

Layered DRT

Bart Geurts & Emar Maier
{bart.geurts|e.maier}@phil.kun.nl

University of Nijmegen

November 7, 2003

1 Introduction

The information conveyed by any utterance is a motley ensemble. Utterances carry content about the world as it is according to the speaker, but also about speakers' attitudes, the way they speak, what has been said before, and so on. There are many kinds of information that are conveyed by way of language, and differences in kind correlate with differences in status. Presupposed information exhibits a distinctive projection behaviour; conversational implicatures are cancellable in a way that asserted information is not; a pronoun's gender may help to determine a referent, but is otherwise truth-conditionally inert; and so on.

Interpreting utterances is as much a matter of integrating these various kinds of information as of keeping them apart. This much is uncontroversial. As far as we are aware, however, no attempts have been made thus far to devise a fully general framework within which processes of information integration can be modeled. There are partial theories, to be sure. For example, there are quite a few well-developed analyses of the interaction between presupposed and non-presupposed content. But to the best of our knowledge the problem of information integration as such has not been addressed before. So that is what this paper is about: a general framework for representing and integrating all and sundry kinds of information that can be conveyed by linguistic means. This may seem like a grandiose project, and perhaps it is, but it is less ambitious than one might think. Our aim in this paper is to develop a framework for *representing* different kinds of linguistic and

para-linguistic information. How this information is processed is a different matter altogether, and not our main concern in the following.

2 Information integration

In order to explain what we mean by information integration, we will discuss a few concrete cases. It will be seen that our examples are quite diverse, but this is to be expected in view of the broad aim of this paper.

2.1 Presupposition

Our first example concerns the representation of presuppositions. In DRT, presuppositions are treated on a par with anaphoric expressions. Presuppositions prefer to link up to an antecedent, and if no suitable antecedent is available, they are interpreted by way of accommodation (van der Sandt 1992; Geurts 1999). Here is an example:

- (1) Perhaps White met the Chinese Empress today.

The initial representation of this sentence is as follows:

- (2) $[y: \text{White}(y), \diamond[x: \underline{\text{Chinese-Empress}}(x), \text{meet}(y, x)]]$

The underlining in this DRS reflects the fact that the definite NP ‘the Chinese Empress’ triggers the presupposition that there is a Chinese Empress. Assuming that this presupposition doesn’t have a suitable antecedent, it may be construed by way of accommodation (provided that the hearer is prepared to accept that there is a Chinese Empress), which means that the presupposition is added to the principal DRS, yielding the following representation:

- (3) $[x, y: \text{Chinese-Empress}(x), \text{White}(y), \diamond[: \text{meet}(y, x)]]$

This correctly captures what is expressed by (1), viz. that there is a Chinese Empress, and that White may have met her. Note that in the final representation of (1) the distinction between presupposed and asserted information is obliterated. For many purposes this is fine, for we mainly need that distinction in order to account for the fact that presupposed material is interpreted in its own special way. But as it turns out, the distinction between presupposed and asserted information remains active after a sentence

has been processed. To illustrate this, consider what would happen should another speaker object against (1) as follows:

(4) No, he had an encounter with the Japanese President.

Intuitively, this response only corrects what (1) asserts; the presupposition that there is a Chinese Empress remains unscathed. It is not a good idea, apparently, to discard the division between presupposed and non-presupposed material once the mechanism of presupposition projection has performed its duty.

The moral of the previous discussion is obvious and quite independent of the theory of presupposition we happen to prefer. It is simply that presuppositions will have to be separated from other types of information, because they have a special status: presuppositions are interpreted in their own way and once they have been interpreted they continue to be treated differently, as witness (4).

2.2 Implicatures

What has been just said about presuppositions holds good for implicatures, too. By way of example, consider what is generally regarded a ‘scalar implicature’:

(5) The porridge is warm.

This sentence presupposes that there is porridge, it asserts that the porridge is warm, and on the orthodox neo-Gricean view it implicates that the porridge is not hot; so the lexical meaning of ‘warm’ does not by itself rule out that the porridge is hot. According to this analysis, the information that the porridge is warm is of a different kind than the information that the porridge is not hot, and one of the stock-in-trade arguments in favour of the distinction is that the implicature is cancellable in a way the assertion is not:

- (6) a. The porridge is warm. As a matter of fact, it is hot.
b. ?The porridge is warm. As a matter of fact, it is cold.

(6a) shows that the implicature is cancellable, and the infelicity of (6b) suggests rather strongly that asserted information is more robust.

The upshot of these observations is analogous to that of the presuppositional case. We need to separate implicated information from other infor-

mation conveyed by an utterance, and it will not do to discriminate between presupposition, assertion, and implicature only for the duration of the process of sentence interpretation (as e.g. Gazdar (1979) does in his formal pragmatics); for the subsequent discourse may need these distinctions, too.

2.3 Secondary meaning

The heading of this rubric is somewhat tentative. What we have in mind is such phenomena as metaphor, metonymy, irony, and so on: non-truth-conditional content that is clearly part of the speaker's message, but may be at odds with its literal meaning. Even if such phenomena are to be treated in terms of conversational implicature, we prefer to distinguish them from run-of-the-mill cases of implicature, which merely add to the literal meaning of an utterance. Especially stark cases in point are irony and sarcasm. Suppose a connoisseur of modern art volunteers (7), demonstrating what is obviously a fumbling attempt at self-expression:

(7) That is a beautiful painting.

