

TEMPORAL REFERENCE INSIDE AND OUTSIDE PROPOSITIONAL ATTITUDES

Hans Kamp, University of Stuttgart, Germany

1 Summary

In this paper I discuss one application of a theory whose general features have been presented elsewhere.¹ The core of this theory is a formalism, with well-defined syntax and (model-theoretic) semantics, for the representation of propositional attitudes, complex mental states consisting of several connected attitudes, attitudinal change (i.e. the change from one attitudinal state to another), and attitudes shared by several agents. A second, though closely related purpose of the formalism is that it can serve as part of a semantics of attitude attributing sentences of natural language. In designing the formalism special attention was given to the internal connections between the contents of different attitudes. In the context of attitude attribution these connections manifest themselves as problems for the semantics of discourses consisting of several attitude attribution sentences, which either attribute different attitudes to the same agent at the same time, or different attitudes to the same agent at different times, or attitudes to different agents at the same or different times.

The one example to which the present paper is devoted is a “mini-discourse” consisting of two sentences which ascribe attitudinal changes to the same agent at two distinct times. The example illustrates two recurrent features of multi-sentence attitude attributing discourses. On the one hand it exemplifies the by now familiar fact that in a sequence of two or more attitude attributions the later attributions tend to rely for their interpretation on the attributions which precede them. This is a phenomenon that also arises when the attributions concern one and the same time (as well as one and the same agent) and it is for such cases that it has been identified and discussed in the literature.² But in addition the example illustrates a number of issues that have to do with temporal reference and time. Some of these arise at the level of the single sentence. This is true in particular of the question how the tenses of the complement clauses of attitude attributing matrix verbs are semantically related to the tenses of the matrix verbs themselves. But there are also time-related questions that concern the way in which the two attitude attributing sentences are connected, as parts of a single cohesive piece of discourse. On the one hand these have to do with the circumstance that the sentences of our example attribute attitudinal *changes* and on the other with the fact that these changes are said to have occurred at

¹ See Kamp (1990, 2003); van Genabith et al. (in press).

² See Stalnaker (1988).

different times. These temporal issues add a new dimension to the general problem how earlier attributions can provide interpretation contexts for later ones.

The theory which I will apply to the example of the paper is an extension of Discourse Representation Theory as it is presented in for instance Kamp & Reyle (1993). More accurately it is a combination of two extensions. One of these concerns the representation of attitudinal states and the semantics of that part of language which is used to describe such states; the other is presupposition. This second extension too has been presented elsewhere.³ Abridged presentations of both extensions can be found in van Genabith et al. (in press).

I will assume some basic knowledge of DRT, roughly corresponding to Chapters 1, 2 and 5 of Kamp & Reyle (1993). But in Sections 3.1 and 3.2 below I give brief introductions to the two mentioned extensions, which should give the reader enough to understand the treatment that follows in Section 4 of the example that this paper is about. The remaining sections of the paper are all quite short. Section 2 presents the example and lists the problems connected with it on which we will focus. Section 5 sums up and restates the principal morals.

2 The Example and the Issues it raises

The example which that is the topic and focus of this paper is given in (1):⁴

- (1) On Sunday Bill heard that Mary was in Paris.
On Tuesday he learned that on the previous day she had left.

The issues raised by this example that will preoccupy us in what follows all have to do with time. Temporal matters are relevant to the interpretation of (1) for three distinct reasons.

- (i) Both attributions are attributions of attitudinal change – i.e. of a change from one attitudinal state to another – expressed by the verbs *hear* and *learn*. As opposed to verbs like *believe*, *want*, *intend*, *regret* and others which are used to attribute a single attitude to their subject at or over a single period of time, such attitudinal change verbs have not received much attention either in the philosophical or the linguistic literature on attitude attribution.
- (ii) Like any other attitude attributing sentences with a matrix verb and a *that*-complement, the sentences in (1) raise the problem how the time of the event or state of the complement clause of a matrix verb relates to the time of the matrix event or state. It is a problem which strictly speaking arises for any combination of tenses in main clause and complement clause. But when both clauses are in the present tense, as has typically been the case for the examples discussed in the philosophical literature, the problem can be side-stepped by treating all predications as holding at the utterance time. In the linguistic literature, on the other hand, this problem has been discussed extensively. It is an issue to which I will have nothing new to add in this paper, and I will implicitly rely on the proposals that are documented in the cited publications.⁵

³ (Kamp 2001a,b).

⁴ A closely similar example is also discussed in van Genabith et al. (in press), though not as extensively as here.

⁵ See for instance Ogihara (1994, 1996); Abusch (1997); Kratzer (1998); von Stechow (2002).

- (iii) Each of the sentences of (1) attributes to Bill an attitudinal change: after the change Bill believes something which he didn't believe before it. As usual in cohesive discourse, the interpretation of the second sentence requires the interpretation of the first as context. Moreover, the interpretation of the second sentence must rely on the interpretation of the first in a way which is specific to sentences of this particular kind: The interpretation of the complement clause of the second sentence, which describes the content of the newly acquired belief, relies on the interpretation of the complement clause of the first sentence, which describes the content of the belief acquired on the occasion spoken of in that sentence. But in the case of (1) there is a further complication. The contextual background that is needed to interpret the second complement clause must provide information about the beliefs which the subject Bill holds at the time – some time on Tuesday – when the belief described in the second complement clause is acquired. But the first sentence doesn't tell us anything about Bill's attitudinal state at that time; it only says something about his attitudinal state on the preceding Sunday. Thus, in order to obtain the kind of background information that can support the interpretation of the second complement clause it is necessary to extrapolate from what the first sentence tells us about Bill's beliefs on Sunday to what his beliefs may be expected to have been at the relevant time on Tuesday. Like other bridging inferences the ones we will be considering have a good deal of plausibility. But their plausibility rests on discourse-independent knowledge, and so they introduce into the analysis of examples like (1) an element which is not found in cases where all attitudes attributed to a given subject by a succession of attributing sentences are attributed at one single time.

Sentence combinations like (1) are interesting not simply because they illustrate these different temporal aspects of the semantics of attitude attributing discourse, but also because of the ways those aspects interact, which each other as well as with non-temporal interpretation aspects. Such interactions between different interpretation principles are among the most difficult challenges for semantic theory. One of my aims in the discussion of (1) in Section 4 will be to make this interactional dimension of the example clearly visible.

3 Preliminaries

This section outlines the two extensions of “standard DRT” that were mentioned in the Summary. The extension required for the representation of attitudinal states will be outlined in 3.1, that needed for the treatment of presupposition in 3.2. The presentations will proceed by example and remain largely informal and illustrative. For formal definitions and further technical details the reader should consult the documents that are cited in the Summary.

3.1 DRT-based Representation of Attitudinal States

Consider the following scenario. Our subject Bill goes deer hunting in the company of his father-in-law. After reaching the forest where they hope to find what they are looking for they separate, taking different forks of the path which they have been following. At some point Bill sees (or at least he thinks he sees) a deer hidden in the undergrowth and barely visible. The effect of this visual experience is the belief that there is a deer in the location on which his eyes are focused; and this belief gives rise – not surprisingly given the purpose of the expedition – to the intention

to shoot the deer he sees (or thinks he sees). Belief and intention are propositional attitudes, each with its own “propositional content”. But these contents are *referentially connected* in that they are targeted on the same individual – the deer that Bill takes himself to be seeing.

The connection between these two contents is reminiscent of the referential connections between the contents of successive sentences of discourse or text, where the later sentences may be anaphorically linked to the earlier ones. When two “propositional” contents are linked in this manner, it is in general not possible to understand them as each determining a self-contained proposition. In the original versions of DRT this difficulty was finessed by focusing on the propositions expressed by the conjunctions of sentences with referentially linked contents. But in the present context this ‘solution’ is not open to us: since the two attitudes are of different “mode” – one is a belief and the other intention; thus they play distinct roles within the mental state of which they are components, and each plays the role it plays partly because of its own individual content – their contents must be separately identified. Kamp (2003) and van Genabith et al. (in press) present a formal semantics for such representations in which linked content representations receive separate but nonetheless referentially connected semantic values.) Here I will only show how the formalism is applied to the case of hunter Bill.

Combinations of propositional attitudes are represented as sets of representations of the individual attitudes, where the representation of each individual attitude consists of a ‘mode indicator’ – which determines whether the represented attitude is a belief or an intention etc. – and a DRS representing its content. Referential links between the attitudes that make up such a combination are represented by discourse referents that are shared between different content representations. For example, the combination of the described belief and intention of Bill can be represented as in (2)

$$(2) \quad \left\{ \left\langle \text{BEL}, \begin{array}{l} x \quad s \quad l \\ \text{deer}(x) \text{ loc}(l) \\ \quad \quad P(l) \\ n \subseteq s \\ s: \text{IN}(x,l) \end{array} \right\rangle, \left\langle \text{INT}, \begin{array}{l} e \\ e: \text{shoot}(i,x) \end{array} \right\rangle \right\}$$

Legenda: The content of the belief representation is that currently there is a deer x at a certain location l , which is further specified as satisfying certain characteristics P , about which more below. That the state of x being in location l is a current state is indicated by the condition “ $n \subseteq s$ ”, which says that the state s includes the time indicated by the indexical discourse referent n – as a constituent of the representation of the content of the given belief, n refers to the time at which this belief is being entertained. Note that it follows from this that if an agent entertains an attitude of such a form over some extended period of time, then the truth-conditional import of its content will change as time goes on. Suppose for instance that t_1 and t_2 are two distinct times within this period during which Bill entertains the belief represented in (2). Then at t_1 his belief has the content that the deer x is at the location l at t_1 , whereas at t_2 its content is that x is in l at t_2 . This phenomenon, of changing truth conditions determined by a constant representational form, will be important in the analysis of our example (1).

The content of the intention of (2) is that Bill, represented by the ‘self-representing’ discourse referent i , shoots x . The referential link between intention and belief is captured by their sharing

the discourse referent x . In this case the semantics of Kamp (2003) and van Genabith et al. (in press) treats the content of the belief as determining an independent proposition and the content of the intention as presupposing the belief content. The intuitive idea behind this is that it is only in combination with the belief that the intention has any real meaning.⁶

With regard to the visual experience which produces in Bill the two attitudes represented in (2) we can distinguish three possibilities: (i) there is a particular object he sees – which thus is the actual cause of his visual perception – and that object is a deer; (ii) there is an object which Bill is seeing, but that object is not a deer; (perhaps it is Bill's father-in-law, whose path has reconverged with Bill's and who is using the bushes for cover in pursuit of his own quarry); (iii) there isn't any object at all that is the perceptual cause of Bill's visual experience; the experience is an optical illusion, not a proper object perception in which the perceived object is at one end of the causal-perceptual chain and the resulting representation in the mind of the perceiver at the other end. Note that Bill's experience is in principle compatible with each of these three possibilities. But for someone who describes the case from his own, external point of view it may nevertheless be possible – and also important – to distinguish between them. (Which of the three possibilities is true can be quite important for what happens when Bill turns his intention into action).

Our representation formalism is designed to capture these distinctions. They are represented by means of *internal* and *external anchors*. In case there is an object which produces in Bill the experience of seeing it (through the kind of causal interaction that is typical of veridical visual perception), then the resulting discourse referent x is said to be *externally anchored* to this object; and in this case the content of any attitude which contains x as part of its content representation will be *de re with respect to x* – it will be a *singular* proposition, which attributes the conditions of the content representation in which x occurs to the object to which x is externally anchored. For example, if there is an external anchor for the discourse referent x in the intention representation in (2), then this representation determines the (doubly) singular proposition that there is an event of Bill shooting the object he is seeing.

In the third case, where the experience involves an optical illusion, there is no external anchor.

Bill's impression that he is connected by a relation of causal perception to an object which he thinks he is seeing is, we noted, compatible with each of the three mentioned possibilities. It is a further assumption of the present theory that whenever this impression is present, – i.e. irrespective of which of these three possibilities applies – it confers on the discourse referent which results from the perceptual experience a special psychological status, that of being an entity representation which is the causal product of a perception of the very object that it serves to represent. In the present theory this special status is represented in the form of internal anchors: discourse referents with this status are accompanied by internal anchors, those without this status do not. For instance, if as we have been assuming Bill attributes this special status to the discourse referent x which is the result of his visual experience, then his representation will involve an internal anchor for x (irrespective, to repeat, which of the possibilities (i)–(iii) happens to apply). In the cases (i) and (ii), where there is an external as well as an internal anchor, the propositional content of the belief will be *de re with respect to the object to which x is externally anchored*.

