DEFERRED INFORMATION: THE SEMANTICS OF COMMITMENT

LÁSZLÓ KÁLMÁN

Research Institute for Linguistics
Hungarian Academy of Sciences
Budapest I., P.O.Box 19.
H-1250 Hungary

0. INTRODUCTION

0.1. THE PROBLEMS

Imagine that the following pieces of discourse are uttered and heard without any particular linguistic or extra-linguistic antecedents:

(1) *Joe got married yesterday.* The minister *spoke very harshly.*
(2) *Joe got married yesterday.* The dog *barked very loudly.*

The radical difference between (1) and (2) is that, under standard assumptions on the cultural background, the hearer will know what minister the speaker has in mind in (1), whereas (s)he will have no idea of what *the dog* in (2) refers to. Clearly, the first sentence of these pieces of discourse *licenses* the use of *the minister* but not *the dog*.

There are at least three problems inherent in the phenomenon illustrated above:

1. The fact that the sentence *Joe got married yesterday* licenses the definite description *the minister* but not *the dog* does not follow from the truth conditional properties of the expressions involved. In particular, the truth of *Joe got married yesterday* does not *entail* the existence of a minister (which *the minister* presupposes), nor does it exclude the existence of a dog (which *the dog* presupposes). Consequently, a formal distinction between 'entailing a proposition' and 'licensing a presupposition' is to be made.

2. In addition to the fact that the existence of a minister in the context of (1) becomes a valid inference after uttering both sentences, other entailments, which do not follow from either the first or the second sentence (or their

László Kálmán and László Pólos (eds.), *Papers from the Second Symposium on Logic and Language*. Akadémiai Kiadó, Budapest, 1990.

conjunction) also arise (these entailments will be called pieces of *implicit information* in the following). In particular, the fact that Joe's marriage involved a church ceremony follows from the two sentences taken together as a sequence. Traditional model theoretic semantics is not able to account for this sort of interactions between consecutive sentences.

3. It seems obvious that the 'licensing' relation introduced above holds between the marriage event in the first sentence and the definite description *the minister* in the second. That is, the phenomenon in (1) is to be explained on the basis of the lexical/encyclopedic knowledge related to 'marriage' and 'minister'. The question is, now, how that knowledge is to be represented in order to account for the licensing relation. Moreover, it is to be examined whether only lexical/encyclopedic knowledge can license presuppositions without entailing them, or other types of information, such as discourse information, behave in the same way in certain cases.

0.2. STRUCTURE OF THE PAPER

In Section 1 of this paper, I will examine the problems illustrated in (1) and introduce some basic concepts that a solution should rest upon. In particular, I will define the *deferred information* content of an expression, that is, those pieces of meaning which are not entailed by the expression but are licensed as subsequent presuppositions. In Section 2, I will outline the basics of a theory, unfortunately dubbed *D.I.R.T.* for *Deferred Information Representation Theory*, which offers a way of formally representing deferred information and processes involving such representations. Section 3 is devoted to the application of D.I.R.T. to some celebrated linguistic problems. In particular, certain properties of *specific* vs. *non-specific indefinites* and *'scope ambiguities'* will be examined under D.I.R.T.'s approach. Finally, I will suggest that lexical meanings are to be represented in the same format as pieces of discourse.

1. ACCOMMODATION AND LICENSING

1.1. DEFERRED INFORMATION

Let me start with the examples in (1–2) above.

(1) *Joe got married yesterday.* The minister *spoke very harshly.*
(2) *Joe got married yesterday.* The dog *barked very loudly.*

As I said in the Introduction, *the minister* in (1) is licensed, whereas *the dog* in (2) is not. In this section, I will elaborate on the exact content of licensing and introduce some related concepts to make the problem clearer.

The first thing to call to attention is that both *the minister* and *the dog* have to be licensed because they are definite descriptions used anaphorically. Such definite descriptions presuppose the existence of exactly one individual with certain properties in the relevant universe of discourse. For example, *the minister* in (1) presupposes the existence of exactly one relevant minister, and *the dog* in (2) presupposes the existence of exactly one relevant dog. By 'relevant' I mean that a definite description requires that there be a single way to *connect* the entity whose existence is presupposed to the previous context. In (1), the presupposed entity is *m* (the minister), and the connection to be established is, presumably, '*m* conducted Joe's wedding ceremony'.

According to Lewis (1979), such presuppositions spring into existence automatically whenever a definite description is uttered. He calls this process the *accommodation* of presuppositions. In his opinion, it is a purely pragmatic issue whether a presupposition can be accommodated. That is, he considers the difference between (1) and (2) above a purely pragmatic one. I want to propose an entirely different view of this issue. Although it may be a pragmatic issue that the piece of information that licenses *the minister* in (1) is present in the first sentence (whereas no information is present that could license *the dog*), the way in which that piece of information is put to work cannot be a pragmatic issue, for the following reason. In principle, an information state in which a given presupposition is licensed can be described in terms of propositions under standard assumptions about semantics, irrespective of *where* those propositions come from. Assume that we have access to a set of propositions which represent *all* information relevant to a given licensing process, of pragmatic as well as lexical origin. Now, if we still cannot explain the licensing properties of the information state, which seems to be the case for example (1) if truth conditions alone are considered, then our view of propositions, our theory of semantics is not correct. It does not matter what our theory of pragmatics is like, provided that it furnishes those relevant pieces of information that characterize a given information state; all other duties, including the explanation of how a presupposition 'springs into existence', and how the licensing connection becomes implicit information, rest upon semantics as soon as propositions rather than situations are involved.

Obviously, then, the problem that modern truth conditional semantics faces with regard to (1–2) is twofold.

(i) As I have just pointed out, definite noun phrases are to be 'connected' to the previous context; but they do not necessarily correspond either to individual names occurring in the previous context or to individual variables bound by some operator. The latter alternative, i.e., conceiving of definite noun phrases as free variables which can be bound by operators in the previous context, would be feasible in a dynamic theory such as that of Heim (1982), Kamp (1981) or Groenendijk and Stokhof (1989). But

(ii) neither the first sentence of (1) nor the 'pragmatic' knowledge associated with it *entails* the existence of a church ceremony or a minister in the truth conditional sense of the word and, consequently, that sentence has no power to 'bind the variable' in the second sentence.

My conclusion is that, in order to account for the difference between (1) and (2) above, irrespective of its source, we have to rely on a theory of semantics that is both dynamic and not strictly truth conditional. The theory we need is to be dynamic, so as to allow cross-sentential 'binding' relations, and it is to be non-truth-conditional, so as to allow to distinguish between the entailments of a proposition and its licensing power. For example, we should be able to say that *the minister* in (1) is bound, so to say, by something in the first sentence, and we should be able to say that 'marriage' does not entail the existence of a minister, but it somehow licenses the presupposition that there is a minister, namely, the one who conducts the marriage ceremony. That is, we should also be able to explain that the 'binding' of *the minister* in (1) triggers the *implicit information* that Joe's marriage involved a church ceremony.

Before proceeding, let me summarize the concepts introduced so far:

(3) **Licensing**
An expression E is *licensed* by a piece of discourse D iff the presuppositions of E can be accommodated in the context created by D alone.

(4) **Implicit Information**
A piece of information i belongs to the *implicit information* content of (the meaning of) a piece of discourse D if D is a sequence of the pieces of discourse D_1 and D_2 such that "$D_1 \wedge D_2$" does not entail i, and D does entail i.

(5) **Deferred Information**
A piece of information i belongs to the *deferred information* content of (the meaning of) a piece of discourse D iff D does not entail i, and if i is the only presupposition of an expression E, then E is licensed by D.

1.2. DEFERRED INFORMATION AND UNDERSPECIFICATION

This section is devoted to the question what exactly belongs to the deferred information content of a particular expression. I will examine the knowledge associated with 'marriage' as assumed in examples (1–2) above, and suggest some general conclusions. In particular, I will argue that the *disjunctiveness* and the *generic character* of pieces of information lead to deferment, and claim that this is related to a more general concept of *underspecification*.

If we try to characterize the meaning of *marriage* relevant to the explanation of the phenomenon in (1) in terms of 'meaning postulates', the following seems a good first attempt:

(6) *Information related to 'marriage' (informal)*
  (i) 'in each model in which a marriage event *e* takes place, *e* includes a marriage ceremony event *c* as a rule';
  (ii) 'in each model in which a marriage ceremony event *c* takes place, *c* is either a church service or a common-law marriage as a rule';
  (iii) 'in each model in which a church service *c* takes place, there is a clergyman who conducts *c* as a rule'.

If we conceive of contexts as models of worlds, then — according to the above meaning postulates — the existence of a minister who conducts the marriage ceremony follows if certain conditions are met. Namely, the marriage ceremony is to be a religious one, and the context is to be 'normal' (according to the proviso 'as a rule' in each clause above).