Under the circumstances, this statement is probably intended to convey the opposite from what it literally says. It is, mildly put, something of a mystery how such 'double meanings' are computed, and we don't have anything new to offer in this regard. However, we do have a proposal as to how different levels of meanings can be represented in such a way that some bits of information are shared (e.g. the reference of the pronoun is shared between the literal and the non-literal meaning of the sentence), while others are segregated.

2.4 The former and the latter

Consider the following example:

(8) If a beggar meets a bishop, then the latter will bless the former.

On the face of it, anaphoric devices like 'the former' or 'the latter' do not seem particularly troublesome. On reflection, however, they add an interesting wrinkle to the problems posed by definite NPs. As the wrinkle will appear, in some way or other, no matter what our theoretical predilections concerning definites are, we will here follow our own. According to the DRT treatment of definite NPs that we favour, definites are presuppositional expressions,

which is to say that prefer to link up to a contextually given antecedent. In this respect, everything is fine in the present example, since ‘the latter’ as well as ‘the former’ have suitable antecedents: the former refers back to ‘a bishop’, the latter to ‘a beggar’. However, problems begin to emerge when we ask ourselves how exactly these expressions manage to link up to their antecedents. To see this, consider how the story about ‘the prelate’ in (9) would go:

(9) If a beggar meets a bishop, then the prelate will bless him.

In this sentence, ‘the prelate’ establishes an anaphoric link in much the same way as ‘the latter’ does in (8). In this respect there isn’t much of a difference between the two expressions. However, the descriptive content of ‘the prelate’ is very different from that of ‘the latter’—so different, in fact, that some people would say the information contained in ‘the prelate’ is part of the sentence’s truth-conditional content, whereas nobody would want to claim that the descriptive content of ‘the latter’ has an effect on the truth conditions of (8).

So the wrinkle is this. Since they appear to be just a special case of presuppositional (or anaphoric) expressions, we would like to analyse ‘the former’ and ‘the latter’ as being on a par with any other definite NP, except of course that they constrain the process of interpretation by referring not to *what* has been said but *how* it was said. Qua presuppositional expression, ‘the latter’ presents its referent as given in the same way ‘the prelate’ does; it is just that we have two rather different modes of givenness, so to speak. The problem is, therefore, how we can distinguish two very different kinds of information—about the discourse and about the world—and have a uniform account of definites at the same time.

Note, finally, that in some languages ordinary pronouns may be used in the same way as English ‘the former/latter’:

(10) Am Ende besteht ein wesenhafter Unterschied zwischen [dem Erfassen des Ganzen des Seienden an sich]_i und [dem Sichbefinden inmitten des Seienden im Ganzen]_j. Jenes_i ist grundsätzlich unmöglich. Dieses_j geschieht ständig in unserem Dasein.

[Martin Heidegger, Was ist Metaphysik?]

We will not attempt to render this passage in colloquial English. Suffice it to say that it illustrates how in German the distal and proximal demonstrative

pronouns ‘jenes’ and ‘dieses’ are used precisely as ‘the former’ and ‘the latter’ would be used in English.

2.5 Grammatical gender

Up to a point grammatical-gender pronouns function not unlike ‘the latter’ and ‘the former’, as witness the following example from German:

- (11) Braun hat {einen Wagen/ein Auto} gekauft. {Er/Es} ist grün.
Braun has bought {a car_{neut}/a car_{masc}}. {Pro_{neut}/Pro_{masc}} is green.

On pains of unintelligibility, the pronoun in the second sentence has to agree in gender with its antecedent in the first, and although the term ‘grammatical gender’ may suggest otherwise, this is not a grammatical phenomenon. Neither pronoun in (11) is bound syntactically: they are perfectly ordinary referential anaphors, whose duty it is to retrieve a discourse referent from the context. What makes these pronouns special is the requirement that, to a first approximation at least, the last mention of their referents should have employed an expression of the same gender.

That this is not quite right yet appears from the fact that gender pronouns may be used deictically, that is, without a linguistic antecedent. For instance, a Frenchman watching someone trying to get a table (la table_{fem}) into his car might volunteer the following remark (Tasmowski-De Ryck and Verluyten 1982):

- (12) Tu n’arriverais jamais {la/*le} faire entrer dans la voiture.
You’ll never manage to get {pro_{fem}/*pro_{masc}} into the car.

Here the pronoun has to agree in gender not with an earlier expression, for there was no previous mention of the table, but rather with the noun would have been used by default to refer to the table. Observations like this highlight the fact that grammatical gender, too, depends on the non-linguistic context for its interpretation, although part of the information they carry is of a linguistic nature, and must therefore be represented on a different level, presumably.

2.6 Direct reference

Since Kripke (1972) and Kaplan (1989) it has become widely accepted that certain types of singular terms, especially proper names and indexicals, ‘refer directly’. This is not to deny that these terms have descriptive content. For it is obvious that, for example, ‘I’ carries the information that the speaker is referring to himself qua speaker, and that the name ‘Brown’ refers to someone who is called ‘Brown’. It is just that, intuitively speaking, this sort of content is not truth-conditional content.