From a semantic point of view internal anchors act as presuppositions that their discourse

⁶ In general, when two content representations $K1$ and $K2$ share a discourse referent α which occurs in the universe of $K1$ while occurring in $K2$ only in argument positions of conditions, then the content of $K2$ is treated as presupposing the content of $K1$.

referents are anchored externally. When all internally anchored discourse referents of a propositional representation have corresponding external anchors, the representation will determine the singular proposition described above. On the other hand, representations containing discourse referents with an internal but no external anchor are cases of presupposition failure and do not determine any real proposition at all. This would be the case in particular for both the belief and the intention representation in (2) in the case of possibility (iii) (that of Bill's visual experience being an optical illusion).

As we have told the story of Bill's visual experience it is reasonable to assume that the resulting attitude complex involves internal anchors for both x and l . On this assumption, the representation of this complex in our revised formalism is the one given in (3)

$$(3) \quad \left\{ \left\langle [\text{ANCH}, l], \begin{array}{c} l \\ \text{loc}(l) \\ P(l) \end{array} \right\rangle, \right. \\
\left. \left\langle [\text{ANCH}, x], \begin{array}{c} x \quad s \\ n \subseteq s \\ s: \text{IN}(x, l) \end{array} \right\rangle, \right. \\
\left. \left\langle \text{BEL}, \begin{array}{c} \text{deer}(x) \end{array} \right\rangle, \right. \\
\left. \left\langle \text{INT}, \begin{array}{c} e \\ e: \text{shoot}(i, x) \end{array} \right\rangle \right\}$$

Legenda: In representations of this sort internal anchors are treated as separate components of the represented attitudinal state. Each representation of an internal anchor specifies in its first component not only that what it represents is an anchor (as opposed, say, to a belief or an intention; this information is given by 'ANCH'), but also which discourse referent that anchor is an internal anchor for. The second component gives information about how the entity represented by the anchored discourse referent is conceived – as the object that is being perceived, in some particular way, and with the properties that are attributed to it as an intrinsic part of the perception. Exactly what information this second component should contain, and in what form, are questions about which I will say next to nothing here. To give a sense of the complexity of this issue let me mention just one particularly problematic case. It concerns the internal anchor for l . The conditions which make up that component should capture the spatial relations between the location represented by l and the perceiving agent Bill. How agents represent locations in their

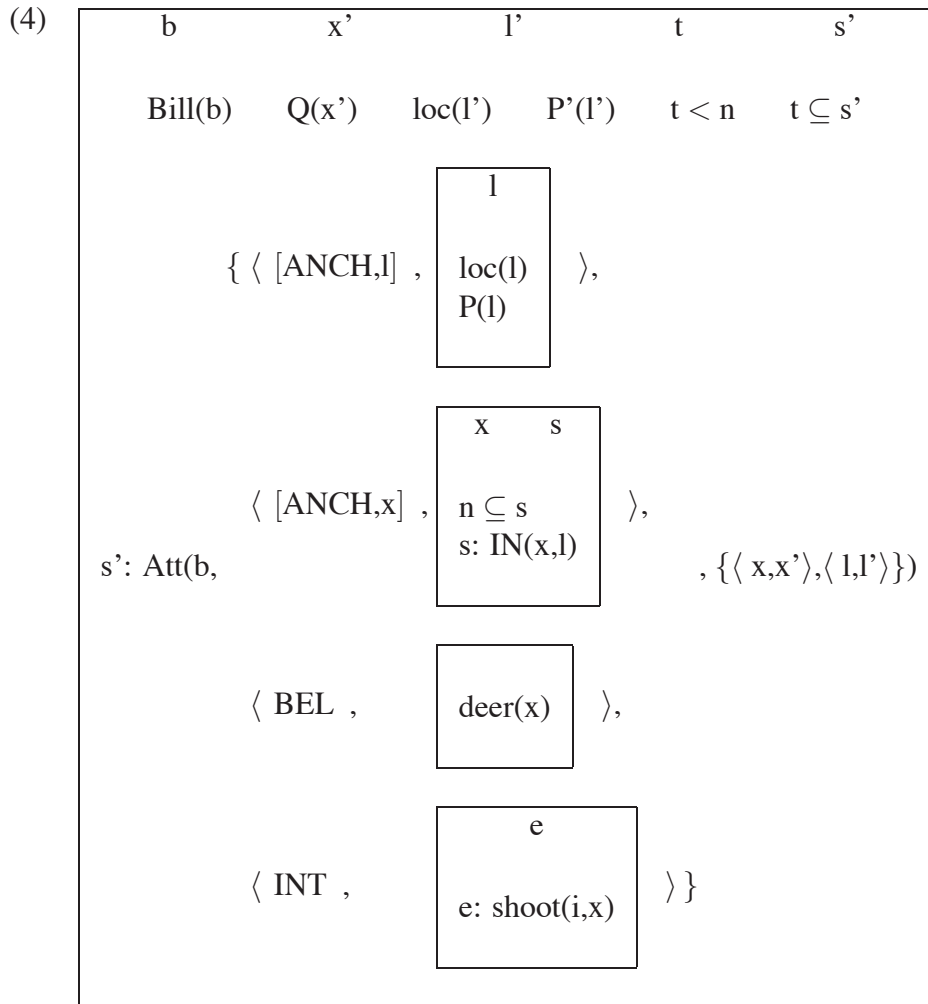
environment so that they are able to keep track of them when they move around themselves and could find their way to the represented location if and when they want to go there, is a matter with many intriguing cognitive aspects to it, few if any of which have been fully understood. I have finessed this particular problem here by assuming that in the internal anchor for *l* in (3) all this information is covered by the unanalysed predicate *P*.⁷

Attitude complexes like the one represented in (3) should be attributable to any agent at any time. So our formalism should have the possibility of expressing who it is that a given attitudinal state is the attitudinal state of, and when this individual is in the given state. In our formalism this is accomplished with the help of a 3-place predicate ‘Att’, whose arguments are (i) an agent; (ii) a part of the agent’s attitudinal state, specified as a set of connected attitudes (e.g. the one given in (3)), and (iii) a set of external anchors for all or some of the internally anchored discourse referents which are part of the second argument. The time at which a given attitudinal complex is being entertained is captured in the way familiar from the treatment of temporal reference in DRT: ‘Att’ is a stative predicate; thus, any predication involving ‘Att’ takes the form of a condition ‘*s*: Att(α, β, γ)’. This condition says that *s* is a state which consists in α being in a mental state which minimally contains the components specified in β , and where all or some of the internally anchored discourse referents of β have the external anchors given in γ . The time at which α is in the mental state thus characterised can then be specified via further conditions involving *s*. In particular, that α was in this state at a given time *t* can be captured by the condition ‘ $t \subseteq s$ ’. Thus, the proposition that Bill was at *t* in the state described in (3), with external anchors for both *x* and *l*, can be represented as in (4), where *x*’ and *l*’ represent (from a perspective external to Bill) the entities to which *x* and *l* are anchored.⁸

⁷ The question what information should be represented in the content representations of internal anchors has many ramifications. For ‘perceptual’ anchors, as we find in (3), most of these ramifications are closely related to problems for the theory of direct perception. But as I have proposed elsewhere, internal anchors also arise for discourse referents whose origin is not one of (purported) perception, and the information content of such anchors raises questions of quite different sorts. The general problems of the use and form of internal anchors is closely connected with issues of direct reference, in thought and – derivatively – in language.

A more detailed account of the propositional content of internal anchors is needed before it will be possible to draw a motivated line between information that should be represented as part of the anchor itself and information which should be part of the belief or beliefs that come about as a result of the purported causal relation between the agent and the entity that he takes the anchored discourse referent to represent. In this paper I will distribute information between anchors and beliefs on an ad hoc basis. In some cases the division I adopt may be only one of several intuitively plausible alternatives. The matter will be of no importance to what concerns us in this paper.

⁸ From the way in which I described the case represented in (4) the predicate *Q* cannot be determined. In fact, *Q* will vary depending on whether we are dealing with a case of possibility (i) or of possibility (ii). In the first case *Q* will be consistent with the property which Bill attributes to *x*’, whereas in the second case it won’t be. (In that case *Q* could be, for instance, ‘i’s father-in-law’). In case of possibility (iii) there will be no external anchor for *x* and thus also no condition involving *Q*. Given the way in which I have described the mental episode represented in (4) it seems reasonable to assume that irrespective of which of the possibilities (i)–(iii) applies with regard to *x*, there will be an external anchor *l*’ for *l*. The predicate *P*’ is to be seen as short for one or more conditions which fix the location *l*’ in terms that are independent of the attitudinal subject Bill.



3.2 Presupposition

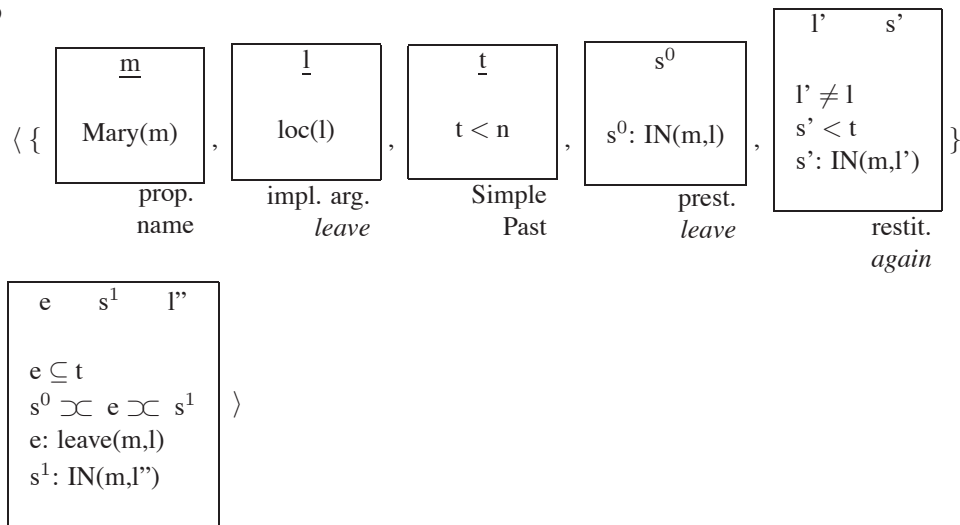
The second DRT extension we will need concerns the treatment of presuppositions. I here give just one example which will indicate the general features of this treatment. The sentence involved in this example – it is given in (5a) – is closely similar to the *that*-complement of the second sentence of (1). Like most sentences of English and other natural languages it generates several presuppositions, and these presuppositions are ordered hierarchically in the sense that some of them presuppose others in their turn. The triggers of the presuppositions of (5a) which are explicitly represented in the preliminary representation (5b) are the following:

- (i) The proper name *Mary*.
- (ii) The verb *leave*. *Leave* triggers two presuppositions. One of these is connected with the implicit direct object argument of the ostensibly ‘intransitive’ use of *leave* in (5a). The other presupposition is a “prestate presupposition”, an instance of the general kind that is associated with all verbs of change. The prestate of an occurrence of *leave* consists in the subject being in the place that she is described as leaving and the prestate presupposition is to the effect that this state obtains at the start of the leaving event described by the given occurrence of the verb.

- (iii) The past tense of *left*. In (5b) I have followed the view according to which certain tense occurrences carry a presupposition to the effect that the time of the described state or event is identifiable from the given context. (One of the intuitions supporting such a presuppositional account of tense is that out-of-the-blue utterances of a sentence like (5a) are perceived as strange or awkward because they lack a context which indicates what event time is intended.)
- (iv) Finally, (5a) contains the presupposition trigger *again*. The occurrence of *again* in (5a) is one of those which allows for both a *repetitive* and a *restitutive* interpretation. In (5b) I have opted for the restitutive interpretation, according to which the result state of the described event also held at some time before the event (and the event thus has the effect of ‘restituting’ that state). In the case of *leave* the restituted state is that of not being in the place that the described event is a leaving from.

(5) a. Mary left again.

b.⁹



Legenda: The top tier of (5b) consists of the presuppositions of (5a) while the lower tier is the non-presuppositional part. Note that some of the discourse referents occurring in the universes of DRSs which represent presuppositions are underlined. Such underlinings signify that justification of the presupposition involves finding an ‘antecedent’ for the underlined discourse referent. In DRT terms this means: finding a discourse referent in the context representation with which the anaphoric discourse referent is formally identified. In order that the chosen antecedent provide a proper justification for the presupposition containing the anaphoric discourse referent, the context must entail that it satisfies the condition or conditions contained in the presupposition. Presuppositions without underlined discourse referents have a purely propositional status. Such presuppositions are justified iff the context entails the propositions they express.