It is obvious that, given the 'meaning postulates' in (6i–iii), the existence of a church ceremony or a minister does not *follow* from the fact that a marriage takes place. On the other hand, we independently know from the phenomenon of licensing that, in a case when only a marriage event is present in the context, there is a fundamental difference between assuming that the marriage involves a clergyman, on the one hand, and assuming that it involves a dog, on the other. It seems that the 'regular' character of the context or the truth of a member of a disjunction can be assumed, if necessary, without further provisos, whereas the presence of dogs in the context, which should be assumed as a premise on its own, cannot. In sum, I submit that a distinction should be made between propositions absent in a context altogether, on the one hand, and propositions in the scope of a *generic* one[1] ('as a rule') or embedded in a *disjunction*, on the other. The latter constitute pieces of *deferred information*.

Indeed, it can be easily checked that the existence of the minister who conducts the wedding ceremony will not follow from (6i–iii), not even if we make those closes non-generic (i.e., omit 'as a rule'), provided that (a) (6ii) remains a disjunction, or (b) (6i) or (6iii) is/are put in disjunction(s). Similarly, if there was no disjunction in (6i–iii), but at least one clause was generic, then the presence of a clergyman would no longer follow. On the other hand, it seems obvious that the licensing power of 'marriage' illustrated in (1) would not change if we altered its meaning in any of the above ways. Note that the question whether any of these variants constitutes at least partially appropriate meaning postulates for 'marriage' is a pragmatic issue and immaterial indeed from the point of view of the present discussion.

In sum, in the 'meaning postulates' discussed above, the existence of a clergyman is a piece of deferred information if and only if it is embedded in a consequent within the scope of a generic 'operator' ('as a rule') or within a member of a disjunction. Thus, I submit that at least the following sorts of pieces of information count as deferred:

(7) *Sources of Deferred Information*
   (a) a member of a disjunction;
   (b) a generic proposition;
   (c) a piece of information embedded in a piece of deferred information.

It should be noted that surface *or* does not necessarily express 'disjunction' in the sense of (7a) above. If the speaker refers to a marriage event, he may or may not know whether that marriage was a religious one. If he does, he would not use *or* if he was more explicit; if he does not, then he would not presuppose either alternative in subsequent discourse. Accordingly, after uttering a sentence like *You leave or I kill you* (or, alternatively, *If you don't leave, I kill you*), neither 'you leave' nor 'I kill you' are licensed presuppositions. Overt generic sentences, on the other hand, do behave like the pieces of information mentioned in (7b), i.e., they transfer deferred information only. It is a common-place in the literature that the truth conditions of generics are hard to establish. For example, it would be difficult to tell what property a model should have in order for *Ravens are black* to be true in it. This sentence certainly does not entail

$$\forall(x : \text{RAVEN}(x))[\text{BLACK}(x)].$$

On the other hand, it clearly makes it possible to accommodate the presupposition *x is black* for every raven x in subsequent discourse. For example, in

(8) *Ravens are black. ...*
   a. ... *An eagle chased a raven.* The poor black thing *could not escape.*
   b. ... Although *this bird is a raven, it is not black,*

the possibility of using the definite description in (a) or *although* in (b) shows that the presupposition that a raven is black can be accommodated in the context created by the first sentence.

In conclusion, I would like to suggest that a distinction should be made between those pieces of information that the meaning of an expression is unrelated to, on the one hand, and those pieces of information with respect to which it is *underspecified* (i.e., which belong to its deferred information content), on the other. For example, the meaning of *marriage* is unrelated to dogs, but underspecified with respect to the church ceremony; the meaning of *Ravens are black* is underspecified with respect to the set of ravens that it quantifies over. This interpretation of underspecification is parallel to the *phonological* interpretation of the word. For example, a consonant is unrelated to, or undefined for, say, tone, whereas a vowel which bears neither high nor low tone is underspecified with respect to tone in a language where tone matters. That is, underspecification is tightly linked to *abstractness* inasmuch as a vowel that is neither high nor low is abstract (or inasmuch as Locke's triangle, which is neither rectangular nor equilateral, is abstract).

Underspecification is not only opposed to unrelatedness, but also to *ambiguity*. One could argue that the word *marriage* might as well be ambiguous between

'religious marriage' and 'register marriage'. That is, we would have two separate lexical entries for *marriage*, which would mean that there are two words that accidentally coincide in terms of their phonetic shape. Such an assumption, however, would be weird for several reasons.

First, it seems reasonable to assume that, when a speaker utters an ambiguous expression, (s)he actually has one single reading in mind. It would be extremely strange to utter *Flying planes can be dangerous* without either having the intention of claiming that *Flying planes are dangerous* or meaning that *Flying planes is dangerous*. (Moreover, it is unlikely that you utter *Flying planes can be dangerous* and you mean either the conjunction or the disjunction of its two readings.) On the other hand, it is perfectly conceivable to speak about a marriage without knowing whether it was a religious one.

Second, the accidental coincidence of two meanings in *Flying planes can be dangerous* can be made explicit in terms of facts about English which are completely independent of the meaning of *-ing* forms, namely, the coincidence of the singular and plural forms of *can* (for example, replace *can* by *is/are* as in the previous paragraph). On the other hand, the alleged ambiguity of *marriage* cannot be made visible using pairs of linguistic contexts which differ in something independent of the meaning of *marriage*.

Finally, there may be empirically testable differences between underspecified meanings and ambiguous expressions. The different meanings of an ambiguous expression are related to each other in accidental ways. On the other hand, a meaning which is underspecified in a given respect is *entailed* by all more specific versions. That is, underspecification can be 'resolved' by further specification, i.e., by providing more information, which is not generally true for ambiguity. If I say *I bought a plane*, and you think I bought an aircraft, and I add *but its blade is all blunt*, then you will have to 'backtrack' to fix your meaning representation. This is probably more costly than, say, processing (1), which needs an extra assumption (i.e., the assumption that there was a church ceremony), but no backtracking. It remains to be seen whether this difference can be measured in terms of processing time.

2. D.I.R.T.

2.1. KEY CONCEPTS

In this section, I present the leading ideas that motivate the basic features of the semantic representation language of D.I.R.T., to be introduced in Section 2.2. In particular, I will argue that D.I.R.T.'s representations are to grasp the *hearer's* rather than the speaker's meaning representation; that D.I.R.T.'s representations will be built up in a *compositional* way in a particular sense of this term; I will explain the *limited representationalism* of D.I.R.T. and the assumption that meaning representations are *homogeneous*; finally, I will speculate on the *semantics* of D.I.R.T.'s language.

Hearer's meanings

The essential concept underlying D.I.R.T. is the *dynamic view of meaning*:

(9) **The Dynamic View of Meaning**
The meaning of an expression is best characterized by the change of information state that it can bring about in the hearer.

That is, the speaker's intentions do not appropriately characterize the meanings of the expressions that (s)he utters. In particular, the semantic representation of an expression whose meaning is underspecified in a given respect must grasp that underspecified meaning even though the speaker may have a more specific reading in mind. This is most obvious for meanings such as 'marriage' in (1) (cf. Section 1): even though the speaker has a religious marriage in mind, as is clear from the second sentence of (1), the semantic representation of the first sentence should not include that information.

Compositionality

As is clear from the previous section, I assume that pieces of information relevant to licensing subsequent presuppositions are to be present in meaning representations even if they are not significant in the truth conditional sense. Accordingly, the regularities of licensing are to be captured at the semantic representation level rather than attributed to 'pragmatic' factors or to interpretations. That is, these regularities follow from rules of combining meaning representations of sentences with representations of the previous contexts. The combination operation performing this will be called *incorporation* (of a representation into another one, or of a piece of discourse into another one) in what follows.

Therefore, D.I.R.T. differs from the two existing families of dynamic semantic theories, namely, versions of DRT (Discourse Representation Theory, cf. Kamp 1981) and DPL/DMG (Dynamic Predicate Logic, cf. Groenendijk and Stokhof, 1989, and Dynamic Montague Grammar, cf. Groenendijk and Stokhof, in this volume). Most versions of DRT are not *compositional* in the sense that the meaning

of a piece of discourse is not calculated from the meanings of the sentences that constitute it. DPL/DMG, on the other hand, are compositional, but conceive of meanings as identical to truth conditions. D.I.R.T. obeys *weak compositionality*, which means the following:

(10) **Weak Compositionality**
The meaning representation of a complex expression is calculated from the meaning representations of its constituent expressions.

Since meaning representations may contain information which does not influence the truth conditions of the expression (i.e., deferred information), weak compositionality does not entail truth conditional compositionality. Note that most versions of DRT are not even weakly compositional, because they do not assign meaning representations to individual sentences.

Limited Representationalism

Discourse Representation Theory is sometimes said to be *representationalist* because it recognizes the possibility of directly accessing and modifying the representation of the previuos context during the processing of subsequent sentences. DPL/DMG, on the other hand, are *anti-representationalist* because they do not allow for such operations. According to the latter theories, we have only access to the truth conditional aspects of the previous context. As suggested in the above, D.I.R.T. will be characterized by a certain *representationalism* because the non-truth-conditional features of the context are relevant to the explanation of licensing phenomena. On the other hand, the representationalism of D.I.R.T. is constrained in several ways.