Direct reference poses a problem for theories of meaning that treat all descriptive content alike. Kripke and Kaplan have argued against such mono-semantic accounts, observing that a sentence like (13a) does not have the same truth-conditions as (13b); the first is a contingent truth, while (13b) is necessarily—even logically—true.

- (13) a. Brown is called ‘Brown’.
b. Brown is Brown.

The same point can be made with indexicals:

- (14) a. I am the speaker.
b. The speaker is the speaker.

Whereas the proposition expressed by (14a) might have been false, (14b) is, on one of its readings at least, necessarily true.

Kaplan’s well-known analysis of direct reference involves dividing the Fregean notion of sense into two components, which he calls ‘character’ and ‘content’. The character of expression is its linguistic meaning, which in a given context determines the expression’s truth-conditional content. The descriptive content of indexical is unlike that of a definite description in that it remains at the level of linguistic meaning, and doesn’t enter the truth-conditional level. We will see later on how this distinction can be captured in our representational framework.

3 Binding problems

We have discussed a handful of phenomena that illustrate the mundane truth that different kinds of linguistic and para-linguistic information need to be kept apart. But although we must separate between different kinds of con-

tent, the separation had better be not too strict. This is the lesson taught by a problem first noted by Karttunen and Peters (1979), which has come to be known as the ‘binding problem’ of presupposition projection (we will shortly see, however, that the problem is quite general). Karttunen and Peters’ example is the following:

(15) ?Someone managed to succeed George V on the throne of England.

This sentence has a question mark because it is pragmatically infelicitous: it suggests that the person who succeeded George V found it difficult to do so, which can hardly be the case (at least not in the sense intended here; George V’s successor may have had problems adjusting to his new station, but he obtained it without effort). Apparently, the presupposition triggered by the verb ‘manage’ fails in this case. The problem is that many theories of presupposition (including Karttunen and Peters’ own) cannot account for this kind of infelicity, because they strictly separate between asserted and presupposed information, as a consequence of which the content of (15) is predicted to have the following components:

(16) Assertion: Someone succeeded George V on the throne of England.
Presupposition: It was difficult for someone to succeed George V on the throne of England.

Unfortunately, the presupposition triggered by ‘manage’, thus construed, comes out true: practically everybody would have had a hard time succeeding George V. This is a problem not only for Karttunen and Peters’ own treatment of presupposition, but for many theories of a younger vintage, too. The problem arises, obviously, because presupposed and asserted content are separated too strictly, and it is the opposite from the problem discussed above: the DRT treatment of presupposition does not run into the binding problem because it keeps presuppositions and assertions together, though for other reasons they should be differentiated more than they currently are, as we have seen.

Although there has been some discussion of the binding problem in the literature (Krahmer 1998; Beaver 2001), it hasn’t been noted before, as far as we are aware, that the problem is not confined to presuppositions. In fact, binding problems are liable to crop up whenever a strict separation is made between different kinds of information with interdependencies between them. For example, what has been discussed by von Stechow (1994), Rooth

(1995), and Krifka (2001) as a problem of ‘requantification’ may be seen as just another binding problem, which arises when the interpretative effects of focusing are treated separately from ordinary semantic content.

(17) A dog is usually intelligent_F.

The most obvious interpretation of (17) is that most dogs are intelligent. The requantification problem arises for such theories that treat focused information separately from other content expressed by the sentence. How exactly the problem manifests itself is different from one theory to the next, but the following treatment is representative. Suppose that the backgrounded information in (17) simply adds a restriction to the domain of ‘usually’ along the following lines:

(18) Usually (if a dog is either intelligent or moderately clever or dumb or ...)(then a dog is intelligent)

The problem with this is that the indefinite ‘a dog’ is used twice, which, together with the widely accepted idea that indefinites introduce new referents, entails that the dogs introduced in the restrictor and the nuclear scope of ‘usually’ cannot be identical. And this is clearly wrong, because what (17) should mean is something like the following:

(19) Usually (if a dog_{*i*} is either intelligent or moderately clever or dumb or ...)(then it_{*i*} is intelligent)

As in the case of presuppositions, the problem seems to be that two kinds of information that clearly have to be separated come too far apart, and again the problem doesn’t arise for those DRT analyses of focusing that represent backgrounded information on the same level as assertions and presuppositions (Geurts and van der Sandt 2003).

Another variety of binding problem arises in connection with implicatures. Consider the following example:

(20) Green thought that some of the madrigals were nice.

This may be heard as implicating that none of the madrigals in question made a profound impression on Green. Let us assume, if only for expository purposes, that this is a scalar implicature; that is, the use of the scalar adjective ‘nice’ implicates but does not entail ‘not very good’. How can we

represent this implicature in the case of (20)? One solution that comes to mind is the following:

(21) Green thought that some of the madrigals were not very good.

This is correct as far as it goes, but it fails to capture what we are after, because there is no guarantee that (21) is about the same madrigals as (20). Another attempt then:

(22) According to Green the madrigals she thought were nice were not very good.

Suppose, however, that Green found one of the madrigals really superb. Then (20) might still be true, but (22) would be false (assuming, of course, that ‘superb’ entails ‘nice’). Intuitively speaking, what we would like to have is something like this:

(23) Green thought (assertion) that [some of the madrigals]_i were nice, and she thought (implicature) that they_i were not very good.

And we don’t need attitude contexts to bring out the problem. Suppose that after the concert Green reported as follows:

(24) Half of the madrigals were nice.