All presuppositions in (5b) come with labeling subscripts. These are shorthands for constraints, determined by the type of their trigger, on the form their justification can take.¹⁰ A

⁹ (“ \supset ” denotes *abutment* of two eventualities or periods of time. Thus, for instance, “ $s^0 \supset e$ ” means that s^0 ends at the very instant when e begins.)

¹⁰In particular, constraints on presuppositions with anaphoric discourse referents may specify in what form their antecedents must be available in the context, or may be made available by it: Should the context explicitly contain

proper taxonomy of justification constraints is one of the outstanding problems of current presupposition theory. In what follows I will, with few exceptions, ignore these constraints. In line with this the constraint labels will mostly be omitted.

The effect of presupposition justification on a ‘preliminary’ representation such as (5b) is that the presuppositions disappear from it, though the contributions they make to the truth conditions of the representation remain in force, since these are now taken over by the context (either as it was given at the outset or altered by accommodation). The non-representational part of the preliminary representation, which is all that remains of the preliminary representation after justification, constitutes the new information which the represented sentence contributes to the discourse.

In (5b) all presuppositions have been collected into a single set. This flat structure ignores presuppositional dependencies between the different presuppositions. Between the presuppositions of (5a) there are several such dependencies. For instance, the discourse referents *m*, *l* and *t*, all of which require resolution in context, also occur as arguments of conditions in the *again*-presupposition; so the *m*-, *l*- and *t*-presuppositions are not only presuppositions of the non-presuppositional part of (5b) but also of the *again*-presupposition. These dependency relations could be made explicit by further structuring of the presupposition set. But for the purposes of this paper there is no advantage in this, and so we will make do with simple set-representations like that in (5b).

Another simplifying feature of our example is that all presuppositions of (5b) are located at the same attachment site of the representation. This entails that when it comes to presupposition justification, all presuppositions have access to the same contextual information. In general this is not so. Preliminary representations may also contain presuppositions that are found in different attachment sites, and the justification of these presuppositions can make use of information that is expressed by parts of the representation that have scope over these respective sites.¹¹


4 The Example

4.1 The Representation of the First Sentence of (1)

At last we are ready to tackle example (1). Our central concern will be the interpretation of its second sentence (1.ii). The first sentence (1.i) is of interest primarily because of it provides the context for the interpretation of (1.ii) and we will pass over most of the questions of detail that arise in connection with the construction of its representation. (However, many of these same questions will also come up when we discuss the second sentence in Sections 4.2 and 4.3. Having read those sections, the reader should be able to fill in most of the details which our analysis of (1.i) in the present section leaves out.)

the antecedent as it is given? Or is it enough if it can be expanded with entailments which introduce the antecedent discourse referent? In that case, what kinds of entailments are permitted for this purpose? But these are only some among the many different constraints on presupposition accommodation which our current knowledge of these matters suggests, but most of which have not yet been identified with any precision.

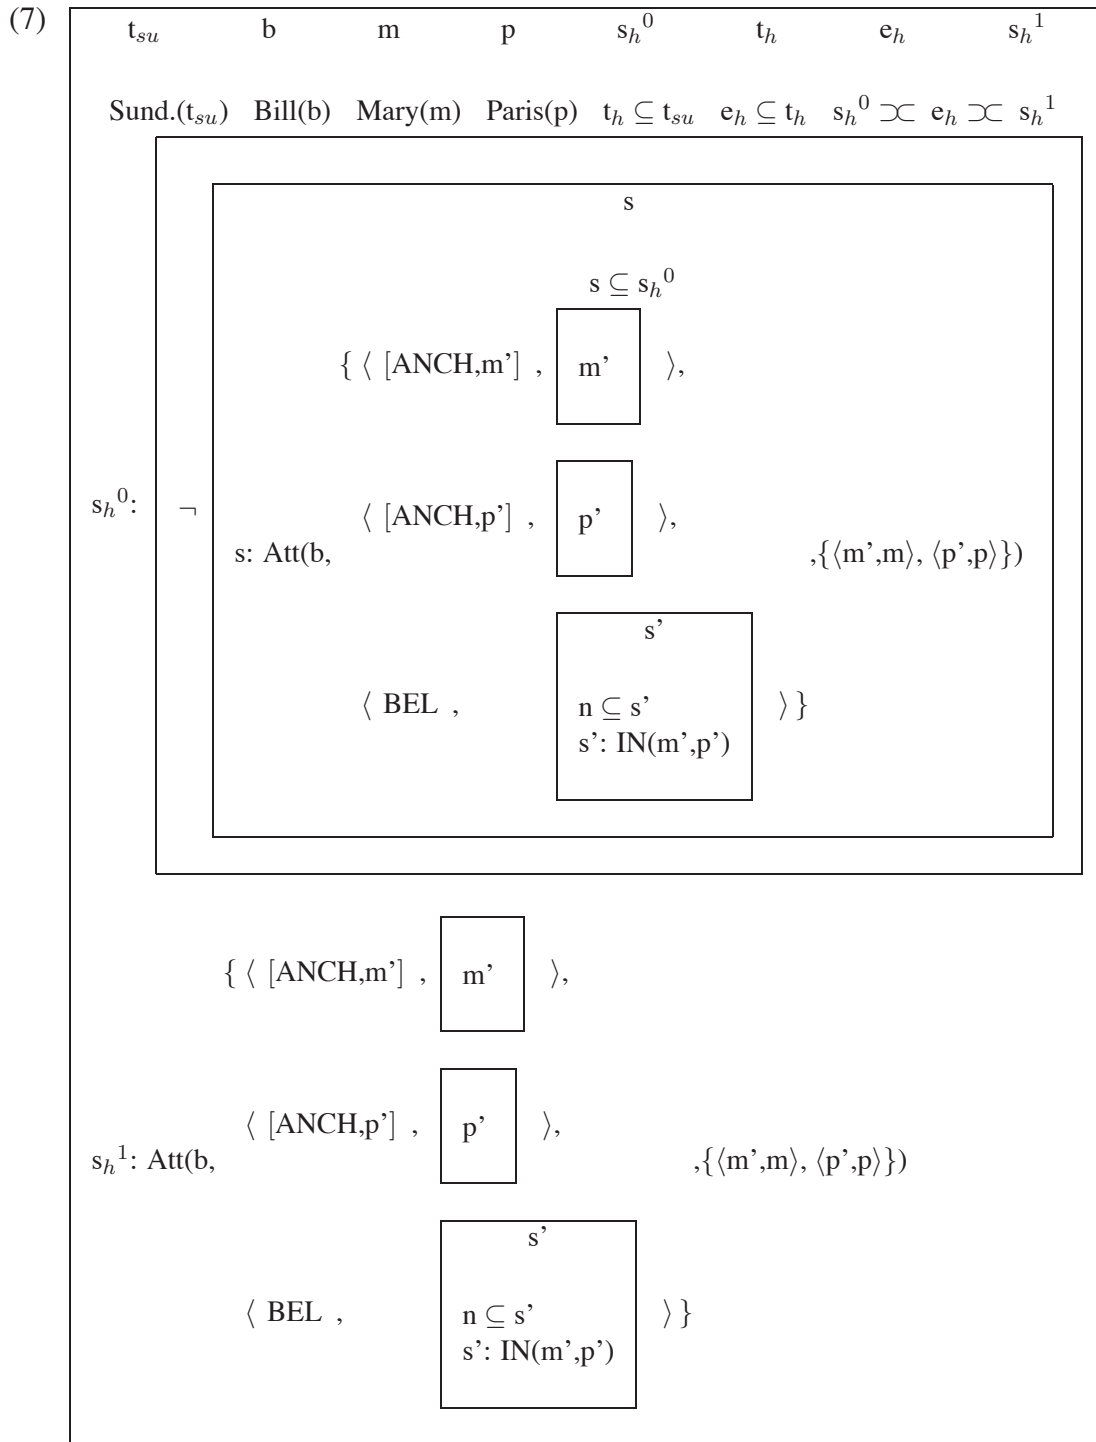
¹¹It is along these lines that dynamic theories of presupposition account for projection phenomena: a locally triggered presupposition is not perceived as a presupposition of the entire sentence in which the trigger occurs provided it can be justified on the basis of ‘contextual’ information provided by the sentence itself. See e.g. Heim (1983); van der Sandt (1992); Beaver (1995/2001) and Geurts (1995, 1999).



(6) represents Bill as having internally and externally anchored representations for the entities denoted by the proper names *Mary* and *Paris* which occur within the *that*-complement of (1.i) – that part which expresses the information of which (1.i) asserts that Bill acquired it on the given Sunday. This appears to be the only coherent interpretation of proper names occurring inside the descriptions of attitudinal states. It is important to note the difference between the two presupposition-like contributions of these name occurrences to (6). On the one hand they trigger – like any other occurrences of names, e.g. those of the proper name *Bill* and the name-like NP *Sunday* in (1.i) – the proper name presuppositions that are displayed in the upper tier of (6). But they are special in that they also contribute the mentioned anchors. Put in more general terms: the occurrence of a proper name (or other ‘directly referential’ NP) in the description of the content of an attributed attitude has two distinct presupposition-like effects:

- (i) Like any other occurrence of a directly referential NP it carries the presuppositional requirement that the context provide identification of its referent.
- (ii) As an occurrence within an attitude description it attributes to the possessor of the attitude an anchored representation of its referent.

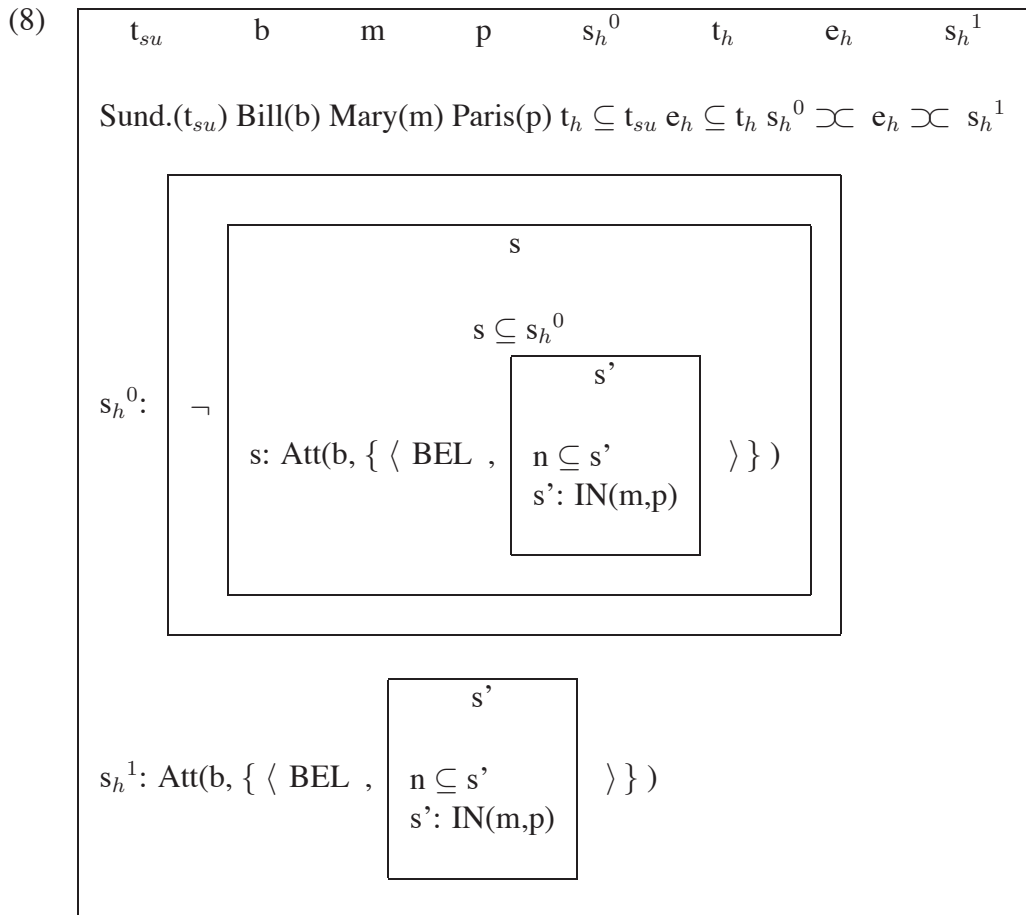
I assume that all presuppositions of (6) can be justified in the context in which (1) is uttered, either directly or after accommodation, and that (7) is the representation which results from their justification. In (7) all information that is represented in the various presuppositions of (6) has been included as non-presupposed (and thus as on a par with the non-presuppositional part of (6)). The motivation is that after justification all information contained in the presuppositions will be available as non-presupposed information, either because it was part of the context already or because it has been made available through accommodation.