The most important restriction relates to the use of deferred information in subsequent discourse. As we have seen in **Section 1**, deferred information may license presuppositions through licensing connections between an expression to be accommodated and the previous context. Establishing a licensing connection may result in implicit information, i.e., additional *commitment* (for example, with regard to the existence of a church ceremony in the example discussed earlier), but this can never 'override' previously established truth conditions. That is, D.I.R.T.'s context representations will be *uncommitted* with respect to many details of a context as a rule, and may specify possible ways of further commitment in the form of deferred information. So the idea is that all legitimate 'representationalist' operations will serve to further specify representations that are underspecified with respect to certain details. That is, operations defined over discourse representations must *monotonically increase the information content* of a context representation. One of those operations is the *incorporation* operation mentioned earlier; the other, called *unpacking* (cf. **Section 3.2**), will 'move' certain pieces of deferred information (occurring in the scope of a generic or in a disjunction) from their embedded positions, triggering additional commitment (i.e., it will enrich the truth conditions).

Homogeneity

Another idea which plays a central role in D.I.R.T.'s philosophy is that it would be extremely costly to assume that pieces of information originating from different sources (such as explicit, lexical, encyclopedic, pragmatic etc. information) are represented by different means within the representation of a piece of discourse and, in general, are processed by different mechanisms. Therefore, I will adopt the following principle, formulated on the basis of Bartsch (1987:2):

(11) *Homogeneity of Lexical and Discourse Representations*
Using the same syntax for discourse representations and for lexical frames secures a unified semantic theory of lexicon and syntax, a necessary condition of building up sentences from lexical items in a compositional manner. Lexical representations encode the contribution of lexical items to sentence meanings in the same way as sentence representations encode the contribution of sentences to discourse meanings or the contribution of smaller pieces of discourse to the meanings of larger ones.

According to this principle, lexical meaning representations are to be of the same form as discourse representations. Since lexical meanings typically express predicates and relations, there must be a way to embed them into sentence (and discourse) representations in a predicative function. As will be explained in Section 2.2, applying *predicative* (n-argument) discourse representations to discourse referents as arguments will be the only way of (recursively) embedding a discourse representation into another one in D.I.R.T.

D.I.R.T. Semantics

For the time being, D.I.R.T.'s representation language has no proper model theoretic interpretation, and it is not my ambition to present one in this paper. In Kálmán and Szabó (forthcoming), we present an interpretation which ignores deferred pieces of information altogether, but is suitable for making the meaning of D.I.R.T. representations explicit in all other respects. A full interpretation of D.I.R.T. representations would require a significant enrichment of the simple models used there. As a matter of course, a full interpretation of D.I.R.T. representations would include both

(i) assigning truth conditions to deferred pieces of information and
(ii) defining a rule-to-rule semantics for representational operations.

In the remaining part of this section, I will outline a possible way of solving problem (i) above.

The solution rests on the assumptions that

(i) the semantic value of a formula can be *undefined* in certain possible worlds;
(ii) there is an accessibility relation over possible worlds; and

(iii) the meaning of a sentence is best captured by its *updating function* rather than its truth conditions.

Assumptions (i) and (ii) make it possible to make a distinction between

(a) formulae which have no semantic value in a world w because they 'make no sense', in which case they will have no semantic value in any world accessible to w, either; and
(b) formulae which have no semantic value in a world because that world is not specified for them, in which case they will be true or false in some accessible world.

That is, I want any possible world w' accessible to a possible world w (written: $w' \sqsubset w$) to be 'more specific' than w in some relevant respect. For example, if *Joe got married yesterday* is true in a world w, but the value of *Joe had a religious marriage* is undefined, then there must be a possible world $w' \sqsubset w$ in which the latter sentence is true or false, but there need not be one in which a relevant dog participated in Joe's wedding.

I will assume that the accessibility relation "$\sqsubset$" is irreflexive, antisymmetric and transitive, i.e., it is a strict partial ordering over the class of possible worlds. If a proposition p is true (false) in a given possible world, then it will be true (false) in all worlds accessible to it and, for any $w_2 \sqsubset w_1$, there must be a proposition q whose semantic value is undefined in w_1 and true or false in w_2.

On the other hand, assumption (iii) above makes it possible to conceive of propositions *not* as the entire class of possible worlds in which they are true, but only a subclass of them. As I have suggested above, while truth conditions impose constraints on the class of possible worlds in which a sentence is true, pieces of deferred information determine what possible worlds must be accessible to those worlds. That is, the updating function of sentences can be described as follows:

(12) **Sentence Meaning**
The meaning of a sentence is a mapping from classes W of possible worlds to classes W' of possible worlds such that
(i) W' has to meet the truth conditions of the sentence; and
(ii) if p is a piece of deferred information carried by the sentence, then $\Diamond p$ has to be true in each member of W'. That is, if $w_1 \in W'$, then there has to be a possible world $w_2 \sqsubset w_1$ such that p is true in w_2.

For example, the sentence *Joe got married yesterday* is certainly true in a possible world w^* where Joe had a religious marriage, but w^* will not be in W' since, presumably, there will be no world $w^{*\prime} \sqsubset w^*$ in which Joe had a common-law marriage, although the latter proposition also belongs to the deferred information content of the sentence under scrutiny. (For the sake of simplicity, I ignore the fact that one and the same marriage event may involve both a church ceremony and a civil ceremony.) In sum, the updating function associated with *Joe got married*

yesterday will pick out some subclass of W in which the semantic values of both deferred pieces of information are undefined.

Note that this modal semantic framework would make it possible to make a distinction between cases in which the semantic value of a formula is undefined in a given possible world, on the one hand, and genuine 'truth value gaps', i.e., cases in which a formula makes no sense in a given world, on the other. In the latter cases, the semantic value of the formula in question will be undefined *in each and every accessible possible world*. I will rely on the latter sort of undefined value in what follows, but I will not deal with the possible-world semantics of D.I.R.T. in this paper. As a matter of course, if a modal calculus for D.I.R.T. along the above lines was available, then there would be a chance to make D.I.R.T. compositional in the strong sense. Assessing this option, however, requires further research.

2.2. THE REPRESENTATION LANGUAGE

In this section, I will more or less informally outline the representation language of D.I.R.T. I will characterize the semantics of this language in a very superficial manner only (cf. Kálmán and Szabó (forthcoming) for more details). In particular, I will ignore the interpretation of deferred pieces of information for the reasons mentioned above.

Given that D.I.R.T. considers inter-sentential licensing, the main source of implicit information, of extreme importance, it makes a syntactic distinction between that part of the meaning of a sentence which is to be *anchored* to the previous context, i.e., which specifies the target of the change to be effected in an information state, on the one hand, and all other pieces of information, on the other. As in Discourse Representation Theory, the formulae of the language are called *discourse representations*, which are interpreted as the conjunction of the *conditions* that constitute them. Moreover, D.I.R.T. formulae are split into two parts, a set of *anchor conditions*, and a single *predicate condition*.

The semantic value of a formula can be *true*, *false* or *undefined*. The undefined semantic value (or *semantic value gap*) is assigned to formulae whose presuppositions cannot be accommodated. The undefined value will be assigned, among others, to statements concerning non-existing objects, or context-changing instructions whose target is not in the current information state. Therefore, anchor conditions are presuppositional, i.e., the truth of their conjunction is a necessary condition for a formula to be either true or false. If some of the anchor conditions are not true, which means that anchoring is not possible, then the formula will have no semantic value. Otherwise, its semantic value will be the same as that of the predicate condition.

Conditions are of two sorts, namely, an n-ary predicate applied to n arguments (variables) is a condition, and the conjunction of a finite number of conditions is a condition. The semantic value of an n-ary predicate is a pair $\langle \Phi, \Psi \rangle$ where Φ

and Ψ are disjoint sets of n-tuples over the universe. The predicate is true for n arguments if the n-tuple consisting of the values of those arguments is in Φ (i.e., the *extension* of the predicate); it is false for them if the n-tuple of values is in Ψ (i.e., the *anti-extension* of the predicate); it is undefined otherwise. The conjunction of a set of conditions is true if each member condition is true; it is undefined if at least one member is undefined; it is false otherwise. We define the negation of a predicate; if the semantic value of the predicate is $\langle \Phi, \Psi \rangle$, then the value of its negation is $\langle \Psi, \Phi \rangle$. The semantic interpretation of disjunction is related to deferment, so I will ignore it here.