Suppose furthermore that this is taken to implicate that the madrigals in question were not very good. Then, again, the problem is how this implicature can be captured. For the same reasons as discussed the foregoing it will not do to say that half of the madrigals were not very good, nor is it right to say that, according to Green, the madrigals that were nice were not very good.

Binding problems of an especially hairy sort enter the picture when we take non-literal meanings into account. Consider, for example, sentence (25) as said by a father to his 15-year-old son:

(25) Someone used my after shave this morning.

This may be understood as implying that the addressee is the culprit, and if it is, the referents of ‘someone’ and ‘you’ will have to coincide at some level of representation.

The examples discussed in the foregoing all present potential binding problems in the sense explained at the beginning of this section: in each case there are good reasons for wanting to partition the information expressed by an utterance, but the separation shouldn't be so strict that all connections are severed.

In this section and the last one we have discussed various phenomena illustrating that different kinds of information need to be kept apart but in such a way that certain interdependencies between them are captured. In the remainder of this paper we present a unified account that attempts to accomplish this.

4 Layered DRT

The basic idea underlying Layered DRT (or LDRT for short) is straightforward enough. It is that a discourse representation should admit of more than one layer of information. All the information that is exchanged between speakers will go into the same representation, but within this representation we want to distinguish between information that is asserted, presupposed, implicated, and so on. So within an LDRS there will be layers for assertions, presuppositions, implicatures, grammatical features of utterances, formal properties of the discourse, and so on. In many cases, information will reside on a single layer, but occasionally the same information can be on more than one layer. This holds, in particular, for discourse referents, which may be seen as inter-layer communication hatches.

Formally, layers are implemented as sets of labels on discourse referents and conditions. Every layer has its own label, and as the same piece of content may be on several layers at once, discourse referents and conditions will be assigned sets of labels. In the following, we will show how to add layers to the standard DRT language.

4.1 The LDRT language

The vocabulary of the LDRT language extends the standard DRT language with a set of layer labels. So we start, as usual, with sets of discourse referents, predicates, and logical constants. All conditions in a layered DRS (LDRS) will bear zero or more labels; discourse referents will be labeled, too, but only when they are introduced, not when they occur as arguments.

Taking as given inventories of discourse referents, predicates, logical constants, and layer labels, the following clauses simultaneously define the set of LDRSs, labeled discourse referents and labeled conditions:

- (26) Let L be a set of layer labels. Then:
- a. An LDRS φ is a pair $\langle U(\varphi), Con(\varphi) \rangle$, where $U(\varphi)$ is a set of labeled discourse referents and $Con(\varphi)$ is a set of labeled DRS-conditions.
 - b. If u is a discourse referent, then u_L is a labeled discourse referent.
 - c. If P is an n -place predicate and u_1, \dots, u_n are discourse referents, then $P_L(u_1, \dots, u_n)$ is a labeled condition.
 - d. If u and v are discourse referents, then $u =_L v$ is a labeled condition.
 - e. If φ and ψ are LDRSs, then $\neg_L \varphi$, $\varphi \vee_L \psi$, and $\varphi \Rightarrow_L \psi$ are labeled conditions.

Officially, LDRSs are set-theoretic constructs, but inofficially we will employ the following notation, which we find easier to read. First, instead of $\langle \{u_1, \dots, u_m\}, \{\varphi_1, \dots, \varphi_n\} \rangle$ we will write $[u_1 \dots u_m : \varphi_1 \dots \varphi_n]$. Secondly, if $\{\alpha_1, \dots, \alpha_n\}$ is a set of layer labels we will write $\alpha_1 \dots \alpha_n$; hence x_{abc} , $P_{abc}(x_1, \dots, x_n)$, etc. Thirdly, if a discourse referent or condition resides on all layers, we will omit the label set; for example, if we have only three labels, x_{abc} reduces to x .

To illustrate how the LDRS language may be used, the interpretation of example (5), repeated here as (27a), can be rendered in LDRT as (27b):

- (27) a. The porridge is warm.
 b. $[x_p: \text{porridge}_p(x) \text{warm}_a(x) \neg_i[: \text{hot}_i(x)]]$

The LDRS in (27b) has three layers: a, p, and i, which contain asserted, presupposed, and implicated material, respectively. The intuitive interpretation of (27b) is that it is presupposed that there is some x that is porridge, that x is asserted to be warm, and that x is implicated not to be hot. Note that the asserted and implicated parts use a discourse referent that is presupposed: we simply cannot say what (27a) asserts or implicates without referring to whatever it is that is being presupposed by the subject term. This is to say that structures like (27b) cannot be interpreted as consisting of three independent layers of information; rather, discourse referents should be allowed to liaise between layers.