I already noted the problems connected with the propositional content of internal anchors. For the present application (as for many others) the exact content of the internal anchors hardly matters. For such applications a notation which abstracts from the content of internal anchors will be adequate. Such an alternative notation can be substantially simpler than the one we have used so far, provided there also is no need to worry about discourse referents which have internal but no external anchors – in other words, if it can be assumed that all internally anchored discourse referents are externally anchored as well. This is an assumption that can be made in connection with our example (1) without giving away anything of what makes the example of

interest to us. In this alternative notation anchored discourse referents are eliminated, together with their internal and external anchors, and are replaced by the discourse referents representing the entities to which they are externally anchored. In other words, where the old notation has an anchored discourse referent x' with an internal anchor IA and an external anchor $\langle x', x \rangle$, the new notation will have neither IA nor $\langle x', x \rangle$, while the occurrences of x' in argument positions of conditions are all replaced by x . The new notation thus permits a discourse referent x which is ‘declared’ outside an Att-predication – i.e. which occurs in a universe that is not within the scope of the given occurrence of Att – to occur in positions that are part of this predication. The semantic import of such occurrences is essentially the same as the complexes of the old notation which they replace: Any attitude whose content representation contains occurrences of discourse referents that are declared in positions external to it is to be understood as determining a singular propositional content, which is *de re* with respect to the entities represented by these discourse referents.

Using the new notation we can simplify (7) to (8).



s'

$n \subseteq s'$

$s': \text{IN}(m,p)$


4.2 An interpretation for the *that*-complement of (1.ii)

The *that*-complement of the second sentence of (1) is repeated in (9)

(9) (that) she had left the previous day

I assume the following syntactic tree for (9)


(10)



Our next task is to construct a preliminary representation of (9) on the basis of (10). More accurately, we will construct a semantic representation for the subtree of (10) whose root is the IP *she left the previous day*. As always when semantic representations are constructed from syntactic trees, the construction requires lexical entries for each of the words and semantically relevant features that occur at the leaves of the tree. So we begin by listing these entries. In doing so we follow the left-to-right order in which the lexical items occur in the IP of (10).

Our first item is the pronoun *she*. The entry which we will use is given in (11.a)

(11) a. $\langle \textit{she}, \text{NP} \rangle$
 α



A few comments before we proceed to the next entry. The upper tier of (11.a) states that *she* is an expression of category NP and that it introduces some discourse referent α . The semantics is given in the lower tier. It consists of a representation that is composed of (i) a set consisting of a

¹³In (10) the ‘intransitive’ use of *leave* in (9) has been analysed as one that is syntactically transitive but involves a direct object that is phonologically null. At least on one occasion when I presented the material on which this paper is based this analysis has met with opposition. The opposition may well be justified: It may be that on balance we will obtain a better over-all theory of the syntax and semantics of ‘intransitive’ uses of verbs which also allow for a corresponding transitive use if we assume that these uses are genuinely intransitive in the sense that there is no direct object node in the syntactic structure of the clauses containing them. For my present purposes this question is not important. Anyone who prefers a syntactic analysis for (9) without a direct object node will find no difficulty in restating what I will have to say about the sentence below in a form which fits such an analysis.

Another decision which I have made here and which may offend syntactic persuasions is to describe the arguments of verbs (and other lexical predicates, although these play no role here) as ‘NPs’ rather than ‘DPs’. In fact, the nomenclature matters little here since the internal structure of such phrases plays no role in the problems we will be discussing.

single presupposition, and (ii) a non-presuppositional part. The presupposition is anaphoric with respect to its anaphoric discourse referent α which represents the pronoun's referent, and requires of α 's antecedent that it be a female person. (We ignore the complications connected with the possibility of using English *she* to refer to ships, mares and certain other kinds of non-persons.) The subscript 'pron.' is short for constraints on the possible resolutions of personal pronouns. (Compare the remarks in 3.2.)

The non-presuppositional part of the semantic tier of (11.a) is a shorthand for something more complicated. The story that it abbreviates is too long to tell in full here. Since it is of marginal relevance to the central issues of the paper, I confine myself to a few hints. The condition " $\lambda P.P(\alpha)$ " as non-presuppositional component of the semantics of referential NPs such as pronouns would be just what is wanted in a theory of semantic form construction which treats the combination of a noun phrase and a predicate to which it is an argument as involving a succession of two λ -conversions, in the manner first proposed by Montague. In a theory of this form the meaning of the sister node of the NP-node (i.e. the node of the predicate which has the NP for an argument) would be assigned a semantics of the form $\lambda\beta .K(\beta)$, with β a variable of the same type as α and $\lambda\beta .K(\beta)$ of the same type as P. Such an account works fine so long as the outermost λ -bound variable in the semantic form $\lambda\beta .K(\beta)$ of the sister to the NP node can be counted on to occupy within K the argument position of the NP. Making sure that this is always so can be awkward, however, especially for languages in which there is much variation in clause-internal word order and argument phrases often occur in positions other than their 'canonical' or 'base' positions. Some form of *linking theory* is required (most evidently in such languages but also in others) to determine which NPs fill which argument slots of the predicates with which they cooccur. Once such a linking theory is in place, it can be (and has to be) used as a guide in the construction of semantic representations to determine in which slot of the argument frame of the predicate we should insert the variable (or discourse referent) that is introduced by the NP. Under such conditions it proves more convenient to represent the semantic contribution of the sister nodes of NPs as involving stores which hold the argument positions of the represented predicate that still have to be filled and to describe the operation which combines the two contributions as involving the insertion of the discourse referent (or variable) introduced by the NP into the argument slot indicated by Linking theory, while removing the slot that is thereby filled from the store. A detailed account along these lines requires a fully explicit syntax as well as an explicit semantics. Sketches of a DRT-based syntax-semantics interface in which linking theory plays the indicated part can be found in several places, among them Kamp & Roßdeutscher (1994) and Kamp (2001a).¹⁴

Our next entry is for the past perfect. I have opted here for a treatment of the past perfect as a single tense form, instead of analysing it into smaller components (e.g. as the combination of a perfect operator and a simple past tense). The semantics of the past perfect proposed in this entry is based on the proposal for the essentially Reichenbachian treatment of the English tenses that can be found in Kamp & Reyle (1993), according to which the past perfect locates

¹⁴Had I assumed a syntax-semantics interface in which predicates and their arguments are combined in the manner of Montague the scare quotes around " $\lambda P.P(\alpha)$ " would have been unnecessary. Their presence is to remind us that the intended interface is strictly speaking not of this form. (I haven't bothered with scare quotes in the λ -terms in the following entries, even though similar caveats apply to their use in DRS construction. An elegant solution to the problem of "getting inside" a sequence of operators which prefixes a matrix representation is developed in Dekker (1993) within his version of Dynamic Semantics. But this is not the place for trying to adapt his ideas to the DRT-framework we are using here.

the predication expressed by the verb whose tense it is as having occurred at a time which is in the past of some “T(emporal) P(erspective) point”¹⁵ and where this TP point is itself situated in the past of the speech time n . The account assumes that the interpretation of an occurrence of the past perfect requires finding a suitable past time in the context which can play the role of TP point and then locating the predication as lying somewhere in the past of that time.¹⁶ Thus the status of the TP point is presuppositional. The subscript “TPpt” in the entry (11.b) stands for special constraints to which the choice of TP points is subject.

$$(11) \quad b. \langle \textit{past perfect}, \quad \text{VP-predicator} \rangle$$

$$\langle \{ \begin{array}{|c|} \hline \underline{tp} \\ \hline tp < n \\ \hline \text{TPpt} \\ \hline \end{array} \}, \lambda P. \begin{array}{|c|} \hline t \\ \hline P(t) \\ t < tp \\ \hline \end{array} \rangle$$

To simplify matters further we also treat the temporal adverbial *the previous day* as if it were an indivisible lexical unit, instead of analysing it as a complex expression with internal semantic structure (i.e. as a definite description with *day* as its nominal head, modified by the prenominal adjective *previous*).¹⁷ One important feature of the semantics of this phrase is the implicit argument of *previous*: In order to interpret an occurrence of this adverbial it is necessary to recover a day d from the context so that the referent of *the previous day* can be identified as the day immediately preceding d . Thus the semantics of *the previous day* also involves a presuppositional component, connected with this implicit argument.

$$(11) \quad c. \langle \textit{the previous day}, \quad \text{VP-operator} \rangle$$

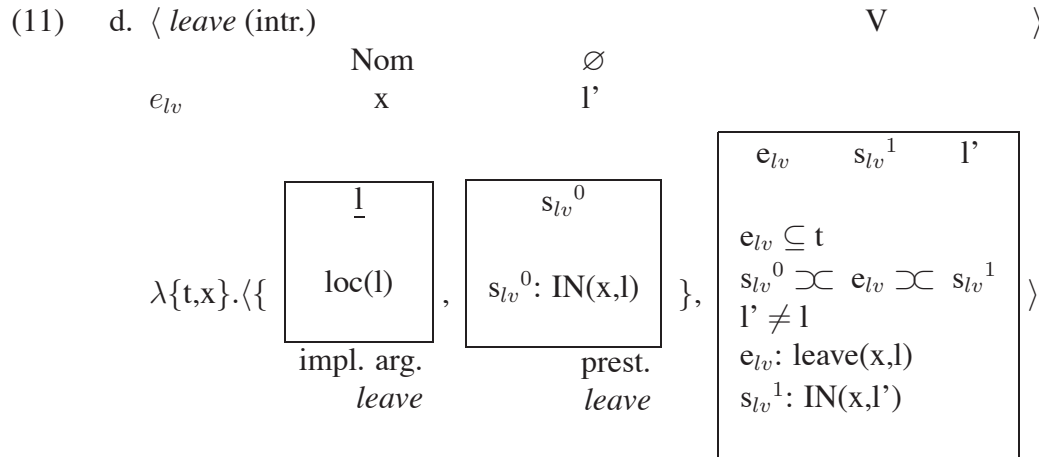
$$\lambda t. \langle \{ \begin{array}{|c|} \hline \underline{t'} \\ \hline \text{day}(t') \\ \hline \text{impl. arg.} \\ \textit{previous} \\ \hline \end{array} \}, \lambda P. \begin{array}{|c|} \hline t'' \\ \hline \text{day}(t'') \\ t'' \supseteq t' \\ t \subseteq t'' \\ P(t) \\ \hline \end{array} \rangle$$

Our last entry is that for the verb *leave*. We only specify a semantics for the intransitive use of *leave* which is found in (9). (A general entry would contain this use as one among several alternatives, with varying argument frames and varying meanings.) This ‘intransitive’ use comes with two presuppositions: (i) the prestate presupposition which this use of *leave* shares with its other uses (and, speaking more generally, with other change-of-state verbs) and (ii) the presupposition connected with its implicit argument.

¹⁵In this analysis of the past perfect the Temporal Perspective point plays the same role that is played by what Reichenbach calls “Reference Time” in his account of the past perfect. (Reichenbach 1993).

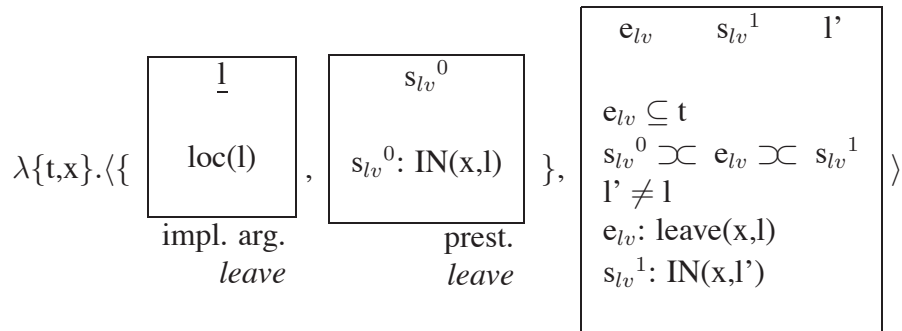
¹⁶This formulation fails to account for complications having to do with “sequence of tense”; since these are irrelevant to this paper, they are being ignored.

¹⁷For a discussion of temporal definite descriptions like this one, which takes the compositional aspect of such phrases seriously, see Kamp & Schiehlen (2002).