The variables which occur free in the conditions constituting a formula are all understood as existentially quantified. (We do not need individual names if we postulate a distinguished set of individual predicates with one-element extensions.) These variables are called the *discourse referents* of the formula. We require that each discourse referent occur free in at least one anchor condition. The notation of discourse representations is as follows. If $a_1, \ldots, a_k$ are the anchor conditions and p is the predicate condition, then

$$[a_1, \ldots, a_k][p]$$

is a discourse representation. If the variables occurring free in such a representation are $r_1, \ldots, r_m$, then the above formula can be equivalently written as

r_1	...	r_m
a_1		
$\vdots$		
a_k		
p		

Discourse representations in predicative function correspond to λ-*abstractions*. They differ from formulae in that they contain certain variables distinguished as λ-*variables* (furthermore, these need not occur free in any anchor condition). If the number of λ-variables is n, then the resulting λ-abstraction will qualify as a 'normal' n-ary predicate of the language. λ-abstractions will be written as

$\lambda v_1 \ldots v_n[s_1, \ldots, s_k][p] =$

λv_1	...	λv_n	r_1	...	r_m
s_1					
$\vdots$					
s_k					
p					

where $v_1, \ldots, v_n$ are the λ-variables, and $r_1, \ldots, r_m$ are the discourse referents (i.e., variables not bound by λ and occurring free in conditions).

The interpretation of λ-abstractions is rather complex, and will be only sketched in what follows. Two subsets of referents are to be distinguished among the discourse referents of a λ-abstraction. These are called *floating* referents and *existential* referents, respectively. To define these subsets, I will refer to a graph consisting of referents and λ-variables as vertices, and edges labelled by conditions. An edge labelled by a given condition runs between two vertices just in case both vertices occur free in that condition. Now, a referent is existential iff all paths which lead to λ-variables from it go through an edge labelled by the predicate condition. That is, an existential referent has no other connection 'to the external world', i.e., the arguments to which the λ-abstraction is applied, than the predicate condition of the λ-abstraction. Both the existence of an existential referent and the conditions in which it occurs are *claimed* by the λ-abstraction (and interpreted accordingly). The existence of floating referents, on the other hand, is *presupposed* by the predicate, together with the anchor conditions in which those referents occur. Accordingly, if those presuppositions are not true when the predicate is applied to a given list of arguments, then that list of arguments will not belong to either the extension or the anti-extension of the λ-abstraction. So 'embedded' presuppositions can give rise to undefined semantic values, although they have no role in anchoring.

For example, the λ-abstraction corresponding to

(13) *x saw a unicorn*

contains an existential referent corresponding to *a unicorn*:

(13′) $\lambda x[\text{UNICORN}(u)][\text{SAW}(x,u)] =$

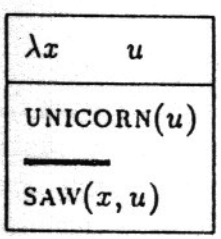

The existence of an object corresponding to u above (together with the condition "UNICORN(u)") is claimed by this predicate (or any condition in which it is applied to an argument), because the only condition that establishes a relationship between u and the λ-variable x is the predicate condition. On the other hand, the λ-abstraction corresponding to

(14) *x beats x's donkey*

is as follows:

(14′) $\lambda x[\text{DONKEY}(d), \text{OWNS}(x, d)][\text{BEATS}(x, d)] =$

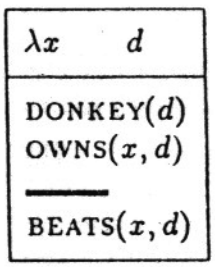

Here d, i.e., the referent corresponding to *a donkey*, is a floating referent, because it is related to a λ-variable (namely, x) through a condition which is not the predicate condition (namely, "OWNS(x, d)"; the predicate condition is "BEATS(x, d)"). Therefore, the donkey's existence is presupposed. Anchor conditions which contain free occurrences of floating referents are called *floating conditions*.

In D.I.R.T., quantification is treated by second-order predicates and *distributive λ-abstractions*. A positively (or negatively) distributive λ-abstraction is a λ-abstraction in which one or more arguments are marked by a special λ-operator λ^+ (or λ^-). If the kth λ-variable of a λ-abstraction is marked by λ^+ or λ^-, then the kth argument to which the λ-abstraction is applied has to be a group referent (or else a semantic value gap arises). Such an application is interpreted as if the non-distributive counterpart of the predicate applied to each member (or no member) of the group represented by the kth argument. For example, the sentence

(15) *Every farmer saw a unicorn*

would be represented as follows:

(15′) $\left[\text{EVERY-FARMER}(F)\right]\left[\lambda^+ f[\,][\lambda f'[\text{UNICORN}(u)][\text{SAW}(f', u)](f)](F)\right] =$

$=$

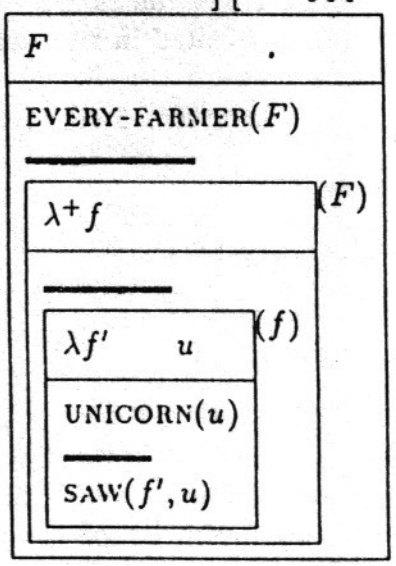

Here EVERY-FARMER is a second-order predicate. The embedded λ-abstraction is positively distributive. Therefore, its application to F is interpreted as 'for each $f \in F$, ...'. The representation of *No farmer saw a unicorn* would be

analogous, with "λ^-" replacing "λ^+" in the embedded box (which would be interpreted, then, as 'for no $f \in F \ldots$'). Positively and negatively distributive λ-operators cannot co-occur in λ-abstractions. For example, a λ-abstraction of the form $\lambda^+x\lambda^-y[a_1,\ldots,a_k][p]$ is ill-formed.

3. D.I.R.T.'S APPLICATIONS

In this section, I will examine certain linguistic phenomena in the light of D.I.R.T.'s principles. In **Section 3.1**, the difference between *specific* and *non-specific indefinites* will be presented and explained. In **Section 3.2**, I will deal with the two representational operations, i.e., *incorporation* and *unpacking*, which I mentioned in **Section 2.1**, to shed light on so-called *'scope ambiguities'*. Finally, in **Section 3.3**, I will use these concepts to analyze the first problem addressed in this paper, i.e., the licensing of definite descriptions, and speculate about lexical meaning representations.

3.1. SPECIFIC AND NON-SPECIFIC INDEFINITES

In dynamic theories of semantics, *indefinite expressions* are translated as discourse markers that are *referentially independent of previous discourse entities* (indefinites 'introduce new discourse markers', as it were). In this section, I will first show that this is not always the case. In particular, certain indefinites (which I will refer to as *specific*[2]) are to be licensed by the context in a way similar to anaphoric expressions. Other indefinites (which I will call *non-specific*) need not be licensed. The two types of indefinites are represented in two different ways in D.I.R.T.

The two types of indefinites that are to be distinguished can be illustrated by the following sentences:

(16) A nun *has not arrived.*
'one of the nuns', 'one of them, who is a nun'

(17) *There is* a lion *in the garden.*
*'one of the lions', *'one of them, who is a lion'

A nun in (16) is *specific*, for it has to be licensed by the previous context, whereas *a lion* in (17) is *non-specific*, because it does not have to. Note that, as the glosses show, *a lion* in (17) cannot be interpreted as specific at all. This is due to the semantics of *there is*. On the other hand, the specificity of the noun phrase *a nun* in (16) most probably follows from its being the *topic* of the sentence (or its 'Aristotelian subject', i.e., what the sentence is about).

As the above examples show, specific indefinites are to be licensed by *containment* rather than identification. That is, they are to be considered as *contained*

in some contextually available entity (typically, they are members/parts of groups or effects of events/processes). If anaphoric definite noun phrases presuppose the existence of just one suitable antecedent, specific indefinites presuppose the existence of a whole range of entities which may contain the object that they refer to.[3]

Non-specific indefinites, on the other hand, do not presuppose the existence of anything. Instead, they *claim* the existence of a given object. For example, the existence of a lion is part of the main claim of (17) (cf. example (13) in Section 2). Now, given the interpretation of D.I.R.T. representations sketched above, the difference between specific and non-specific indefinites falls out immediately if we guarantee that the former correspond to discourse referents in the outermost discourse universe, whereas the latter be represented by existential referents within the λ-abstraction in the predicate condition.

In discourse representation theories, the existence of an object is expressed by its presence in the discourse universe. On the other hand, D.I.R.T. makes a distinction between presupposing vs. claiming the existence of an object, as we have seen earlier. As is clear from the above, this distinction is somehow related to 'topicness' in that an indefinite Aristotelian subject has to be specific. In other words, I will stipulate that topics have an anchoring function. Moreover, it will have to be stipulated that those discourse referents in the outermost discourse universe of the meaning representation of a sentence which are not anaphoric are to be licensed by containment rather than identity. For example, in terms of what has been said so far, the representation of (16) above is as follows:

(16) *A nun has not arrived.*

(16′) [NUN(n)][HAS-NOT-ARRIVED(n)] =

n
NUN(n)
HAS-NOT-ARRIVED(n)

In order to incorporate this representation into a context representation, the incorporation algorithm is to find a (group) referent N in the previous context such that the condition $n \in N$ can be accommodated. Therefore, *a nun* will be implicitly interpreted as 'one of them, who is a nun'.