4.2 Semantics of LDRT

The idea underlying our semantics for the LDRS language is simply that, instead of specifying what is the truth-conditional content of an LDRS φ , we have to define what is the truth-conditional content of a selection L of layers in φ . That is to say, if φ contains a condition ψ_K , where K is the layer set associated with ψ , ψ_K is to be ignored unless K and L overlap; i.e. unless $K \cap L \neq \emptyset$. In the standard semantics of DRT, an embedding function f is said to be extended by another function g , with respect to a given DRS φ , iff $f \subseteq g$ and $dom(g) = dom(f) \cup U(\varphi)$. Here we impose an additional restriction on the discourse referents that may be added to the $dom(f)$ in the process of building g : they must be adorned with a label set K that overlaps with a given set L :

$$(28) \quad f[\varphi]_L g := \\ f \subseteq g \text{ and } dom(g) = dom(f) \cup \{u \mid \exists K: K \cap L \neq \emptyset \text{ and } u_K \in U(\varphi)\}$$

Given a world w , a label set L , and an embedding function f , the principal semantic object associated with an LDRS φ is $\|\varphi\|_{L,w}^f$, which, if defined, is the set of embedding functions g that extend f and make the L -part of φ true at world w . If $\|\varphi\|_{L,w}^f$ is a non-empty set, the L -part of φ is true at w ; if $\|\varphi\|_{L,w}^f = \emptyset$, the L -part of φ is false at w ; and otherwise $\|\varphi\|_{L,w}^f$ is undefined. If φ_K is an LDRS-condition, $\|\varphi\|_{L,w}^f$, if defined, is either \top or \perp .

(29) Let $M = \langle D, W \rangle$ be a model, where D is a domain of individuals and W is a set of interpretation functions ('worlds'); $w \in W$; f is a partial function from the set of discourse referents into D ; and L is a set of layer labels:

LDRSs: definedness and interpretation

Let φ be an LDRS. Then:

- a. $\|\varphi\|_{L,w}^f$ is defined iff $\exists g: f[\varphi]_L g$ and $\forall \psi \in \text{Con}(\varphi): \|\varphi\|_{L,w}^g$ is defined.
- b. If defined, $\|\varphi\|_{L,w}^f = \{g \mid f[\varphi]_L g \text{ and } \forall \psi \in \text{Con}(\varphi): \|\psi\|_{L,w}^g = \top\}$.

LDRS-conditions: definedness

- c. $\|P_K(u_1, \dots, u_n)\|_{L,w}^f$ is defined iff $\{u_1, \dots, u_n\} \subseteq dom(f)$.
- d. $\|u =_K v\|_{L,w}^f$ is defined iff $\{u, v\} \subseteq dom(f)$.
- e. $\|\neg_K \varphi\|_{L,w}^f$ is defined iff $\|\varphi\|_{L,w}^f$ is defined.
- f. $\|\varphi \vee_K \psi\|_{L,w}^f$ is defined iff $\|\varphi\|_{L,w}^f$ and $\|\psi\|_{L,w}^f$ are defined.

- g. $\|\varphi \Rightarrow_K \psi\|_{L,w}^f$ is defined iff $\|\varphi\|_{L,w}^f$ and $\|\varphi \oplus \psi\|_{L,w}^f$ are defined, where $\varphi \oplus \psi = \langle U(\varphi) \cup U(\psi), Con(\varphi) \cup Con(\psi) \rangle$ (LDRS-merge).

LDRS-conditions: interpretation

If φ_K is an LDRS-condition, $\|\varphi_K\|_{L,w}^f = \top$ iff $\|\varphi_K\|_{L,w}^f$ is defined and either $K \cap L = \emptyset$ or one of the following holds:

- h. φ_K is of the form $P_K(u_1, \dots, u_n)$ and $\langle f(u_1), \dots, f(u_n) \rangle \in w(P)$.
 - i. φ_K is of the form $u =_K v$ and $f(u) = f(v)$.
 - j. φ_K is of the form $\neg_K \psi$ and $\|\psi\|_{L,w}^f = \emptyset$.
 - k. φ_K is of the form $\psi \vee_K \chi$ and $\|\psi\|_{L,w}^f \cup \|\chi\|_{L,w}^f \neq \emptyset$.
 - l. φ_K is of the form $\psi \Rightarrow_K \chi$ and $\forall g \in \|\psi\|_{L,w}^f: \|\chi\|_{L,w}^g \neq \emptyset$.
- $\|\varphi_K\|_{L,w}^f = \perp$ iff $\|\varphi_K\|_{L,w}^f$ is defined and $\|\varphi_K\|_{L,w}^f \neq \top$.

If φ is a DRS, $\|\varphi\|_{L,w}^f$ is the set of embedding functions that extend f and that make the L -part of φ come out true in w . Shifting perspective somewhat, the following clauses define the set of worlds in which φ 's L -layers are true:

- (30) $\|\varphi\|_L^f = \{w \mid \|\varphi\|_{L,w}^f \neq \emptyset\}$, if $\exists w: \|\varphi\|_{L,w}^f$ is defined; undefined otherwise.
- (31) $\|\varphi\|_L = \|\varphi\|_L^{f_0}$, where f_0 is the empty function.

In LDRT, every choice of labels engenders its own sort of information. For example, $\|(27b)\|_{\{p\}}$ is the set of worlds that contain porridge; $\|(27b)\|_{\{p,a\}}$ is the set of worlds with warm porridge; and $\|(27b)\|_{\{p,a,i\}}$ is the set of worlds containing porridge that is warm but not hot. Note that, for example, $\|(27b)\|_{\{a\}}$ and $\|(27b)\|_{\{i\}}$ are undefined. This is because the a - and i -layers of (27b) use a discourse referent that is introduced in the p -layer. The undefinedness of $\|(27b)\|_{\{a\}}$ and $\|(27b)\|_{\{i\}}$, as opposed to the definedness of $\|(27b)\|_{\{p\}}$, is due to the fact that the assertion and implicature of (27b) depend on what the sentence presupposes, but not vice versa. What (27b) asserts or implicates can only be specified relative to a given value (or range of possible values) of the discourse referent x .