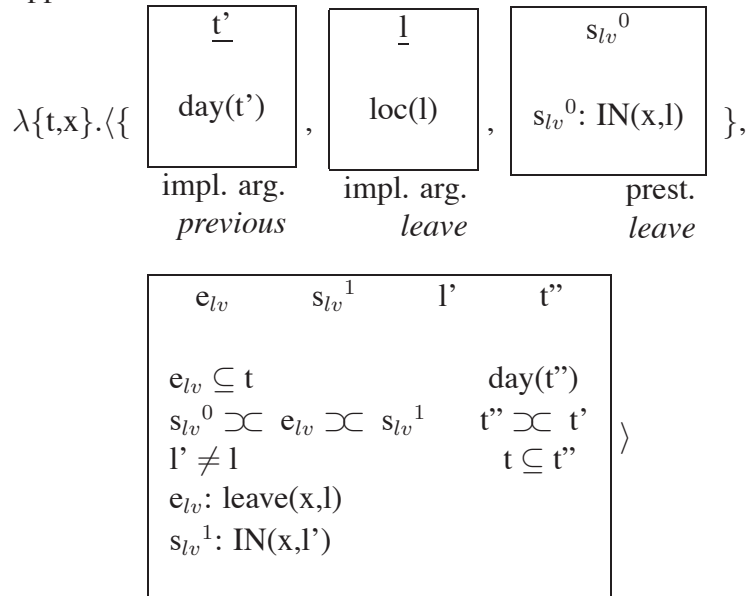


(12) presents the semantic representations of the complex nodes of the IP of (10). We proceed bottom-up, from the representation for the lower VP node in (12.a) to that of the IP node in (12.d). I have refrained from stating the composition principles according to which these representations are obtained from those of their daughter nodes and refer the reader to relevant literature.¹⁸

(12) a. lower VP node:

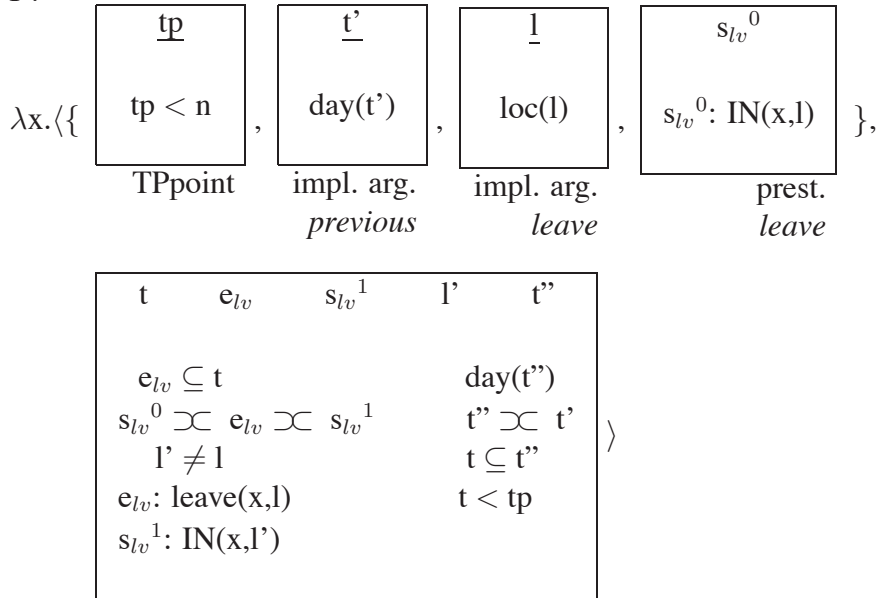


b. upper VP node:

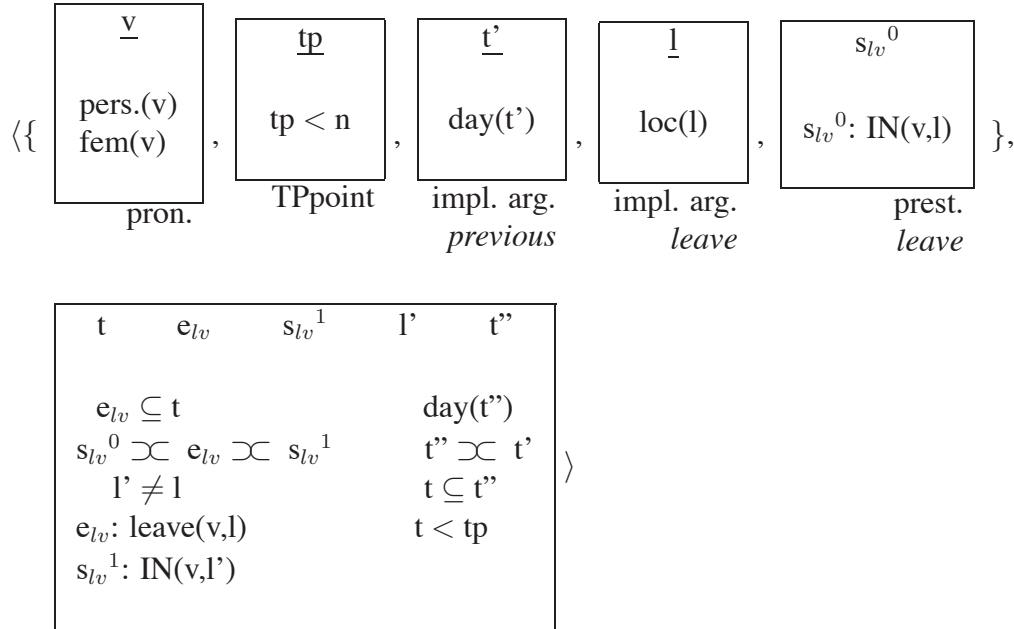


¹⁸There is to my knowledge no place in the existing literature where these principles are given in exactly the form in which they are needed here. Discussions of similar principles can be found in Kamp (2001a).

c. I':



d. IP:



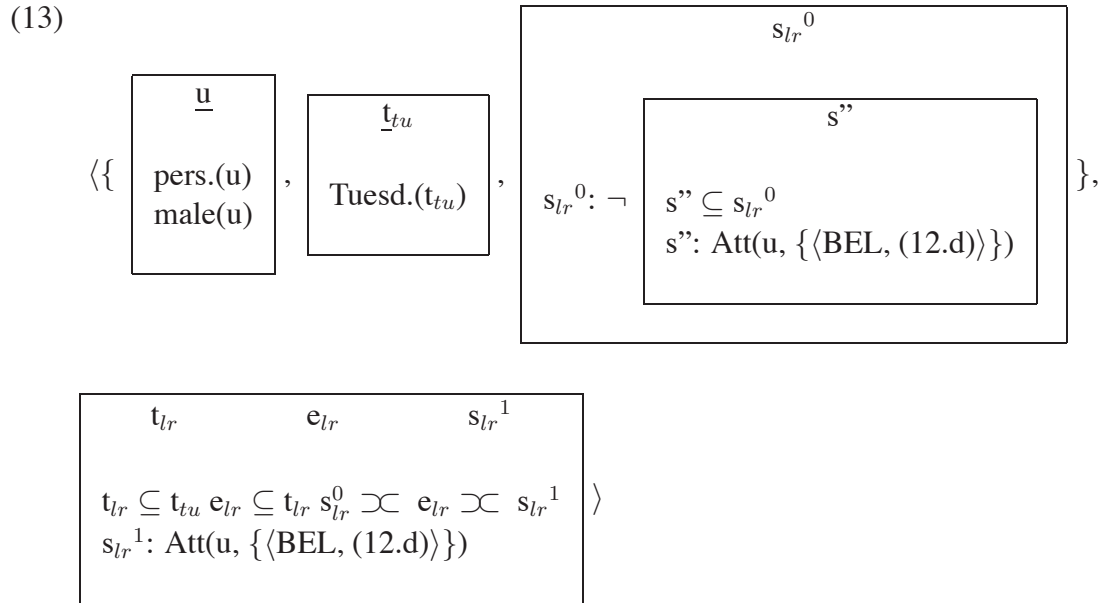
(12.d) will be used as representation of the complement clause of (1.ii) when we construct its preliminary representation in the next section.

4.3 Preliminary representation of (1.ii) and its resolution

The preliminary representation for (1.ii) has the same general form as that for (1.i). Once again we are dealing with a sentence which reports a change in Bill's attitudinal state. So once again the representation of the *that*-clause, i.e. (12.d), enters into the representation at two places – into the representation of the belief that is part of Bill's resulting attitudinal state and into the representation of the prestate of *learn*, which consists in Bill not holding that belief.

I will forego the syntax-semantics interface details involved in obtaining this preliminary representation. The remaining steps which lead from the representation (12.d) of the embedded

IP to that of the full (1.ii) involve little that has not yet been illustrated in the construction of (12.d). There is one exception, viz. the binding of the embedded tense by the tense of the matrix verb. But as announced in Section 2, this is an issue where I am relying entirely on existing work of others and a proper exposition of how the embedded tense is bound by the matrix tense according to any of the existing proposals would involve us in much extra detail that would detract from our real concerns. So you have decided to let this matter rest, and we proceed straight to the preliminary representation of (1.ii).



(13) has explicit representations of a total of eight presuppositions, the three displayed in (13) itself and the five of the embedded (12.d).^{19,20} It is important to note that the presuppositions

¹⁹Note well that we are dealing with 8 presuppositions, not 13. The fact that (12.d) occurs in (13) in two different places does not mean that it introduces 2×5 presuppositions into (13). Since the two occurrences of (12.d) in (13) stem from the same constituent of (1.ii), each of (12.d)'s presuppositions counts as only one, which demands a single resolution. In particular, there must be single resolutions for the anaphoric discourse referents in these presuppositions, with antecedents which provide a justification of the presupposition in question that are valid for both positions in which it occurs. In fact, the matter is more complicated than this description suggests. For more on this see Kamp (2001a).

²⁰It could be argued that there is one presupposition of (1.ii) which (13) fails to display. The verb *learn*, it might be said, is factive in the sense that the new information which *learn* describes its subject as getting is presented as true (or 'as fact'). Factivity of stative attitude verbs such as *know* or *regret* is widely regarded as presuppositional, and much the same considerations that have been adduced in support of that view equally support a presuppositional treatment of the factivity of the 'inchoative' epistemic verb *learn*. Personally I am inclined to think that *learn* is factive, but have nevertheless chosen to ignore this aspect of its semantics here.

The factivity of *learn* is connected with another aspect of its meaning. "x learns that p" seems to entail that the information that p did not just become available to x, but that x accepted this information, adopting the belief that p is true. In my analysis of (1) I have assumed this acceptance aspect as part of the meaning of *learn*. It could be argued that the acceptance aspect of *learn* and its factivity go hand in hand. If that is so, the policy I have adopted here – of taking the one aspect on board but not the other, is of course not quite right. But it is a decision which has no serious consequences for the issues that concerns us.

The factivity issue might have been raised also in connection with the verb *hear* as it occurs in (1.i). According to my own intuitions *hear* (in the sense of 'come to be informed about' in which it is used in (1.i)), is less clearly factive than *learn*, but quite possibly it has a factive use as well.

are found in two different attachment sites. For justification of the three displayed in (13) the only contextual information available is the primary context provided by (8). In contrast, the justification of the five presuppositions of (12.d) may resort to ‘secondary’ context information as well. However, as noted in Section 1, the question how (8) can provide the needed secondary context information is one of those which embody the interest which (1) holds for us and which we will have to look into carefully.

Six of the eight presuppositions of (13) are anaphoric, with the anaphoric discourse referents u , t_{tu} , v , tp , t' and l . For two of these, u and v , the context (8) offers only one possible resolution:

- (i) $u := b$;
- (ii) $v := m$;

The story about the other four is a little more complicated, and it is a somewhat different story for each. First t_{tu} . Adverbial uses of names of days of the week, such as that of *Sunday* in (1.i) and of *Tuesday* in (1.ii), are anaphoric in roughly the following way. The interpretation of such a name occurrence involves the identification of some other day d which can serve as the ‘origin’ from which the reference of the given occurrence can be computed. The day referred to by the name is then either the first day after d that fits the name (e.g. the first Tuesday after d) or the last one before d . Which of these two is the intended referent depends on further factors, from which it is possible to tell whether the state or event described by the clause in which the name occurs is located after d or before it. Often the origin is the day of the utterance time. But it can also be some other contextually salient day. In the case of the *Tuesday* of (1.ii) the natural choice for d is the *Sunday* referred to by *Sunday* in (1.i).²¹ (What day that is, is yet another question. It is one which we passed over when discussing the interpretation of (1.i), and it is a matter which cannot really be settled, the interpreter of (1) must simply accommodate the referent by assuming that whoever produced (1) must have been talking about some particular *Sunday*.) I am assuming further that the event described in (1.ii) follows the one described in (1.i).²² On these assumptions the referent of *Tuesday* is the Tuesday immediately following the day referred to by *Sunday* in (1.i), and that is what we will take it to refer to throughout the remainder of the paper.