Notice that, under the assumption just made, referents corresponding to domains of quantification, such as the referent F in (15′), qualify as specific indefinites. Therefore, they are to be licensed by the previous context by larger domains. That is, in contradistinction to traditional approaches to universal quantification, *every farmer* in (15) will be interpreted as 'every farmer in some domain X available in the context'. In other words, a condition of the form $F \subseteq X$ is to be accommodated in order to incorporate the representation (15′) into previous contexts.

Obviously, the same holds for negative quantification (involving λ^{-}-abstractions, e.g., *No farmer...*). The consequences of this lie outside the scope of the present paper.

As for non-specific indefinites, they are to be translated by existential referents within the λ-abstraction of the predicate condition.[4] In some cases, such as

(18) *Joe saw a unicorn,*

(13′) $[\text{JOE}(j)]\,[\lambda j'[\text{UNICORN}(u)][\text{SAW}(j', u)](j)] =$

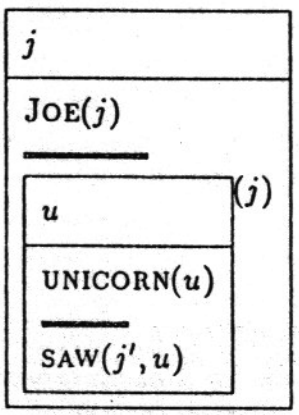

this is obviously the case. Given the syntactic structure of (18), there is no reason why we should put u in the outermost discourse universe, nor is there any reason why u should be a floating referent (cf. the representation in (13′)).[5] On the other hand, it could be objected that a sentence like (19) below is ambiguous:

(19) *A nun arrived.*
  a. 'one of the nuns' (or 'one of them, who is a nun') 'arrived'
  b. cca. 'there arrived a nun'

I submit that the appropriate representation of the reading in (19b) is the following:

(19′) b. $[\text{LOC}(e)]\,[\lambda e'[\text{NUN}(n)][\text{ARRIVED}(e', n)](e)] =$

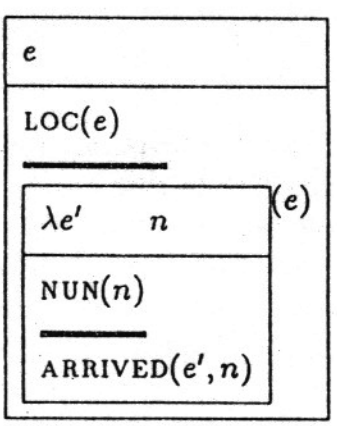

where the condition "LOC(e)" intends to express that e is a referent corresponding to a spatio-temporal location (or a 'Davidsonian' referent in the sense of Kratzer, 1988). As a matter of fact, such referents should be present in (13′–16′) as well, where I omitted them for the sake of convenience. In the case of (19′b), however,

this referent plays an important role because, intuitively, the incorporation of this sentence requires that its (indefinite) spatio-temporal location be licensed by the previous context (and, as opposed to the reading in (19a), only that referent has to be licensed in this case).

To sum up, I assume that the two readings of (19) is a true ambiguity, i.e., we have to assign them two different representations, one similar to (16′) above, on the one hand, and (19′b) above, on the other. I would be inclined to say that the two readings correspond to two different syntactic structures, namely,

(19) *A nun arrived.*
  a. [SUBJ *A nun*] [PRED *arrived.*] (specific indefinite reading);
  b. [SUBJ] [PRED *A nun arrived.*] (non-specific indefinite reading).

The assumption that the two readings of (19) correspond to two different syntactic structures is supported by the fact that, in the case of (19b), only the noun phrase bears primary sentence stress.

3.2. INCORPORATION AND UNPACKING

In what follows, I will first outline some basic features of the incorporation operation. In particular, I will argue that sentences exhibiting a *subject shift* (with respect to the previous discourse) are incorporated differently from those which do not. (This is the topic of **Section 3.2.1**, which is based on Szabó's, 1989, presentation.) Next, in **Section 3.2.2**, I will introduce the *unpacking* operation, which leads to *commitment*, i.e., gives rise to implicit information. Finally, in **Section 3.2.3**, I will suggest that so-called *'scope ambiguities'* are due to the effects of non-specific indefinites and unpacking.

3.2.1. INCORPORATION

As was mentioned earlier in this paper, *incorporation* is an operation that merges two discourse representations. That is, 'incorporation' is actually a cover term for the extremely complex rule system which may be called 'discourse syntax' (cf. Scha and Polanyi, 1988), or even 'syntax' in general, provided that the construction of sentence meaning representations is also described in terms of 'incorporation' rules. In this section, I will concentrate on one feature of incorporation only, namely, the one which captures the dynamics of discourse processing. This feature lies in the stipulation that, as soon as a sentence is incorporated into a piece of discourse, the claim of that sentence becomes the claim of the entire piece of discourse, while the previous claim becomes presuppositional. That is, as soon as a speaker utters a new sentence, (s)he presupposes all previous claims. There are two basic sub-cases of this process, namely, when the *topic* of a new sentence differs from the topic of the previous claim (topics will coincide with the subjects in the examples that

follow), on the one hand, and those cases in which there is no 'subject shift', on the other.

The piece of discourse in (20) below is an example of 'subject shift':

(20) *Ivan is walking.* [IVAN(i)][IS-WALKING(i)] *John is whistling.* [JOHN(j)][IS-WHISTLING(j)]

i
IVAN(i)
IS-WALKING(i)

j
JOHN(j)
IS-WHISTLING(j)

The incorporation of the second representation into the first proceeds as follows. First, the anchor condition of the second representation is to be satisfied within the first representation. Let me assume that this is possible in (20), i.e., *John* is licensed in the context. As a result, the anchor condition "JOHN(j)" will be an anchor condition in the resulting representation. Second, in terms of what has been said above, the claim of the context (i.e., the first representation in (20)) becomes presuppositional, that is, is 'moved' among the anchor conditions, and is replaced by the new claim "IS-WHISTLING(j)":

(20′) *Ivan is walking. John is whistling.* [IVAN(i), JOHN(j), IS-WALKING(i)][IS-WHISTLING(j)]

i j
IVAN(i) JOHN(j) IS-WALKING(i)
IS-WHISTLING(j)

If the topic of the new sentence is the same as the previous one, the incorporation process is somewhat more complex. Consider:

(21) *Ivan is walking.* [IVAN(i)][IS-WALKING(i)] *He is whistling.* [HE(x)][IS-WHISTLING(x)]

i
IVAN(i)
IS-WALKING(i)

x
HE(x)
IS-WHISTLING(x)

Assume now that the incorporation mechanism is able to retrieve the referent i in the first representation as the antecedent of x in the second, and to decide that there is no topic shift. Now, there is evidence to the effect that the resulting representation contains a *complex predicate* (which means, roughly, 'is walking and whistling'). This happens when all the discourse anchors of the second sentence, including the ones left implicit in (21), such as location and modality, agree with the previous sentence. In those cases, verb phrase anaphora (e.g., *Joe is, too*) may refer back to the complex predicate as well as just the last claim (i.e., *Joe is, too* can be interpreted either as 'Joe is walking and whistling, too' or 'Joe is whistling, too').[6]

Accordingly, the predicate condition of the representation of the piece of discourse in (21) will contain the previous claim, i.e., "IS-WALKING(i)". On the other hand, the general condition that previous claims become presuppositional is to be satisfied as well. This problem can be solved by assuming that the resulting predicate condition contains a λ-abstraction in which the previous claim is a *floating condition*:

(21′) *Ivan is walking. He is whistling.*
[IVAN(i)] [$\lambda i'$[IS-WALKING(i')][IS-WHISTLING(i')](i)]

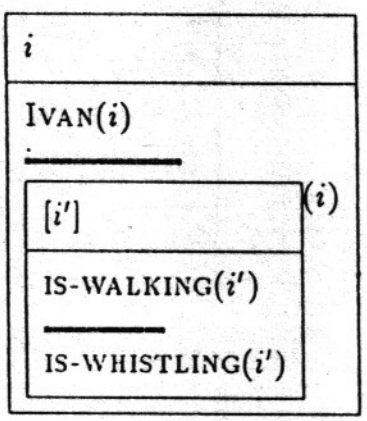

From the technical point of view, this representation can be produced by equivalently converting the primitive predicate IS-WALKING in the first representation of (21) into the abstraction "$\lambda i'$[ ][IS-WALKING(i')]", then incorporating the new claim *into that abstraction* (with i' replacing x). The rules of incorporation are as before, so the condition "IS-WALKING(i')" becomes an anchor condition. As such, it will be a floating condition, because it contains a λ-variable. Therefore, it will be interpreted as a presupposition. If IS-WALKING is a λ-abstraction rather than a primitive predicate, then the new claim is to be incorporated into it in the same way.[7]

3.2.2. UNPACKING

Now that the most essential features of incorporation have been presented, I can go a step forward in its details. We have seen that the referents in the outermost discourse universe are to be licensed within the formula (in the case of subject shift) or λ-abstraction (when the topic does not change) into which the new representation is to be incorporated. We have also seen that, to license a referent, other referents are to be available. This is clearest in the case of anaphoric referents, which want identical antecedents, as in (21) above. In this section, I will explain what happens if the antecedent of an anaphor (or, in general, a piece of information that we want to use for licensing) is not to be found immediately within the discourse universe of the formula or λ-abstraction where the new sentence is incorporated.