If we collect all labels into one set L , then $\|\cdot\|_L$ captures everything that is somehow expressed by a sentence or discourse. We have seen that such content aggregates need not be consistent:

- (32) a. That is a beautiful painting. (= (7))
 b. $[x_k: dem_k(x) beautiful-painting_a(x) \neg_i[: beautiful-painting_i(x)]]$

On the intended reading of (32a), the sentence is meant to convey that the demonstrated object is in fact not a beautiful painting at all. Assuming for the nonce that this is an implicature, and that demonstratives are represented on a special k -layer (more about which in §5), we obtain a representation along the lines of (32b). The literal meaning of (32a) is $\|(32b)\|_{\{k,a\}}$; the implicit message is $\|(32b)\|_{\{k,i\}}$; and $\|(32b)\|_{\{k,a,i\}} = \emptyset$.

It will be clear how LDRT solves the binding problems of §3. Here is an example:

- (33) a. Some of the madrigals are nice. (cf. (24))
 b. $[: [x: \text{madrigal}_p(x)] \langle \text{some}_a x \rangle [: \text{nice}_a(x) \text{ not-very-nice}_i(x)]$
 $[x: \text{madrigal}_p(x)] \langle \text{not-all}_i x \rangle [: \text{nice}_i(x)]]$

Normally speaking, an utterance of (33a) would be construed as implicating, first, that not all of the madrigals are nice, and secondly, that the madrigals that are claimed to be nice are just that, and not exceptionally nice; it is the second implicature that is liable to cause a binding problem. Adopting Kamp and Reyle's (1993) duplex conditions, the problem doesn't arise if we represent (33a) as (33b). (There are other ways of handling the implicature induced by 'some' in (33a). For example, we might consider introducing conjoined quantifiers, which would give us quantifiers like $\langle \text{some}_a x, \text{not-all}_i x \rangle$. The details of such an analysis remain to be worked out.)

5 Indexical content: the k -layer

The LDRS-semantics given in (29) is uniform in the sense that, once a group of layers have been selected, all layers are treated alike. (That is, they are treated alike by the semantics. If layers were alike in every respect, there would be no point in having them in the first place.) In the remainder of this paper we discuss two classes of phenomena which show that this is not quite right, and adjust our semantics accordingly.

Direct reference (cf. §2.6) is a problem for our framework as it stands, because in the case of indexicals, demonstratives, names, etc., the connection between an expression and its referent appears to be rigid. It is a matter of debate whether this rigidity is semantic or pragmatic in nature, but for our present purposes that is as it may be. For now our concern is with the question how this type of context dependence may be represented.

Since the contextual information implicated in direct reference has a spe-

cial status, we will put it on a layer of its own, which we label ‘k’. We already used this layer in the last section (cf. (32b)); another example is the following:

- (34) a. I am the speaker. (= (14a))
 b. $[x_k: \text{speaker}_k(x) \text{ speaker}_a(x)]$

However, it is not enough to just put indexical content in a layer of its own: $\|(34b)\|_{\{k,a\}}$ merely says that there is a speaker, whereas it should say of the individual who *in fact* is doing the talking that he is the speaker. In order to account for this, we follow Kaplan by making the content of an LDRS dependent on the context in which it occurs. A context may be regarded as a small world in the sense that it determines a unique speaker, hearer, etc. In our LDRS-semantics worlds are identified with interpretation functions, so if \mathbf{c} is a context, then $\mathbf{c}(\text{speaker})$, $\mathbf{c}(\text{hearer})$, $\mathbf{c}(\text{now})$, etc., are singleton sets. There are various ways of enforcing this restriction, two of which will be demonstrated in the following.

Relative to a given context \mathbf{c} , we define the indexical content of the L -part of an LDRS φ as follows:

$$(35) \mathbb{I}_{L,\mathbf{c}}(\varphi) = \|\varphi\|_L^\iota, \text{ if } \|\varphi\|_{\{k\},\mathbf{c}}^{f_0} = \{\iota\}; \text{ otherwise undefined.}$$

If $\|\varphi\|_{\{k\},\mathbf{c}}^{f_0}$ is not a singleton set, \mathbf{c} fails to determine unique values for all discourse referents in the k -layer, and $\mathbb{I}_{L,\mathbf{c}}(\varphi)$ is undefined. Otherwise $\mathbb{I}_{L,\mathbf{c}}(\varphi) = \|\varphi\|_L^\iota$, where ι is the unique embedding function determined by \mathbf{c} . For example, if White is the speaker in context \mathbf{c}_{42} , $\mathbb{I}_{\{k,a\},\mathbf{c}_{42}}(34b) = \|(34b)\|_{\{k,a\}}^{\langle x, \mathbf{White} \rangle} =$ the set of worlds in which White is the speaker.