Next the location discourse referent l . It might easily appear as if its resolution is as predetermined as those of u and v ; for these cases are similar in that for each of these three discourse referents there is exactly one antecedent in the discourse context. In the case of l this is p , the only discourse referent in the discourse context representation (8) which satisfies the constraint that what it represents must be a location. Indeed, p seems a natural resolution for l – in other words, the leaving spoken of in (1.ii) is a leaving from Paris – and that is the resolution which I will assume in this section. However, in Section 4.4 we will see that this is not the only possible interpretation for l .

Two anaphoric discourse referents remain, tp and t' . For these resolution is not unequivocal either. t' could be resolved either to t_{tu} or to t_{su} . And in case t' is resolved to t_{su} , tp could also be resolved to either t_{tu} or t_{su} (although this latter choice makes no difference to the resulting

²¹For some more discussion on this see Kamp & Schiehlen (2002).

²²This interpretation is strongly suggested by the tenses of the matrix verbs in (1): the succession of the two simple past tenses *heard* and *learned* of (1) carries a presumption that the event of the second sentence followed that of the first. But many discussions in the literature show how delicate these matters can be. See Lascarides & Asher (1993); Asher & Lascarides (2003).

sentence interpretation). In this section we focus on the option of resolving t' to t_{tu} . This means that *the previous day* is interpreted as referring to the day immediately before the day represented by t_{tu} – in other words, to the Monday between that Tuesday and the Sunday represented by t_{su} . In this case the only coherent resolution for tp is that which identifies it also with t_{tu} .

With these additional decisions our resolution set is extended to:

- (iii) $t_{tu} :=$ “the Tuesday immediately following t_{su} ”;
- (iv) $l := p$.
- (v) $t' := t_{tu}$;
- (vi) $tp := t_{tu}$.

4.4 Other possible Interpretations of (1)

Before we enter more deeply into the reconstruction of the interpretation determined by the resolution equations (i)–(vi), we will quickly review what other interpretations are possible for (1). In doing that I will not question the interpretation of (1.i) given in (8), but only look at alternative interpretations of (1.ii) relative to this interpretation of (1.i). About these alternative interpretations I will have nothing more to say in the sequel. So the reader who isn't particularly interested in this matter and more concerned not to lose the thread of the main argument of this paper, can skip this section and move straight to 4.5.

We noted that for some of the anaphoric discourse referents of the preliminary representation of (1.ii) the context that is provided by (1.i) offers only one possible resolution. For *he* it is 'Bill', for *she* it is 'Mary'. The matter of *Tuesday* is more complicated. I do not think that interpretations other than the one we have been considering can be excluded absolutely. It seems just possible to interpret *Tuesday* as referring to the last Tuesday before the utterance time in cases where that is a different Tuesday from the first one after the Sunday that was mentioned in the first sentence. (Note that when *Sunday* in the first sentence is itself interpreted as the last Sunday before the utterance time, and the utterance time is later in that week – e.g. on Friday or Saturday – then this second interpretation of *Tuesday* will *de facto* coincide with the one assumed in Section 4.3.) A second possibility would be for Tuesday to refer to the Tuesday immediately before the Sunday of the first sentence. As we noted in 4.3, this interpretation is not a very likely given the tenses of the two matrix clauses. But again, it is difficult to exclude it categorically. Since these interpretations seem so marginal, however, I won't consider these possibilities further.

I already mentioned in 4.3 an alternative interpretation for *the previous day*. On this interpretation the phrase would refer not to the day before Tuesday but to the day before Sunday. In order that this option lead to an intelligible interpretation we need a different resolution of the implicit location argument l of *leave*: When the leaving took place the day before Mary was in Paris, then presumably it was not a leaving from Paris but a leaving for Paris, and thus the place that Mary left must have been a place that is different from Paris. Since no place other than Paris has been mentioned, such a resolution of l requires accommodation. It should be stressed in this connection that there are not just speakers for whom this second interpretation of (1) is possible, but even some for whom it is the more salient one. One important implication of this is that the principle, to my knowledge first stated explicitly in van der Sandt (1992), that resolutions of anaphoric discourse referents which do not involve accommodation – cases of *binding* in van der

Sandt's terminology – have preference over resolutions which require it cannot be universally valid in this simple form. (Cf. Beaver 1995/2001, 2005) for earlier indications that there is more to this principle than the simple formulation suggests.) There is no doubt that something like this principle must be true, but evidently a more refined formulation is needed. Such a refined formulation will have to pay attention among other things to the types of different anaphoric presupposition triggers. In particular, a distinction will almost certainly have to be made between overt pronouns and implicit arguments such as that of intransitive *leave*. (See also the remarks in Section 3.2 on trigger-dependent rules of presupposition justification.)

In my estimation this exhausts the range of interpretations for (1.ii).²³

4.5 Assumptions about the persistence of other agents' beliefs

I will take it for granted that the resolutions (i)–(vi) yield full justifications for the presuppositions that contain the discourse referents for which they provide anaphoric discourse referents, i.e. that (8) entails each of the conditions that are contained in these presuppositions when the anaphoric discourse referents are replaced by their chosen antecedents. This leaves us with two remaining justifications, for the non-anaphoric pre-state presuppositions of *learn* and *leave*. It is the justification of the second of these presuppositions, the prestate presupposition of *leave*, which brings into focus those issues which motivated the choice of (1), and it is to this justification that the remainder of this section and all of Section 4.6 will be devoted. Once the justification of this presupposition is in place, the justification of the prestate presupposition of *learn* will be unproblematic.

Given the resolution of *v* to *m* and of *l* to *p*, the prestate presupposition of *leave* says that Mary was in Paris immediately before the occurrence of the leaving event described in (1.ii). We should keep in mind, however, that neither the event nor its presupposition are presented as mere facts but as contents of Bill's belief state. Contextual information which justifies the prestate presupposition must therefore also be information about Bill's belief state, and it must be information about his belief state just before the time when on Tuesday he learned about Maria's leaving. But as things stand (8) provides no direct information about Bill's beliefs at that particular time. The attitude which (8) does ascribe to Bill is a belief which he holds as a result of the event of his hearing on Sunday about Mary being in Paris. Since it was the result of that event, Bill must have had it for at least some time on Sunday. But how long Bill stuck to this belief is something about which the interpreter has no direct information. However, beliefs, like many other propositional attitudes, have a tendency to persist, and so it is a natural assumption that Bill is still in this belief state when the new information reaches him on the following Tuesday. In fact, this assumption is so plausible that an interpreter of (1) would hardly be aware that he was making it. We will see presently that it doesn't really solve the justification problem posed by *leave*'s prestate presupposition, but it is a natural assumption nevertheless and it appears that a discourse like (1) exerts a strong pressure on the interpreter to make it.

The assumption that Bill's belief about Mary's presence in Paris persists until Tuesday may

²³Claims of this sort are notoriously prone to error – one always has to be wary of special contexts, which would open the way to other interpretations, but which are so unusual that a normal interpreter would not assume them unless the context was explicitly presented to him.) Actually demonstrating that a given set of readings exhausts the interpretational possibilities for a given sentence or discourse on the basis of theoretical assumptions is almost always very hard, and well beyond our current capacities. So, as things are, such claims have to be handled with the greatest circumspection.

seem innocent enough at first sight; but what really does this assumption amount to? What *is* the belief about Mary that is supposed to have persisted until the time when Bill got the new information on Tuesday? There are two different ways in which the persistence of Bill's belief about Mary being in Paris can be construed. One possibility is that Bill retains the belief that Mary was in Paris at the time when he was informed about this, i.e. the belief that Mary was in Paris on Sunday. But it is also possible for the belief to have persisted as the belief of Mary being *currently* in Paris. The difference between these two forms of persistence is as important as it is evident. From a truth-conditional perspective persistence of the second kind isn't persistence in the strict sense of the word. As noted in Section 2.1, the truth-conditional content of a belief about what is currently the case evolves with the time at which the belief is entertained. For instance, on Sunday the content of Bill's belief that Mary is currently in Paris is the proposition of Mary being in Paris on Sunday, whereas the content of the "same" belief on Tuesday is that of Mary being in Paris on Tuesday. Note, however, that in the representational format we have been using this evolution of truth conditions goes hand in hand with invariance of representational form. For it is the same content representation, repeated in (14), which carries these different truth conditions depending on whether it represents a belief entertained on Sunday or a belief entertained on Tuesday.

$$(14) \quad \boxed{\begin{array}{l} s' \\ n \subseteq s' \\ s': \text{IN}(m,p) \end{array}}$$

As opposed to this second kind of persistence, the first kind of persistence we mentioned is one which succeeds in preserving truth-conditional content of the belief that Mary was in Paris on Sunday. But this stability of truth conditions can, in the given case, only be achieved at the cost of a change in representational form: If we assume that Bill initially represented the belief he acquired on Sunday in the form in which it is given in (8) and repeated above as (14), its representation must have changed in the meantime if it is still to represent, when entertained at the later time on Tuesday, the proposition that Mary was in Paris on Sunday. In the abbreviated notation we first used in (8) and have been using since, the new, truth condition preserving representation can be presented in a pleasingly compact form:

$$(15) \quad \boxed{\begin{array}{l} t' \quad s' \\ t' \subseteq t_{su} \\ t' \subseteq s' \\ s': \text{IN}(m,p) \end{array}}$$

In this particular instance, however, we do well to remind ourselves of the abbreviatory nature of this notation. As the representation of a propositional attitude, (15) conceals the internal and external anchors that would be explicitly represented in the earlier notation. A representation in that notation would have displayed among other things an internal anchor for a discourse referent t'_{su} together with the external anchor $\langle t'_{su}, t_{su} \rangle$. So there is more to the transition from

the belief represented in (14) to the one represented in (15) than meets the eye in a superficial comparison of (15) and (14).

Yet we must be careful not to overstate this point. It is likely that Bill has an anchored representation for the mentioned Sunday, and quite possibly also one for the particular time on that Sunday when he got the information about Mary being in Paris. In that case most of what is needed to realise the transition from (14) to (15) is already in place and all that remains is to change the indexical link of the state *s*' to the time *n* by the non-indexical link to the anchored representation of the given Sunday (or of the particular time on Sunday when Bill's first belief change occurred).

For the interpreter of (1) there is no conclusive way of telling which of the two persistence assumptions about Bill's belief is the (more) appropriate one. In the next section we will spell out the implications of either option. I will use the remainder of the present section for two rather speculative remarks on the reasons why and the ways in which beliefs about the current state of the world are transformed into beliefs about what the world was like at some earlier time. First, an observation that goes back more than two decades and probably more²⁴: The time it takes for the belief that a certain condition currently obtains to be transformed into the belief that this condition obtained at some earlier time (where this earlier time may be fixed with more or less accuracy) will vary as a function of what the condition is. When I observe that there is a dead blackbird lying on the roof and form the corresponding belief that this is currently the case, then this belief is likely to survive for longer than when I observe that there is a living blackbird sitting on the roof. We humans are equipped with elaborate networks of expectations as to how long different conditions are likely to last (in various circumstances) and these expectations guide us in our doxastic behaviour. And the longer the 'life expectation' we associate with a given condition, the longer it will take for us to abandon the belief that the condition currently obtains after we found out about its obtaining. And what is true for our own beliefs about such conditions also applies to our expectations about the doxastic behaviour of others. For instance, had the first sentence of (1) spoken of Mary being in the kitchen rather than in Paris, we would be much less inclined to make the assumption that Bill maintained his belief that she was still there until Tuesday.²⁵

The second observation concerns the notion of doxastic strength. Some beliefs are held more firmly than others – we are more resistant to giving them up in the light of contrary evidence, we attach higher probabilities to their propositional contents, we are prepared to place higher stakes on outcomes which depend on their being true, etc. (I do not mean to imply that all these criteria must necessarily coincide in what they tell us about the degrees of confidence of which they are the supposed manifestations, but they are all important.) Once we are prepared to distinguish between different degrees of doxastic strength, we are in a position to distinguish between instantaneous transitions from beliefs about current states to beliefs about past states and transitions which are gradual. In particular, it is now possible to conceive of Bill's belief that Mary is currently in Paris as involving, over the period from Sunday till Tuesday, a gradual

²⁴Early work on this topic of which I am aware is that of D. McDermott. But I suspect that in AI and/or in the Philosophy of Mind an awareness of the importance of this factor may go back to an even earlier date.