This happens when a referent in the discourse universe of a more deeply embedded λ-abstraction serves as an antecedent for an anaphor to be licensed higher up in the expression:

(22) *A farmer saw a unicorn.* $[\text{FARMER}(f)]\,[\lambda f'[\text{UNICORN}(u)][\text{SAW}(f',u)](f)]$

It was black. $[\text{IT}(x)][\text{WAS-BLACK}(x)]$

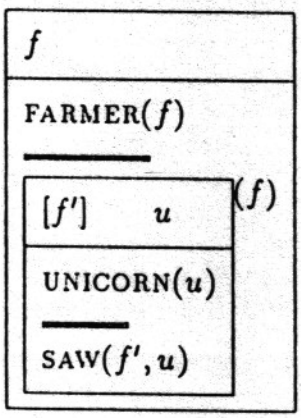

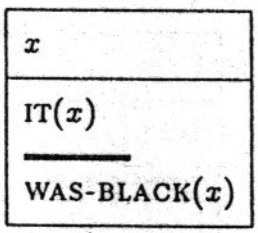

Clearly, there is a subject shift between the two sentences in (22), because f ('a farmer') cannot be the antecedent of x ('it'). Consequently, the referent x is to be licensed directly in the first discourse representation rather than the λ-abstraction in its predicate. Although no referent other than f is immediately present in the discourse universe of that formula, it seems to be the case that the referent u (*a unicorn*), embedded in the λ-abstraction of the predicate, can act as an antecedent of x (*it*).

Assume that the incorporation mechanism can retrieve that antecedent. What happens next? The following representation, which would result from the blind application of the principles mentioned so far, is ill-formed:

(*22′) *A farmer saw a unicorn. It was black.*
*[FARMER(f), $\lambda f'$[UNICORN(u)][SAW(f', u)](f)] [WAS-BLACK($\underline{u}$)]

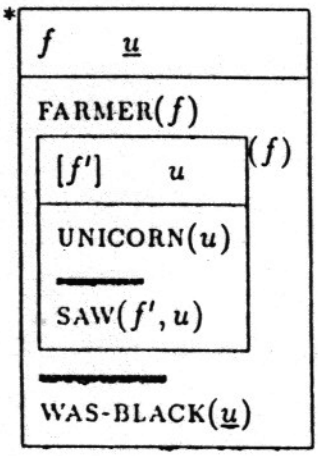

In D.I.R.T., each variable is interpreted as existentially quantified *within its λ-abstraction or formula*. That is, λ-abstractions are referentially closed. (This is the so-called *Principle of Referential Privacy* or *PRP*). Therefore, the underlined occurrences of u in (*22′) above are independent of the non-underlined occurrences of the same variable, since they belong to different λ-abstractions. As a consequence, underlined u is a referent which does not occur in any anchor condition, which makes the representation in (*22′) ill-formed.

To solve this problem, I stipulate that, whenever the antecedent of an anaphor is found in an embedded λ-abstraction, it has to be *unpacked* from there by the incorporation mechanism in such a way that it belongs to the λ-abstraction or formula in which the anaphoric referent resides. Before I go into the details of unpacking, let me show the correct counterpart of (*22′):

(22′) *A farmer saw a unicorn. It was black.*
[FARMER(f), UNICORN(u), $\lambda f'u'$[UNICORN(u')][SAW(f', u')](f, u)]
[WAS-BLACK(u)]

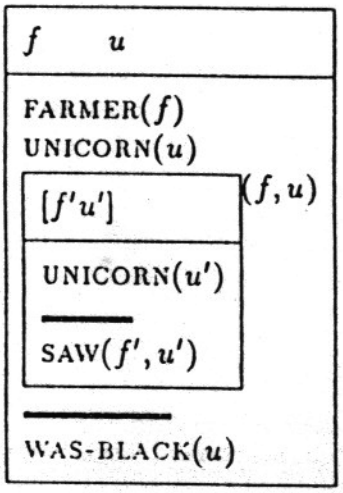

What happened here is the following. First, the anchor condition "UNICORN(u)" has been copied into the anchor condition part of the formula.[8] The internal occurrences of u have been replaced by a 'fresh' variable u' for better readability (variable names are irrelevant because of the PRP); these occurrences have become λ-bound. This stage of the incorporation is represented in (22$_i$) below:

(22$_i$) *A farmer saw a unicorn.*
[FARMER(f), UNICORN(u)] [$\lambda f'u'$[UNICORN(u')][SAW(f', u')](f, u)]

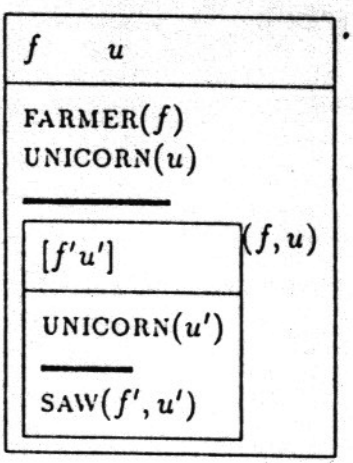

It *was black.*
[IT(x)][WAS-BLACK(x)]

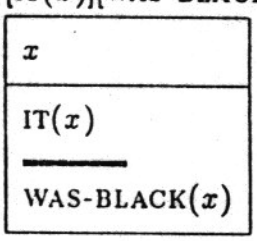

From here, the procedure is analogous to the one seen in (20) (i.e., in the case of subject shift). The result is (22′) above.

Technically, unpacking can be defined with reference to the graph associated with formulae and λ-abstractions which I introduced in **Section 2.2** (i.e., vertices are variables, edges are labelled with conditions, and an edge labelled with a given

condition runs between two variables just in case they both occur free in that condition). This is necessary because, in addition to the target discourse referent r of the unpacking operation (i.e., the referent to be unpacked), it may be necessary to unpack other referents as well. Those referents are said to be *dependent* on r. A discourse referent r' is dependent on r just in case each path from r' to a λ-variable runs through r. The unpacking operation can be described in three steps:

(i) copy the anchor conditions in which the referent r or the referents depending on it occur into the formula or λ-abstraction in which the λ-abstraction occurs;
(ii) convert r and the referents depending on it into λ-variables (and append them to the λ-variables already present) in the λ-abstraction in which they used to be discourse referents (instead of this, I used the fresh variable u' rather than u in (22,) to enhance readability);
(iii) append r and the referents depending on it to the arguments of the λ-abstraction.

The importance of unpacking lies in the fact that it may convert deferred pieces information into non-deferred pieces information provided that its target used to reside in an abstraction representing deferred information. (I only mention here that unpacking — as antecedent retrieval in general — does not have access to λ^--abstractions.) That is, the mechanism sketched in this section accounts for the fact that licensing may lead to additional commitment.

3.2.3. 'SCOPE AMBIGUITY'

In this section, I will argue that so-called 'scope ambiguity' phenomena can be accounted for in the framework presented so far without postulating multiple representations. I consider 'scope ambiguity' a case of underspecification which is not directly related to deferred information and, therefore, it can effectively motivate some of the ideas that D.I.R.T. is based upon.

It is obvious that non-specific indefinites, such as the one in (18) (*Joe saw a unicorn*), can license anaphoric expressions. Consequently, in terms of what has been said above, the existential referents corresponding to them are to be 'moved' into the main discourse universe of a context representation under certain circumstances. In what follows, I will argue that this fact is responsible for 'scope ambiguity'.

