Time: as long as we look backward in time, we always face a linear path. The

past tense operator can thus be defined in the following straightforward way (making no reference to any particular history):

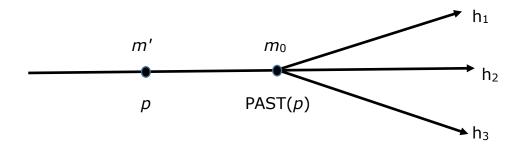- "PAST(p)" is true at m_0 *if and only if* p is true at a moment m' preceding m_0

Fig. 21. The past in Branching Time

Since there is a unique history (or, more precisely, a unique path) going backward from moment m_0, the past tense does not raise a problem of definition in the same way as the future tense does. We only need to look at that unique past history (path) and check whether it contains a moment at which p holds or not.

Given Branching Time, there is a sense in which PAST(p), unlike FUT(p), when true, is not just *contingently* true but true *by necessity*:

- If p is now past, then it is now *necessary* that p is past (thesis of the "necessity of the past").

The relevant notion of necessity is what has been called *historical necessity* (Thomason 1984); taking the symbol "NEC" to denote historical necessity, this notion can be defined as in (HN) and the above thesis of the "necessity of the past" can be formulated as in (NP):

(HN)   "NEC(p)" is true at m *if and only if* "p" is true at m *in every*

history passing through m

(NP)   PAST(p) $\rightarrow$ NEC(PAST(p))

In conclusion, Branching Time provides for a sense in which the past is necessary. This is the way in which this model formally captures the standard view of the non-modifiability (or persistence) of the past. In the next section I introduce a puzzle that challenges this view.

4. The Puzzle of the "Changing Past"

Barlassina and Del Prete (2015) (hereafter, B&DP) present the following puzzle about the past. Preliminarily, some technical vocabulary:

- People utter sentences at *contexts*, represented as pairs <w, t> of a world and a time (the world and the time of the utterance).
- Sentences express *propositions* at contexts.
- A sentence S is *temporally specific* if and only if, for any context c, the proposition expressed by S at c is about a specific time.
- A sentence S is *about the past* in a context c if and only if the proposition expressed by S at c is about a time that precedes c.
- A temporally specific sentence S that is about the past in a context c is *true in c* if and only if *the time* the proposition p expressed by S in c is about *has*, relative to the past of c, *the property that p ascribes to it*.
- A context c' is a *successive same-world context* to context c if and only if:
  the world of c' is the same as the world of c,
  the time of c' follows the time of c.
- A sentence is *context-insensitive* if and only if it expresses the same proposition at all contexts.

Their starting point is what they take to be a platitude: *we say true (false) things about the past and the truth (falsity) of what we say depends on how the past is*. This is made formally precise as follows:

- *Truth About the Past* (TAP)

  Let S, c, p, Q, t be such that: S is a temporally specific sentence that is about a past time in context c, p is the proposition expressed by S at c, and Q is the property that p ascribes to the specific time t. Then, S is true in c if and only if t has property Q relative to the past of c.

To see how TAP works, consider sentence

(26)   Barack Obama was born in 1961

as uttered in context <@, 23rd May 2019> – this consists of our actual world @ and time 23rd May 2019; (26) is temporally specific (it is about the year 1961), it is about the past (the year 1961 is in the past of 23rd May 2019), and the proposition it expresses at <@, 23rd May 2019> (i.e., *that Obama was born in 1961*) ascribes to the year 1961 the property of being a time in which Obama was born. By TAP, (26) is true in <@, 23rd May 2019> since, relative to the past of <@, 23rd May 2019>, the year 1961 has indeed the property of being a time in which Obama was born.

B&DP ask us next to consider the following story (almost true – apart from the fictional character of Frank):

"It is the 23rd of July 2000. Being the rider with the lowest overall time at the end of the last stage, Lance Armstrong is declared the winner of the Tour de France by *Union du*

Cyclisme Internationale (UCI). On <@, 25th December 2002> (hereafter, Context A), Frank utters [(27)]:

[(27)] Lance Armstrong won the Tour de France in 2000.

There is a clear intuition that Frank said something true. Time goes by. Having discovered that Armstrong made use of banned substances, on the 22nd of October 2012 UCI withdraws all of Armstrong's wins at TdF. As Frank is not aware of this fact, he utters [(27)] again at <@, 25th December 2012> (hereafter, Context B). This time, it seems that Frank said something false".

Based on this story, they present the following reasoning, leading to the paradoxical conclusion in F:

(A) (27) is a *temporally specific* sentence: for any context *c*, the proposition expressed by (27) at *c* is about the year 2000;

(B) (27) is *about the past* in both Context A and Context B;

(C) (27) is *context-insensitive* (it does not contain any context-sensitive element – demonstratives, indexicals, or the like), hence it expresses the same proposition at both Context A and Context B, namely, *that Armstrong won the TdF in 2000*;

(D) (27) is *true in Context A*, hence it follows from *TAP* that, relative to the past of Context A, the year 2000 has the property of being a time in which Armstrong won the TdF;

(E) (27) is *false in Context B*, hence it follows from *TAP* that, relative to the past of Context B, the year 2000 does not have the property of being a time in which Armstrong won the TdF;

(F) Context B is a *successive same-world context* to Context A (Context A and Context B are located in the same world, @),

hence in moving from Context A to Context B *the past* (of the actual world) *has changed*: the year 2000 had a certain property on Christmas 2002, but did not have that property on Christmas 2012 any longer.