²⁵It is worth noting that this variant of (1) has a peculiar ring to it. Apparently the form of (1) impels the interpreter to assume that the current state belief acquired on Sunday persists until the time of the belief change on Tuesday. For the just mentioned variant the persistence of this belief over a period of the given duration is inherently less plausible than it is for (1) itself and this seems to be the reason why the variant sounds odd in a way that (1) does not. See also Section 4.7.

decay rather than a sudden switch from belief to non-belief. In fact, the more plausible picture would now be that Bill's belief that Mary is currently in Paris coexists with his belief that she was in Paris at the time when he was informed about this on Sunday and that the strength of the former slowly decays while the latter continues in full force.

Taking degrees of doxastic strength into account would arguably lead to a deeper and more natural analysis of the interpretation of (1) than I am able to give in this paper. But degrees of belief cannot be represented in the formalism we are using and to modify the formalism so that variations of doxastic strength can be adequately represented is a non-trivial matter. There is no question of developing such a modification here and all I can do is leave this as a task for the future.²⁶

I should add that I am not persuaded that matters of doxastic strength are essential to the *semantics* of discourses like (1), even if it is relevant to the extra-linguistic reflections that their semantic representations are likely to set in motion. The reasons for these doubts will become visible in section 4.7.

4.6 Implications of the two Persistence Assumptions

Let us assume that both persistence hypotheses we discussed in the last section – the hypothesis that Bill's belief at the relevant time on Tuesday has the content given by (14) and the hypothesis that it has the content given in (15) – are possible accommodations which an interpreter might entertain in his attempt to make sense of (1.ii). For all that has been said so far, either one is a candidate. But which one the interpreter adopts makes a good deal of difference to how he will understand what happens when Bill learns about Mary's departure.

First consider the assumption that Bill retained the belief he formed on Sunday in the form given in (14). Then the new information he received on Tuesday was in contradiction with what he believed at that time. For he still believed that Mary was in Paris, which is obviously inconsistent with the discovery that she left Paris the day before. So the interpreter who assumes Bill's belief to have persisted in this form must also assume that the new information which Bill got on Tuesday led him to a belief *revision*: Apparently Mary left earlier than Bill thought. So that thought had to be given up and replaced by one which is compatible with the new information. Note well, though, that the new belief state which results from this revision still involves persistence of the condition represented in (14) in a more limited form. If Mary left Paris on Monday then she must have stayed in Paris until some time on Monday. Thus (14) remained true until some time on Monday, rather than becoming false already on Sunday, which would have been possible too.

It should be emphasised that the belief revision of which we have spoken is a revision of Bill's beliefs, not of the assumptions made by the interpreter. It is Bill who is forced on Tuesday to abandon the belief of the form (14) which he held until then and to replace it by the belief that

²⁶One possibility which comes to mind is that of extending the formalism with an open-ended set of doxastic indicators '[BEL, μ]' where the μ are terms denoting measures of the doxastic strength. (The measures might be real numbers, but they could also belong to some other type of confidence scale.) The further question then is what kinds of terms should be admitted in the extension. In order to state general hypotheses about the decline of belief strength over time it would seem desirable to admit variables ranging over the confidence scale. A further question concerns the types of complex terms that should be allowed, and the operations that may appear in those terms. But whatever the details, it is clear that such a formalism would constitute a very substantial change from the one we are using. And so long as no decisions have been made about its exact properties, there isn't very much that we can say about its impact on semantics, or, for that matter, about its implications for the analysis of (1).

Mary remained in Paris only until some time on Monday. Only if the interpreter himself were to assume, on the strength of his interpretation of (1.i), that Mary was still in Paris on Tuesday would he have had to backtrack and correct his own perception of the facts. But there is no obvious reason why even an interpreter who adopts the persistence hypothesis under discussion would make this additional assumption as well.

The second persistence hypothesis attributes to Bill, at the time on Tuesday when he received the information that Mary had left, the belief that Mary was in Paris on Sunday. This belief was not contradicted by the new information which reached him and so there was no cause for belief revision. It should be observed, however, that this hypothesis too isn't quite what justification of the prestate presupposition of *leave* requires either. This time the belief which it attributes to Bill just before he gets the new information is too weak – it is just the belief that Mary was in Paris on Sunday, not that she remained there until Monday. So once again Bill must accommodate. However, all that is required of him on this hypothesis is an addition to his beliefs, without the need to discard any he had.

The differences in the interpretations which result from the two hypotheses are significant, but they are limited nonetheless. The interpretation processes that are guided by the two hypotheses converge in that they end up attributing to Bill the same final belief state (the one in which he is said to have been as the result of his second belief change on Tuesday). They differ only with regard to what they entail about his beliefs immediately before that. (16) represents what is common between the two interpretations of (1.ii) – Bill's final belief state, together with the condition that up to the time when he received the new information he was not in that state.^{27,28}

²⁷Some may doubt whether prestate presuppositions are presuppositions at all. Or they may – a more moderate disagreement with the position I have assumed here – take it that they are conditions which function as normal entailments of positive declarative occurrences of the change-of-state verbs which trigger them and behave as presuppositions only when the trigger occurs within the scope of some operator, such as negation or a question. (For an early defence of a position similar to this see Wilson (1975).) Someone who holds this view of prestate conditions might either want to maintain that the persistence hypotheses we are discussing are irrelevant to the interpretation of (1) or he would want to describe their role in different terms. I do not think however, that this issue matters. In particular I do not believe that it affects the question whether the hypotheses play any role at all in the interpretation of sentence sequences like (1) could be decided on the strength of the status which we ascribe to prestate conditions. More on this in the next section.

²⁸(16) is dependent on the representation (8) of (1.i) which serves as context for the justification of the presuppositions of (13). I have not repeated (8) here for reasons of space.

$$(16) \quad \begin{array}{l} \text{u} \quad t_{tu} \quad t_{lr} \quad s_{lr}^0 \quad e_{lr} \quad s_{lr}^1 \quad v \quad l \\ \\ \text{u} = b \quad \text{Tuesd.}(t_{tu}) \quad t_{lr} \subseteq t_{tu} \quad e_{lr} \subseteq t_{lr} \quad v = m \quad l = p \\ \quad t_{su} < t_{tu} \quad \quad \quad s_{lr}^0 \supset e_{lr} \supset s_{lr}^1 \\ \\ \boxed{\begin{array}{l} t \\ \neg \text{Tuesd.}(t) \\ t_{su} < t < t_{tu} \end{array}} \\ \\ \boxed{\begin{array}{l} s'' \\ s_{lr}^0: \neg s'' \subseteq s_{lr}^0 \\ s'': \text{Att}(u, \{ \langle \text{BEL}, \mathbf{K} \rangle \}) \end{array}} \\ \\ \text{(K)} \\ \boxed{\begin{array}{l} t_{lv} \quad e_{lv} \quad s_{lv}^1 \quad l' \quad t'' \\ e_{lv} \subseteq t_{lv} \quad \text{day}(t'') \\ s_{lv}^0 \supset e_{lv} \supset s_{lv}^1 \quad t'' \supset t_{tu} \\ l' \neq l \quad t_{lv} < t_{tu} \quad t_{lv} \subseteq t'' \\ e_{lv}: \text{leave}(v, l) \\ s_{lv}^1: \text{IN}(v, l') \end{array}} \\ \\ \text{Att}(u, \langle \text{BEL}, \boxed{\begin{array}{l} s_{lv}^0: \text{IN}(v, l) \end{array}} \rangle) \end{array}$$

In (16) the accommodated belief that Mary remained in Paris till she left on Monday has been integrated into the final belief state ascribed to Bill. **K**, the content representation of the belief specified in the first Att-condition of (16), is the same DRS which also represents the content of the belief that appears in the second Att-condition. The new conditions involving t_{tu} state that the time represented by t_{tu} is the Tuesday immediately following the Sunday represented by t_{su} .

It should be clear how (16) could be extended to a fuller representation of the interpretations that result from the two persistence hypotheses. In either case a further Att-condition would have to be added which describes the relevant part of Bill's attitudinal state just before he got the new information on Tuesday. As nothing of interest is revealed by these full representations, I see no point in displaying them.

4.7 Are the Persistence Assumptions really involved?

Do the persistence hypotheses involved in the two interpretations just described play in the interpretation of (1) the role that I have been attributing to them? There appear to be good reasons

for doubting this. For one thing, isn't (16), which captures the common denominator of the two interpretations described in the last section, all that an interpretation of (1) should yield? And if (16) is all that interpretation of (1.ii) should yield, could this representation not be obtained by simpler means? For instance, couldn't we simply assume that interpretation of (1) proceeds like this:

- (17) (1.ii) says that as a result of what Bill learned on Tuesday he adopted the belief that Mary left Paris on the preceding Monday. This entails the belief that she was in Paris at the point on Monday when she left (since leaving Paris on Monday evidently entails being in Paris at the time on Monday when one leaves, and so you cannot believe the first without believing the second.)

The problem with (17) is that it fails to do justice to the role that the first sentence of (1) plays in the interpretation of the second. And that isn't quite right. It is essential to the interpretation of (1) that the prestate of the event of Mary's leaving Paris which the second sentence speaks of is the same state that is spoken of in the first sentence – or, more precisely, that Bill must end up believing that they are one and the same state. This is something that the strategy described in the last section does capture, but (17) does not.

Persistence of the state of Mary being in Paris does play a role in the interpretation of (1). But as we will see presently, it almost certainly doesn't play the role that it was assumed to play in the account described in Section 4.6.

Let us, as a first step towards this, begin by considering the variant of (1) that is given in (18)

- (18) On Sunday Bill heard that Mary was in Paris.

On Tuesday he learned that on the previous day she had visited the Swiss Embassy.

There is much that (18) and (1) have in common. The *that*-clause of the second sentence of (18) requires, just like the *that*-clause of the second sentence of (1), the context that is provided by the *that*-clause of the first sentence. And like in the case of (1) the context is not enough as it stands, but needs to be amplified by the assumption that the state of Mary's being in Paris persists till the relevant time on Monday. In the case of (18) this assumption is needed in particular to identify the intended referent of the NP *the Swiss Embassy*.²⁹

So much for the similarities between (18) and (1). But there is also an important difference. In (18) the content of the second *that*-clause is *consistent* with the amplified context. (In fact, it is not just consistent with it, but entails it.) Because of this the final belief which (18) attributes to Bill is that on Monday Mary was still in Paris and that, while still in Paris, she went to the Swiss Embassy there.³⁰

As a next step notice the similarity between the interpretation of (18) and that of the discourse (19), in which the contents of the *that*-clauses of (18) are presented as facts.³¹

²⁹There are countless Swiss embassies across the globe, so the interpreter needs further information in order to determine which one all those embassies this occurrence of *the Swiss Embassy* refers to. One kind of information that will serve this purpose is information about where the given embassy is located. And that is precisely the information that emerges when persistence is applied to the context condition that is provided by the *that*-clause of the first sentence: If Mary was still in Paris on Monday, then it was the Swiss Embassy in Paris which she visited.

³⁰So even if we were to assume that Bill's belief about Mary being currently in Paris persists till Tuesday, the new information described in (18.ii) would not require him to revise that belief.

³¹I have changed the *was* of the *that*-complement of the first sentence of (18a,b) – (*that*) *Mary was in Paris* – into *arrived*. The discourse *On Sunday Mary was in Paris. On Monday she visited the Louvre.* is strange because we

(19) On Sunday Mary arrived in Paris. On Monday she visited the Swiss Embassy.

In (19) too the first sentence provides the context for the interpretation of the second sentence, and once again, in order that this context can deliver what the interpretation of the second sentence demands of it, it has to be amplified by a persistence assumption. For sequences of factual sentences like (19) the role of persistence in interpretation is well-known and has been closely studied. These studies have shown that it is a general property of narrative texts (i.e. texts which describe episodes that develop in time) that states which have been introduced at one point in the text should be understood as continuing to hold until the discourse gives notice that this is no longer the case. More fully, if the state s has been introduced by sentence S_n of the discourse and has been asserted by S_n to hold at t , and the sentences S_{n+1}, \dots following S_n describe what is the case at times later than t , then s will be understood as still holding at these later times, unless and until it follows from one of these sentences that s has come to an end at or before one of those later times.³²

The interpretation of (19) involves a straightforward application of this persistence principle. The state of Mary being in Paris that is introduced by S_1 as holding on Sunday is assumed to persist at and beyond the later time described by S_2 unless it follows from S_2 that the state has terminated before or during that time. Since visiting the Swiss Embassy in Paris is compatible with this state, the prediction is that the state still holds at and continues to hold beyond the time of this visit on Monday.