Note first that sentences exhibiting 'scope ambiguity', such as the one that I presented in (15), contain non-specific indefinites:

(15) *Every farmer saw a unicorn.*

(15′) $\left[\text{EVERY-FARMER}(F)\right]\left[\lambda^{+}f[\,][\lambda f'[\text{UNICORN}(u)][\text{SAW}(f',u)](f)](F)\right] =$

$=$

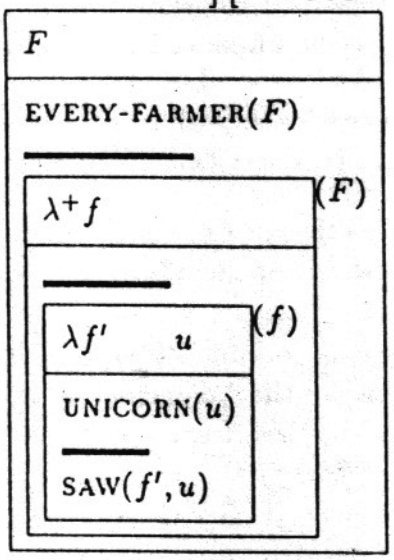

The distributive λ-abstraction in the predicate condition of (15′) claims the existence of a unicorn for each member of the group F. That is, in terms of the semantics of D.I.R.T. representations, the representation in (15′) entails that, for any $f \in F$, SAW-A-UNICORN(f) holds. This means that the representation in question corresponds to the 'narrow-scope existential' reading of (15). As a matter of course, the above representation will be true in models in which all farmers saw one and the same unicorn, but it does not entail the 'wide-scope existential' reading.

On the other hand, if the referent corresponding to *a unicorn* (i.e., u) belonged to the main discourse universe, as if it had been *unpacked*, as shown in (15$_{\prime}$) below, then the representation would entail that, for any $f \in F$, SAW(f, u) holds, i.e., such a representation would embody the 'wide-scope existential' reading:

(15$_{\prime}$)

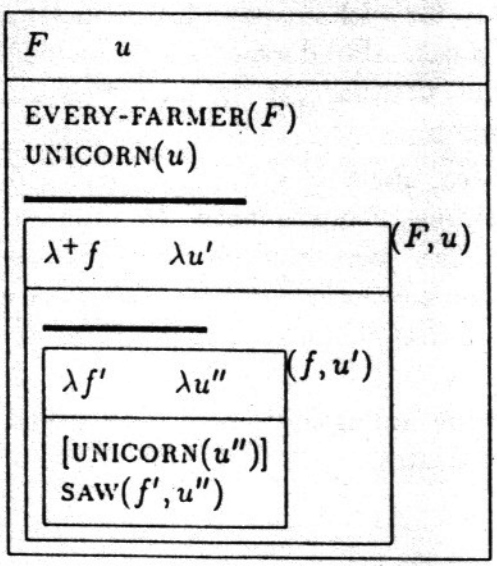

If we were to conceive of the sentence in (15) as ambiguous, then we would have to say that it does not have one single meaning representation, but its two readings can be represented as in (15′) and (15′′), respectively. There is another way of proceeding, however. It can be argued that the 'narrow-scope existential' reading of (15) is *underspecified* as to whether all farmers saw the same unicorn; we could take that reading as basic, and regard the 'wide-scope existential' reading as derived by additional information about the coincidence of the unicorns. This piece of information, however, need not be explicit. For example, consider the following piece of discourse:

(23) *Every farmer saw a unicorn.* The unicorn *was white.*

The definite description in (23) triggers the 'wide-scope existential' reading of (15) without explicitly stating that the unicorns seen coincide. The mechanism at work in this example must be sensitive to the meaning representation of the first sentence, since definite descriptions in general do not have the power of 'collapsing' previous referents. Consider, for example:

(24) *I saw a unicorn$_i$. Joe saw one$_j$, too.* *The unicorn$_{i=j}$ *was white.*

If the piece of discourse in (24) is well-formed at all, the definite noun phrase *the unicorn* does not trigger the reading under which the two unicorns seen coincide.

In sum, I propose that the sentence in (15) has one single meaning representation, namely, the one in (15′), which corresponds to the 'narrow-scope existential', or non-specific indefinite, reading. This is supported by the fact that the indefinite noun phrase *a unicorn* in (15) belongs to the syntactic predicate of the sentence. Furthermore, a continuation like the one in (23) may trigger the 'wide-scope existential', or specific indefinite, reading, represented as in (15′′). The discourse process shown in (23) involves the accommodation of the definite description, which triggers further specification of the context (of the sort we have seen in earlier examples), which leads to the restructuring of the previous context representation.

The unpacking operation and the Principle of Referential Privacy, as its independent motivation, establish a meaningful relationship between the facts that the noun phrase *the unicorn* in (23) above needs an antecedent, on the one hand, and that the first sentence acquires a 'wide-scope existential' reading, on the other, since the unicorn referent of (15′) is to be 'moved' into the main discourse universe (cf. (15′′)) in terms of the rules of incorporation.

3.3. D.I.R.T. AND LEXICAL MEANINGS

In this section, I will briefly outline the ideas that the previous discussion in Section 3 suggests for lexical meaning. I will first outline D.I.R.T.'s approach to problems such as the one presented at the very beginning of the present paper (i.e., the implicit information arising from example (1)), then I will present evidence

to the effect that lexical predicates are to be represented in the same format as discourse representations and, finally, I will make a remark on a possible view of the interface between syntax and semantics.

As is clear from the preceding sections, the rules of incorporation make strong predictions for what pieces of deferred information can and will give rise to implicit information in a piece of discourse. In particular, the rules of unpacking serve to keep graphs connected, therefore, they guarantee that all information which is really necessary for establishing a licensing connection becomes part of the non-deferred information content of a representation. For example, in the example (1),

(1) *Joe got married yesterday.* The minister *spoke very harshly,*

whatever unique connection may exist between the 'minister' referent and the marriage event (e.g., the one involving the the church ceremony and the person who conducts it) somewhere embedded in its representation will be implicit in the meaning of (1). That is, in addition to the minister, the marriage ceremony, the church and, eventually, the particular rite will be part of the 'hard-core' information content of the discourse.

If this approach is correct, then the D.I.R.T. representation of the lexical meaning of the verb *marry* should encode the 'meaning postulates' associated with this lexical item (e.g., those in (6), **Section 1.2**) in a discourse representation format, and that discourse representation (more precisely, that λ-abstraction) has to be present in the representation of (1). Obviously, the referents corresponding to the ceremony or the minister will be existential referents embedded in the lexical meaning representation, since they need not be licensed by the context. For example, the representation of the 'meaning postulate' in (6i) can be written as follows:

(25) MARRIAGE =

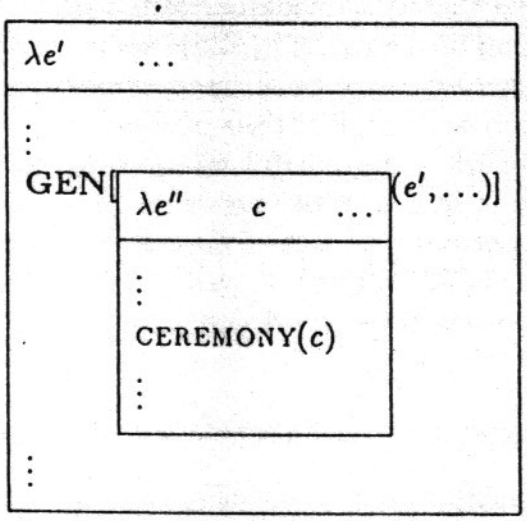

Here GEN is intended to express that the condition which follows it in square brackets is 'generic'. The 'meaning postulates' in (6ii–iii) may be embedded in the predicate called CEREMONY in (25) above. In the case of (1), the referent

corresponding to the minister who conducted the ceremony is to be unpacked in several steps until it reaches the outermost discourse universe of the above representation (and even higher, if so needed). As a side-effect, the referent c, corresponding to the ceremony, will be unpacked as well:

(26) MARRIAGE′ =

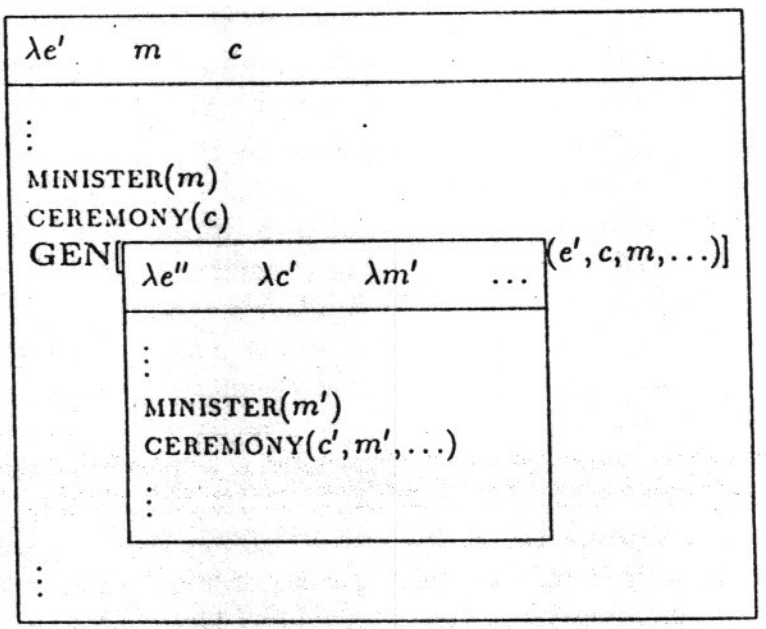

The assumption that lexical meaning representations are of the same format as discourse representations (more precisely, as λ-abstractions) does not entail in itself that their behaviour is exactly the same as that of λ-abstractions originating from explicit discourse information. Nevertheless, the more similarities we discover, the better this assumption is motivated. It is quite interesting, for example, that the 'scope ambiguity' phenomena discussed in Section 3.2.2 can be reproduced using lexical predicates. Compare (23) with (27–28) below:

(23) *Every farmer saw a unicorn.* The unicorn *was white.*
(27) *Every farmer was killed.* The murderer *is still at large.*
(28) *Every farmer was sued.* The trial *will be held tomorrow.*

These 'scope ambiguity' phenomena can be easily accounted for in D.I.R.T. if we assume that the lexical meaning representations corresponding to *was killed* and *was sued* contain embedded existential referents corresponding to the murderer(s) and the trial(s), respectively.