Another way of implementing a Kaplan-style notion of content in LDRT is by building contexts into the definition of $\|\cdot\|$. We have to pass several choice points if we take this line, but the simplest solution we can think of mainly affects the interpretation of atomic conditions, which now comes out as follows:

- (36) a. $\|\mathbf{P}(u_1, \dots, u_n)_K\|_{L,\mathbf{w}}^{f,\mathbf{c}}$ is defined iff $\{u_1, \dots, u_n\} \subseteq \text{dom}(f)$ and either $k \notin K$ or $|\mathbf{c}(\mathbf{P})| = 1$.
 b. $\|\mathbf{P}(u_1, \dots, u_n)_K\|_{L,\mathbf{w}}^{f,\mathbf{c}} = \top$ iff $\|\mathbf{P}(u_1, \dots, u_n)_K\|_{L,\mathbf{w}}^{f,\mathbf{c}}$ is defined and either $K \cap L = \emptyset$ or one of the following holds:
 – $k \notin K$ and $\langle f(u_1), \dots, f(u_n) \rangle \in \mathbf{w}(\mathbf{P})$;
 – $k \in K$ and $\langle f(u_1), \dots, f(u_n) \rangle \in \mathbf{c}(\mathbf{P})$.

If an atomic condition is not on the k-layer, its interpretation is as in (29). If it is on the k-layer, its semantic value is undefined if the current context c fails to assign a unique value to its predicate. Suppose, for example, that $c_{42}(\text{speaker}) = \text{White}$; then $\|(34b)\|_{\{k,a\},w}^{f_0,c_{42}} = \emptyset$ if White is not the speaker at w , or else $\|(34b)\|_{\{k,a\},w}^{f_0,c_{42}} = \{g\}$, where $\text{dom}(g) = \{x\}$ and $g(x) = \text{White}$. Hence, $\|(34b)\|_{\{k,a\}}^{c_{42}} = \mathbb{I}_{\{k,a\},c_{42}}(34b) =$ the set of worlds in which White is the speaker.

So the two methods produce the same result in this case, as they do in many others, but they are not fully equivalent. First, and most importantly, with the second method, all discourse referents and conditions labeled ‘k’ are interpreted at the contextual index. The first method by contrast presupposes, in effect, that all k-material resides in the principal LDRS, or otherwise it will not be interpreted at c . Secondly, whereas the second method requires that each k-marked predicate be unique, the first method is less stringent in this respect, since it requires merely that, between them, the conditions in the k-layer determine unique values for all k-marked discourse referents. At the end of the day, the choice between these methods depends on considerations that go beyond the scope of this paper. We should like to note, however, that on the whole the second method is more versatile, and brings out more clearly the relation between indexical content and what we call ‘formal content’, which is the topic of the next section.

6 Formal content

We explained in §2.4 that we should like to have an analysis of expressions like ‘the former’ and ‘the latter’ which treats them as regular definite descriptions whose content is somewhat special. In this section we shall see that LDRT can provide us with such an analysis, though it doesn’t come for free; for it turns out that ‘the former’, ‘the latter’, and their kin are context-dependent in a way that resembles the context-dependence of indexical expressions.

The LDRS in (37b) is a first stab at capturing the intuition that the expressions ‘the former’ and ‘the latter’ as used in (37a) are presuppositional devices whose content refers to formal properties of the previous discourse:

- (37) a. As the beggar approached the bishop, the latter blessed the former.
 b. $[x, y: \text{beggar}_p(x) \text{ bishop}_p(y) \text{ approach}_a(x,y) \ x \prec_{\text{pf}} y \ \text{bless}_a(y,x)]$

The intended interpretation of condition ‘ $x \prec_{pf} y$ ’ is something like: ‘The last mention of x precedes the last mention of y .’ (There may be other, and perhaps better, ways of rendering the meanings of ‘the former’ and ‘the latter’, but this one will do.) This condition is part of the presupposition triggered by ‘the former/latter’, so it is on the p-layer, and it is also on the f-layer, because it refers to the form of the preceding discourse.

There are two problems with this proposal: a technical one, which isn’t hard to solve, and a deeper one (which, incidentally, is quite general and not an artifact of our framework), for which we shall present a tentative solution. The technical problem is that, as it stands, our LDRS-semantics doesn’t allow us to isolate the content of ‘the former/latter’ from the remaining content of the sentence. In order to compute $\|(37b)\|_{\{p,a\}}$, we have to take into account all discourse referents and conditions whose label sets contain ‘p’ or ‘a’ or both; and ‘ $x \prec_{pf} y$ ’ meets this requirement. Consequently, $\|(37b)\|_{\{p,a\}} = \|(37b)\|_{\{p,a,f\}}$, and there is as yet no way of factoring out in semantic terms the contribution of ‘ $x \prec_{pf} y$ ’—no way, that is, of separating form-related content from the real stuff, so to speak, of presuppositional and asserted content.

The problem is caused by the way our semantics uses label sets to sort the content of an LDRS: $\|\cdot\|_L$ evaluates all discourse referents and conditions whose label sets overlap with L . It turns out that this procedure, though adequate in general, is not always sufficiently selective. What we need, apparently, is a more flexible apparatus for sorting the content of an LDRS. For example, occasionally $\|\cdot\|_L$ should evaluate only those discourse referents and conditions whose label set K is such that $K \subseteq L$. It will be clear, we trust, that it is not too difficult to emend the LDRS-semantics so as to make layer selection more flexible, and we will assume without proof that this can be done.