B&DP consider two objections to the puzzling conclusion in (F). The first one contends that there is no truth value change for (27), hence questioning the joint truth of (D) and (E). One variant of this objection consists in claiming that (27) was *already false in Context A* because Armstrong got the lowest time only by cheating. The other variant consists in claiming that (27) was *still true in Context B* because sincere and informed speakers assert (27), or sentences implying (27), after the revocation of Armstrong's titles (for instance, they assert true historical sentences such as "Armstrong won the TdF seven times from 1999 to 2005. He was later stripped of those titles for doping" …). The second objection contends that (27) has some hidden context-sensitivity (due to the verb phrase 'win the TdF in 2000'), against assumption (C): in Context A, (27) would express the proposition *that Armstrong won the TdF in 2000 according to the declaration of Context A*, but in Context B, it would express the different proposition *that Armstrong won the TdF in 2000 according to the declaration of Context B*. B&DP reply to each objection with detailed linguistic arguments that I cannot recall here. Notice that, if they are right, there might be other cases in which our past can change and this would force us to try to make sense of this.

5. Conclusion

We reviewed some classical approaches, from logic and philosophy, to temporal meaning in natural language, beginning with the subject of temporal location of events in language and moving then to general properties of time as may be revealed by semantic analysis.

Throughout the article, the analytical attitude toward temporal language has been inspired by the sort of logical approach to time and tense that Prior developed in the 1950s and '60s. Building on previous work on historical modalities, we introduced the asymmetry between closed past and open future, showing some linguistic correlates of it, and presented the Branching Time model as a formal tool to capture this asymmetry and to build enlightening semantic analyses of a number of linguistic contrasts. Finally, we considered a paradoxical argument whose conclusion challenges the traditional view of the persistence of the past. If its proponents are right, we should stop asking *whether* the past can change and start to think *how* this could be. One might notice that the argument is based on a special property of winning events, namely their annullability, and on this basis might contend that the puzzle is not generalizable to other sorts of events that are not annullable in the same way – for instance: though it was possible for UCI in 2012 to make it the case that Armstrong did not win the TdF in 2000, how could it be possible for *anyone* after the 2000 TdF to make it the case that Armstrong did not have the lowest overall time at the end of the last stage or that he did not ascend the Hautacam on the 10[th] of July? It seems hard to erase events of the latter sort from history. These speculative remarks raise the question of a distinction between two sorts of events, with events such as winning the TdF being of one sort and events such as ascending the Hautacam on a bicycle being of the other sort. Providing an explicit characterization of the two sorts of events – one which would allow for a neat separation, among the events reported in our historical narratives, between events of one sort and events of the other sort – might prove to be a difficult task.

References

Aristotle. (1984). *Organon*. In J. Barnes (ed.), *The Complete works of*

Aristotle. Princeton: Princeton University Press.

Barlassina, L. and Del Prete, F. (2015). The puzzle of the changing past. *Analysis* 75: 59–67.

Beaver, D. and Condoravdi, C. (2003). A uniform analysis of *before* and *after*. In *Proceedings from SALT 13*, 37–54. Ithaca, NY: CLC Publications.

Belnap, N., Perloff, M., and Xu, M. (2001). *Facing the Future. Agents and Choices in our Indeterminist World*. Oxford: Oxford University Press.

Blackburn, P. (1993). Nominal Tense Logic. *Notre Dame Journal of Formal Logic* 14: 56–83.

Blackburn, P. (1994). Tense, temporal reference and tense logic. *Journal of Semantics* 11: 83–101.

Blackburn, P. and Jørgensen, K. F., (2016). Reichenbach, Prior and hybrid tense logic. *Synthese* 193: 3677–3689.

Braüner, T. (2017). Hybrid Logic. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), URL = <https://plato.stanford.edu/archives/sum2017/entries/logic-hybrid/>.

Comrie, B. (1985). *Tense*. Cambridge: Cambridge University Press.

Condoravdi, C. (2003). Moods and modalities for *will* and *would*. In *Proceedings of the Amsterdam Colloquium 2003*.

Copeland, B. J. (2020). Arthur Prior. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2020 Edition), URL = https://plato.stanford.edu/archives/spr2020/entries/prior/

Dowty, D. R. (1979). *Word Meaning and Montague Grammar*. Dordrecht: Reidel Publishing Company.

Goranko, V. and Rumberg, A. (2020). Temporal Logic. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2020 Edition), URL = <https://plato.stanford.edu/archives/sum2020/entries/logic-

temporal/>.

MacFarlane, J. (2003). Future Contingents and Relative Truth. *The Philosophical Quarterly* 53: 321-336.

Øhrstrøm, P. and Hasle, P. (1995). *Temporal Logic: From Ancient Ideas to Artificial Intelligence*. Dordrecht: Kluwer.

Partee, B. H. (1973). Some structural analogies between tenses and pronouns in English. *Journal of Philosophy* 70: 601–609.

Prior, A. N. (1957). *Time and Modality*. Oxford: Oxford University Press.

Prior, A. N. (1967). *Past, Present and Future*. Oxford: Oxford University Press.

Quine, W. V. O. (1960). *Word and Object*. Cambridge, Massachusetts: MIT Press.

Reichenbach, H. (1947). *Elements of Symbolic Logic*. New York: The Free Press.

Thomason, R. (1984). Combination of tense and modality. In D. Gabbay and F. Guenthner (eds.), *Handbook of philosophical logic: Extensions of classical logic*, 135–165. Dordrecht: Reidel.