The examples in (27–28) above also suggest that the licensing properties of lexical predicates are not always due to the encyclopedic or pragmatic information associated with them. For example, the fact that *was killed* licenses the definite description *the murderer* may be related to the fact that the same entity, namely, the agent of killing, can be referred to by syntactic means, i.e., a *by*-phrase. Under this assumption, an optional argument in the lexical meaning representation of a verb corresponds to an existential referent that licenses inter-sentential anaphors

and, furthermore, has the additional (eventually, idiosyncratic) property of licensing an argument slot. That is, it could be assumed that certain constructions *within* sentences (such as constructions consisting of a verb and its optional argument) are held together by links similar to inter-sentential licensing. In the case of optional arguments, such an analysis would be certainly superior to one that doubles lexical representations (i.e., posits homonymous representations, one with an obligatory argument slot and one with no argument slot at all).

It would be worth examining how far the metaphor of 'construction-internal licensing' can be stretched, that is, to what extent the syntactic properties of a lexical predicate can be accidental with respect to its meaning representation. If the homogeneity principle in (11) (**Section 2.1**) is to be taken seriously, then sentence representations of the kind presented in the previous section are to be constructed from lexical representations of the same sort. Obviously, this means that intermediate phrasal constituents and other constructions between the lexical and sentence levels will also be represented by fomulae and abstractions of the D.I.R.T. language. While the construction of discourse representations is governed by (eventually universal) rules of 'discourse syntax', encoded in the incorporation mechanism, the construction of lower-level representations may be driven by (eventually language specific) rules of 'real syntax' which are associated with combination operations similar to what I have called *incorporation* in the case of the combination of pieces of discourse. The idea is, then, that each grammatical relation (e.g., 'modification', 'specification', 'complementation', or whatever the relevant relations will be) corresponds to a way of constructing complex representations from smaller ones. If those smaller representations are of the same format as higher-level ones, then some of the distinctions made at the sentence representation level (e.g., different types of conditions and referents) may be relevant at the lower levels.

The above speculation directly affects the interface between lexical semantics and syntax. Since the representations which determine what kinds of constituents are licensed in a construction are meaning representations, it is possible in principle that the categorial requirements which the members of a construction are to meet are derived from the meanings of those members. This seems to apply at least to finer-grained sub-category distinctions. For example, whether an adjective can modify a given noun must be derivable from the meanings of the adjective and the noun, and even a 'licensing connection' seems to be necessary when certain discrepancies arise. Take, for instance, the celebrated example of *a quick coffee*. The adjective *quick* inherently modifies processes. Therefore, in order to accommodate it, the meaning of *coffee* is to be searched for a 'relevant' process which presumably is associated with it. Thus, *quick* 'presupposes', as it were, a relevant connection between the process it modifies and the noun it is attached to, in pretty much the same way as *the minister* in (1) presupposes a relevant connection between the minister it refers to and the context it occurs in.

It is not my aim in this paper to explore the predictions inherent in the approach to the meaning of syntactic constructions outlined above. What I intended to show in this section is that the assumption that the meaning representations of pieces of discourse, sentences, phrases and lexical items are of the same sort may shed new light on the semantic factors involved in the build-up of syntactic constructions.

ACKNOWLEDGMENTS

I wish to thank Jeff Goldberg and Jeroen Groenendijk for valuable comments on earlier versions and presentations of this paper. I am grateful to Tamás Mihálydeák, László Pólos and Zoltán Szabó, who contributed to various stages of the development of D.I.R.T., and Anna Szabolcsi, who contributed to various stages of writing this paper.

NOTES

1. In what follows, sentences such as *Ravens are black*, which are not about genera, but express rule-like statements that allow exceptions, will be called *generic* for lack of a better term. This is common practice in the literature.

2. This interpretation of the term *specific indefinite* originates from Szabolcsi (1983). Subsequently, she used the distinction between specific and non-specific indefinite to account for the so-called 'definiteness effect' (Szabolcsi, 1986) and for constraints on complex-predicate formation (Szabolcsi, 1987).

3. Webber (1988) calls such indefinites 'anaphoric'. As is clear from the above, I reserve this term for expressions that want identical antecedents.

4. Although such referents are accessible antecedents of anaphoric expressions just the same as presupposed referents (e.g., n in (16′)) are, the two still behave differently. The arguments for this come from VP ellipsis and VP anaphora cases discussed by Szabó (1989) and Kálmán and Szabó (forthcoming).

5. As can be seen in (13′), the referent u originally belongs to the λ-abstraction corresponding to the predicate (i.e., *saw a unicorn*). I assume that there are certain cases in which such an embedded referent has to 'escape', so to say, from its λ-abstraction (and be replaced by a λ-variable, incrementing the arity of the predicate). In particular, proper names, anaphoric referents and ranges of quantification (such as F in (15′)) have to be in the outermost discourse universe of a sentence representation. Since this stipulation is related to the problem of constructing sentence representations from the meaning representations of constituents, I will not go into it in the present paper.

6. For more details see Kálmán and Szabó (forthcoming).

7. Kálmán and Szabó (forthcoming) analyze certain *modal subordination* phenomena (cf. Roberts, 1987) using this feature of the incorporation operation.

8. The redundancy which arises in this way is harmless in the case at hand, and is indispensable in others (e.g., in the case of downward monotonic quantifiers; cf. Kálmán and Szabó, forthcoming).

9. As a matter of fact, (30b) is grammatical under the reading 'Ivan has sent mail', but that is a specialized use of *write*, which should be considered independent, so to say, of the lexical item(s) in (29) and (30). For example, *write* 'send mail' optionally takes a *to*-phrase argument, while the others do not.

10. Given the D.I.R.T. machinery presented so far, this phenomenon must be related to unpacking, because perfect forms of *write* are committed to the existence of a written output.

REFERENCES

Bartsch, R. 1987. 'Frame representations and discourse representations'. *ITLI* Working Paper, Amsterdam.

Groenendijk, J. and M. Stokhof. 1989. 'Dynamic Predicate Logic'. Ms., to appear in *Linguistics and Philosophy*.

Groenendijk, J. and M. Stokhof. In this volume. 'Dynamic Montague Grammar'.

Heim, I. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. PhD. diss., University of Massachusetts, Amherst.

Kálmán, L. and Z. Szabó. Forthcoming. 'D.I.R.T.: An overview'. In *Proceedings of the Seventh Amsterdam Colloquium*, ITLI, Amsterdam.

Kamp, H. 1981. 'A theory of truth and semantic representation'. In J. Groenendijk, T. Janssen, and M. Stokhof, eds., *Formal Methods in the Study of Language*. Mathematisch Centrum, Amsterdam.

Kratzer, A. 1988. 'Stage-level and individual-level predicates'. Paper read at the Tübingen Conference on Genericity.

Lewis, D. 1979. 'Scorekeeping in a language game'. *Journal of Philosophical Logic* 8, 339–359.

Roberts, C. 1987. *Modal Subordination, Anaphora, and Distributivity*. Ph.D. diss., University of Massachusetts, Amherst.

Scha, R. and L. Polanyi. 1988. 'An augmented context free grammar for discourse'. *Proceedings of COLING '88*, Budapest.

Szabó, Z. 1989. 'DR construction in D.I.R.T.'. Paper read at the Second Symposium on Logic and Language, Hajdúszoboszló, Hungary, September 5–8.

Szabolcsi, A. 1983. 'On the specific/non-specific distinction'. *Nyelvtudományi Közlemények* 85, 83–92. (In Hungarian.)

Szabolcsi, A. 1986. 'From the definiteness effect to lexical integrity'. In W. Abraham and S. de Meij, eds., *Topic, Focus, and Configurationality*. Benjamins, Amsterdam. 321–348.

Szabolcsi, A. 1987. 'Indefinites in complex predicates'. *Theoretical Linguistic Research*.

Webber, B.L. 1988. 'Tense as discourse anaphor'. *Computational Linguistics* 14, No. 2, 61–73.