As it stands, our treatment of the interpretation of ‘the former/latter’ and related devices, such as grammatical gender, presupposes that the properties denoted by these expressions are properties of regular individuals: ‘ $x \prec_{pf} y$ ’ is true iff the last mention of the individual associated with x preceded the last mention of the individual associated with y . Example (8), repeated here as (38a), demonstrates that this is not correct in general:

- (38) a. If a beggar meets a bishop, then the latter will bless the former.
 b. $[\vdash [x, y: \text{beggar}_a(x) \text{ bishop}(y)_a \text{ meet}_a(x,y)] \Rightarrow [\vdash x \prec_{pf} y \text{ bless}_a(y,x)]]$

In order to heighten the dramatic impact of the example, we might suppose that the conditional is given a modal interpretation, and is construed as quantifying over worlds. But whatever the sentence quantifies over, its domain is not confined to states of affairs containing pairs of persons one of whom was mentioned before the other; for the mentioning was done in the actual context.

If the predicate ‘ \prec ’, as used in (38), is not about beggars and bishops, be they possible or real, then what *is* it about? The answer, we would like to suggest, is ‘discourse referents’: the condition ‘ $x \prec y$ ’ states that the most recent use of x preceded the most recent use of y . Hence, a speaker who employs the ‘former/latter’ idiom in effect instructs the hearer to retrieve from the context a pair of recently-used discourse referents.

In order to implement this idea, we propose to expand the interpretation of atomic conditions given in (36) along the following lines:

- (39) a. $\|P_K(u_1, \dots, u_n)\|_{L,w}^{f,c}$ is defined iff the following holds:
- If $f \notin K$ then $\{u_1, \dots, u_n\} \subseteq \text{dom}(f)$.
 - If $\{f, k\} \cap K \neq \emptyset$ then $|\mathbf{c}(P)| = 1$.
- b. $\|P_K(u_1, \dots, u_n)\|_{L,w}^{f,c} = \top$ iff $\|P_K(u_1, \dots, u_n)\|_{L,w}^{f,c}$ is defined and either $K \cap L = \emptyset$ or one of the following holds:
- If $k \in K$ then $\langle f(u_1), \dots, f(u_n) \rangle \in \mathbf{c}(P)$;
 - If $f \in K$ then $\langle u_1, \dots, u_n \rangle \in \mathbf{c}(P)$.
 - If $\{f, k\} \cap K = \emptyset$ then $\langle f(u_1), \dots, f(u_n) \rangle \in \mathbf{w}(P)$;

According to this analysis, the interpretation of f -marked conditions depends on the context, and in this respect f -conditions and k -conditions are alike. But on the other hand, f -conditions are special in that they are about discourse referents, whilst all other conditions, k -conditions included, are about ‘real’ things in the world.

7 Conclusion

Our main objective in this paper was to provide a general framework for representing and integrating all sorts of information that may be conveyed by linguistic means. Our proposal is LDRT. From a syntactical point of view, LDRT is perfectly straightforward. Label sets allow us to separate between different types of information, without severing binding relations. This device is uniform in the sense that, syntactically speaking, the *only*

difference between one type of content and another consists in the labels they bear. But of course different types of content will differ in other ways as well—if they weren't they wouldn't have to be distinguished in the first place. Such differences may be procedural; for example, certain types of content are cancellable whilst others are not. In addition, differences may be semantic; for example, indexical and formal content are context dependent in distinctive ways. But at the representational level, all sorts of content are equal.

References

- Beaver, D. (2001). *Presupposition and Assertion in Dynamic Semantics*. Stanford: CSLI.
- von Stechow, P. (1994). *Restrictions on Quantifier Domains*. Ph. D. thesis, UMass Amherst.
- Gazdar, G. (1979). *Pragmatics: Implicature, Presupposition, and Logical Form*. New York: Academic Press.
- Geurts, B. (1999). *Presuppositions and Pronouns*. Oxford: Elsevier.
- Geurts, B. and R. van der Sandt (2003). Interpreting focus. Ms, University of Nijmegen. www.phil.kun.nl/tfl/bart.
- Kamp, H. and U. Reyle (1993). *From Discourse to Logic, Vol. 1*. Dordrecht: Kluwer.
- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, and H. Wettstein (Eds.), *Themes from Kaplan*, pp. 481–563. Oxford: Oxford University Press.
- Karttunen, L. and S. Peters (1979). Conventional implicature. In C. Oh and D. A. Dineen (Eds.), *Syntax and Semantics II: Presupposition*, pp. 1–56. New York: Academic Press.
- Krahmer, E. (1998). *Presupposition and Anaphora*. Stanford: CSLI.
- Krifka, M. (2001). Non-novel indefinites in adverbial quantification. In C. Condoravdi and G. Renardel de Lavalette (Eds.), *Logical Perspectives on Language and Information*, pp. 1–40. Stanford: CSLI.
- Kripke, S. (1972). Naming and necessity. In G. Harman and D. Davidson (Eds.), *Semantics of Natural Language*, pp. 253–355. Dordrecht: Reidel.
- Rooth, M. (1995). Indefinites, adverbs of quantification, and focus semantics. In G. Carlson and J. Pelletier (Eds.), *The Generic Book*, pp. 265–299. Chicago University Press.

- van der Sandt, R. (1992). Presupposition projection as anaphora resolution. *Journal of Semantics* 9, 333–377.
- Tasmowski-De Ryck, L. and P. Verluyten (1982). Linguistic control of pronouns. *Journal of Semantics* 1, 323–346.