The third point to note is that just as (19) is a close factual counterpart to (18), (20) is a close factual counterpart to (1)

(20) On Sunday Mary arrived in Paris. On Monday she left.

Here too the persistence principle makes the right prediction: The state of Mary being in Paris is assumed to persist until the time at which the discourse indicates its termination. In the present case this means persistence at least until the time of the event which the second sentence describes. This assumption makes it possible to interpret the second sentence as describing an event of Mary leaving Paris. But in this case the event description entails that the state ends there. So the interpretation of (20) which results is that the state of Mary being in Paris, which starts at some time on Sunday, continues until some time on Monday, when it is terminated by the event of her leaving Paris.

The parallels between the interpretations of (18) and (1) on the one hand and those of (19) and (20) on the other are striking – so striking that it is almost inescapable to see them as involving the same principles. In particular, interpretation of the belief attributing discourses seems to make use of the same persistence principle as interpretation of the factual discourses, the only difference being that in the case of (19) and (20) persistence is applied to a primary context while in (18) and (1) it is applied to a secondary context. What this suggests is a theory of interpretation

are inclined to perceive an opposition between *Sunday* in the first sentence and *Monday* in the second – as if Mary was no longer in Paris when she went to the Louvre. This is a rhetorical effect which has nothing to do with what we are after and which we do not get in (1) or (18), where *Sunday* modifies the matrix verb and not the verb of the *that*-clause. The first sentence of (19a,b) matches that of (1) and (18) in this respect, because it too carries no implications about the duration of Mary's stay in Paris.

³²In the discourse literature this persistence property of narrative discourse is also known as *monotonicity*. Analyses of monotonicity and formulations of the “monotonicity principles” which are responsible for this kind of persistence can be found for instance in Caenepeel & Sandstrom (1993) and Reyle & Roßdeutscher (2001).

in which the complement clauses of attitude attributions can be treated as if they combine to form bits of narrative discourse – more precisely: which allows sequences of such embedded clauses to be incrementally interpreted according to the same principles which existing theories of discourse semantics assume for bits of narrative discourse consisting of non-embedded sentences.

I will not attempt to formulate such a theory here, but only make a few informal observations about its implications for the interpretation of discourses like (18) and (1).³³ The first implication is this: When interpreting a sequence of attitude reports like those in (1) and (18), the interpreter builds, as he goes along, not only a primary discourse context, but also a secondary one. Both contexts grow as discourse interpretation progresses and are updated according to the same principles of language interpretation, and these principles include principles of persistence or monotonicity: Both types of context come with a default assumption that stative conditions which are initially said to hold at certain times persist until there is clear evidence to the contrary.³⁴

But of course, even if primary and secondary contexts come about in much the same way, their roles are very different. Primary contexts are what they have been made out to be ever since Stalnaker laid the foundations of dynamic interpretation: They capture the content of what has been interpreted already and provide the context for the utterances that follow. The role of a secondary context is different. What that role is depends, first and foremost, on the status of the sentence or clause that has been responsible for its latest update. For instance, when the clause is the *that*-clause of an attitude attribution, the updated context will serve to characterise the content of the attributed attitude; and if it is the *that*-clause of an attribution of belief change, as in the examples we have been looking at, then the updated context typically serves as specification of the content of the belief which results from the attributed change.

In this theory persistence has regained its place in the interpretation of discourses like (1) which (17) denies it. Superficially this might look like a revindication of the interpretation strategy described in Section 4.6, but even a cursory reflection suffices to see that it really is nothing of the kind. In the new theory persistence is made to play a very different part from the one it was made to play in 4.6. The persistence assumptions that are postulated by the present proposal are *not* assumptions about the persistence of the subject's beliefs. Their status is a more abstract one. They are just what we have described them to be: interpretation principles which are operative not only at the level of primary but also of secondary contexts. Precisely what that amounts to can be assessed only in the light of the role which the secondary context ultimately comes to play in the interpretation of the given discourse. But whatever that may tell us about the persistence principles that are used in the updating of various secondary contexts, it is clear that the applications of these principles to the interpretation of (1) or (18) have nothing directly to do with the beliefs which the subject Bill may have entertained during the time between his first belief change on Sunday and his second one on Tuesday. From this point of view the present theory is anything but a reaffirmation of the strategy of Section 4.6.

³³I believe that formalisation should not be all that difficult. At least this should be so at least for a restricted version, in which there is besides the primary context just one secondary context. In general the context structure can be much more complex, with different secondary contexts for different agents and/or different times. Even within a 'secondary' context which corresponds to a given agent's attitudinal state at a given time there may be a further hierarchical structure of 'subcontexts', e.g. reflecting the different modes of the attitudes of which it is composed. (Recall the remarks about complex attitudinal states in Section 2.1.

³⁴As far as this is concerned, the present approach seems to be much in the spirit of what Stalnaker must have had in mind when he introduced the concept 'secondary context' and coined this term for it.

For some readers it may have been plain from the start that the theory we have just outlined gives the better reconstruction of how people process and understand a discourse like (1). But for someone who approaches the problem presented by (1) from the perspective from which I – for better or worse – came to this problem, the point of view we have now reached is not self-evident. Let me, to conclude this section, briefly recapitulate this perspective and what it implies for the interpretation of a discourse like (1).

In a few words, the perspective was this: To correctly interpret the content of an attitude attribution, the interpreter often needs antecedent information about the attitudinal state of the attributee. When the given attribution is one of a sequence, then this information can usually be obtained, in part or in whole, from the attributions which precede it. Therefore the interpreter must, as he proceeds with his interpretation of the successive sentences of the sequence, build a representation of the attitudinal state of the attributee (or attributees) in question, and use, each time he tackles a new attribution, the representation of the attributee's attitudinal state that he has already obtained to make sense of what it is this new attribution contributes.

From this perspective it seems a natural assumption that in a case like (1), where the new sentence attributes an attitude at a later time than the state of which the interpreter has already obtained a representation, the interpreter should try to infer, on the basis of what he knows about the attributee at the earlier time, what his state should be like at the later time that is concerned in the new attribution. But natural though this extrapolation from the case of several attributions at a single time to different attributions at different times may seem at first sight, it doesn't seem – this is the central moral of this paper – to be what is actually going on. The persistence involved in the interpretation of such attribution sequences is *not* the persistence of any attitudes on the part of the attributee.

This is not to deny that persistence plays no part in the mental life of the attributees which such sequences describe. Surely, if they are at all like you and me, persistence will be as much a part of how their mental states evolve as it is of ours. But whatever assumptions we, as interpreters, may be inclined to make about the persistence of the attitudes that are attributed to them by the sentences of which we are the recipients, it appears that these assumptions are not part of the interpretation process as such. Or at any rate, if they are part of the process at all, they do not seem to be part of it in the manner suggested in Section 4.6.

5 Summary and Conclusions

This paper has focused on the analysis of one example – a brief and fairly simple discourse, consisting of two attitude attributions to a single subject. Like other examples which have been discussed in the literature the two attributions are understood as related to each other in content. The example differs from those other examples in that (i) its attributions are not attributions of attitudes as such, but of attitudinal changes; and (ii) these changes are asserted to have taken place at different times. Because the attitude change attributed by the second sentence occurs some time after the change described by the first sentence, a complication arises for the use that can be made of the first sentence as context for the interpretation of the second. This complication has been the paper's central focus.

In my attempt to deal with this problem I have outlined two interpretation strategies. The first of these is based on the idea that interpretation of a sentence which claims that a subject was in a certain attitudinal state at a certain time t , or undergoes a change of attitudinal state at t often requires antecedent information about the subject's state at or just before t ; so, when the

antecedent discourse only provides information about the subject's attitudes at other times, then it will be necessary to extrapolate from that information to the subject's attitudes at the time to which the new attribution pertains.

The main conclusion of this paper is that plausible as this idea may seem initially, it is almost certainly not correct for discourses like (1). Rather, in successions of attitude attributions the clauses which describe the contents of the attributed attitudes form a kind of subtext, whose dynamic interpretation is subject to the same rules that govern dynamic interpretation of discourse in which no attitude attributions are involved. The difference between this second perspective and the first one doesn't manifest itself when all attributions speak of the subject's attitudinal state at a single time. It becomes visible only when the attributions concern different times (or, but that is a theme for another paper, different agents). It is this which makes sentence sequences of the type of (1) a genuine challenge for the semantics of attitude attribution.

The account of the interpretation of (1) which emerged from this study as the more likely one is still awaiting formalisation. I don't expect that a formalisation would encounter serious obstacles, but we won't be sure until the work has actually been done. Since the work hasn't been done yet, the last part of the story this paper has told had to be kept at an informal level. And because of the informal character of this last part, the much more formal presentation of the earlier parts may in retrospect look like a luxury one could well have done without. But this impression is widely justified at best. For without the attention that we have paid to formal detail in these earlier sections, the issue with which we have grappled in the final section wouldn't even have come into proper view.

Bibliography

- Abusch, D., 1997. Sequence of tense and temporal de re. *Linguistics and Philosophy* **20** (1).
- Asher, N., Lascarides, A., 2003. *Logics of conversation*. Cambridge University Press.
- Beaver, D., 1995/2001. *Presupposition and assertion in dynamic semantics*. Doctoral dissertation, Edinburgh University (1995), CSLI (2001).
- Beaver, D., 2005. Accommodating topics. In: H. Kamp, B. Partee (eds.), *Context-dependence in the analysis of linguistic meaning*. Elsevier.
- Caenepeel, M., Sandstrom, G., 1993. A discourse level approach to the past perfect. In: M. Aurnague, A. Borillo, M. Bras (eds.), *Semantics of time, space and movement*. University of Toulouse.
- Dekker, P., 1993. Existential disclosure. *Linguistics and Philosophy* **16** (6).
- van Genabith, J., Kamp, H., Reyle, U., in press. Discourse representation theory. In: D. Gabbay (ed.), *Handbook of philosophical logic*, 2nd Edition. Kluwer.
- Geurts, B., 1995. *Presupposing*. Doctoral dissertation, University of Stuttgart.
- Geurts, B., 1999. *Presuppositions and pronouns*. Elsevier.
- Heim, I., 1983. On the projection problem for presuppositions. In: WCCFL 2 (Second West Coast Conference on Formal Linguistics).
- Kamp, H., 1990. Prolegomena to a structural theory of belief and other attitudes. In: C. A. Anderson, J. Owens (eds.), *Propositional Attitudes*. CSLI.
- Kamp, H., 2001a. The importance of presupposition. In: C. Rohrer, A. Roßdeutscher, H. Kamp (eds.), *Linguistic form and its computation*. CSLI.
- Kamp, H., 2001b. Presupposition computation and presupposition justification. In: M. Bras, L. Vieu (eds.), *Pragmatic and semantic issues in discourse and dialogue*. Elsevier.

- Kamp, H., 2003. Einstellungszustände und -zuschreibungen in der DRT. In: U. Haas-Spohn (ed.), *Intentionalität zwischen Subjektivität und Weltbezug*. Mentis.
- Kamp, H., Reyle, U., 1993. *From discourse to logic*. Kluwer.
- Kamp, H., Roßdeutscher, A., 1994. Remarks on lexical structure and DRS construction. *Theoretical Linguistics* **20** (2/3).
- Kamp, H., Schiehlen, M., 2002. Temporal location in natural languages. In: H. Kamp, U. Reyle (eds.), *How we say WHEN it happens*. Niemeyer.
- Kratzer, A., 1998. More structural analogies between tenses and pronouns. In: SALT VIII.
- Lascarides, A., Asher, N., 1993. Temporal interpretation, discourse relations and common sense entailment. *Linguistics and Philosophy* **16** (5).
- Ogihara, T., 1994. Temporal donkey sentences and sequence-of-tense phenomena. In: Alternate to SALT 4.
- Ogihara, T., 1996. *Tense, attitudes and scope*. Kluwer.
- Reichenbach, H., 1993. *Elements of symbolic logic*. The Free Press.
- Reyle, U., Roßdeutscher, A., 2001. Temporal underspecification in discourse. In: C. Rohrer, A. Roßdeutscher, H. Kamp (eds.), *Linguistic form and its computation*. CSLI.
- van der Sandt, R., 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* **9** (4), 333–378.
- Stalnaker, R., 1988. Belief attribution and context. In: R. Grimm, D. Merrill (eds.), *Contents of thought*. University of Arizona Press.
- von Stechow, A., 2002. Binding of verbs: Tense, person and mood under attitudes. In: NELS 33.
- Wilson, D., 1975. *Presupposition and non-truthconditional semantics*. Academic Press.