

Situations and the structure of content

François Recanati

► **To cite this version:**

François Recanati. Situations and the structure of content. Kumiko Murasugi, Robert Stainton. Philosophy and Linguistics, Westview Press, pp.113-165, 1999, 4. <ijn_00000175>

HAL Id: ijn_00000175

https://jeannicod.ccsd.cnrs.fr/ijn_00000175

Submitted on 8 Oct 2002

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Situations and the structure of content

F. Recanati

CREA

1 rue Descartes, F-75005 Paris

recanati@poly.polytechnique.fr

I. Introduction

- §1.1 Austinian semantics
- §1.2 The debate over 'what is said'
- §1.3 Austinian propositions and minimal content

II. The 'support' relation

- §2.1 Situations, facts, and worlds
- §2.2 Backgrounding and variable polyadicity
- §2.3 Relativisation

III. δ -structures

- §3.1 Mentioning situations
- §3.2 δ -structures in English
- §3.3 Intentional δ -structures
- §3.4 World-shifting

IV. Complex and schematic facts

- §4.1 Complex facts
- §4.2 Schematic and quantificational facts
- §4.3 Questions

V. Informational structure

- §5.1 Exercised situations as 'topics'
- §5.2 Nucleus vs. minimal content
- §5.3 Topics and truth-conditions

I. INTRODUCTION

§1.1 Austinian semantics

According to Austin (1971: 122), the meaning of an utterance can be analysed in two factors: the situation s the utterance refers to, and the utterance's descriptive content σ . The utterance is true iff σ adequately classifies s , that is, iff the situation referred to "is of a type with which the sentence is correlated by the descriptive conventions" of the language.

Different interpretations of Austin's view have been put forward in the literature. Two main questions arise in this respect: first, what is 'the situation referred to'? Second, what is the utterance's 'descriptive content'? One common answer to the first question must be ruled out from the outset because it is too uncharitable an interpretation of Austin's view: The situation referred to by an utterance such as 'The cat is on the mat' cannot be the state of affairs represented by the utterance, that is, the cat's being on the mat, for then there would be no way for an utterance to be false — the utterance's descriptive content would be built into the situation referred to. A better interpretation, close to the letter of Austin's writings, takes the situation referred to be a complex consisting of all the entities (things, times, places, etc.) referred to in the utterance. Let us assume that 'the cat' and 'the mat' are referring expressions in such a way that 'The cat is on the mat' refers to the cat c , to the mat m and to the time of utterance τ . The situation referred to, therefore, is the sequence $\langle c, m, \tau \rangle$, while the descriptive content of the utterance is a Russellian 'propositional function', expressible by the open sentence ' x is on y at t '. Austin's definition of truth thus entails that:

'The cat is on the mat' is true iff $\langle c, m, \tau \rangle$ satisfies ' x is on y at t '

Such an interpretation of Austin was offered (and elaborated) by Herman Cappelen and Josh Dever in a seminar in Berkeley in 1994.

A third interpretation is due to Barwise and Etchemendy (1987). Instead of construing the descriptive content of the utterance as exclusive of all demonstrative elements, as Austin himself does, they construe it as a proposition in the traditional sense, say, as a Kaplanian content. Thus the descriptive content of 'The cat is on the mat' is the singular proposition consisting of the 'on' relation and the above-mentioned sequence:

$\langle \text{On}, \langle c, m, \tau \rangle \rangle$

The situation referred to, for Barwise and Etchemendy, is the portion of reality the speaker intends to be saying something about. To be sure, this situation is normally indicated by various elements in the sentence, elements which are governed by those 'demonstrative conventions' of which Austin says that they determine the situation referred to. Thus the present tense indicates that the speaker intends to talk about what's going at the time of utterance, and the referring expression 'the cat' indicates that it is the whereabouts of the cat that the speaker is concerned with. The point of the Barwise-Etchemendy proposal, however, is that the situation talked about need *not* be linguistically articulated. They give the following example:

If the sentence "Claire has the ace of hearts" is used to describe a particular poker hand, then on the Austinian view the speaker has made a claim that the relevant situation is of the type in which Claire has the ace of hearts. Notice that such a claim could fail simply because Claire wasn't present, even if Claire had the ace of hearts in a card game across town. (Barwise and Etchemendy 1987: 29)

The particular poker hand the speaker is commenting on is not linguistically articulated; it is given only by the context. That is why the situation talked about cannot be merely that to which the sentence (token) is correlated by what Austin calls the demonstrative conventions of the language.¹ A more significant departure from Austin, however, is the claim that the entities referred to by the demonstrative constituents of the sentence (e.g. Claire, referred to by the proper name 'Claire') can themselves be constitutive of the descriptive content of the utterance. In the present example, Barwise and Etchemendy take the descriptive content to be the singular proposition that Claire has the ace of hearts. That proposition determines a type of situation, viz. the type of situation in which Claire has the ace of hearts. In Barwise's and Etchemendy's framework, the utterance says of the particular poker hand talked about that it is a situation of that type.

Though it significantly departs from the letter of Austin's writings, the Barwise-Etchemendy interpretation is faithful to its spirit. It entails that the content of every utterance, whether indexical or not, depends upon the context (since it is the context which determines what counts as the situation referred to); a view which certainly was Austin's. Be that as it may, I find the Barwise-Etchemendy view interesting in its own

¹ Note, however, that we could include among the 'demonstrative conventions' of the language conventions which concern non-linguistic aspects of the utterance (such as the direction of the gaze of the speaker, and so forth). See Beeson 1979.

right, and it is that view which I will elaborate in this paper.² What I find most interesting in that view is the claim that there are two levels of truth-evaluable content, two propositions which an utterance expresses: a content in the classical, Kaplanian sense (e.g., for the utterance 'Claire has the ace of hearts', the proposition that Claire has the ace of hearts); and an 'Austinian' proposition, to the effect that the situation talked about supports that content.

§1.2 The debate over 'what is said'

One reason why I am interested in elaborating the semantic theory suggested by the Barwise-Etchemendy interpretation of Austin is that it connects up with recent discussions in pragmatics, concerning the notion of 'what is said'. In this section I summarize those discussions.

Anyone who has reflected on the sentence meaning/utterance meaning distinction knows that a simple distinction is in fact insufficient. Two equally important distinctions must be made.

- First, there is the distinction between the linguistic meaning of a sentence-type, and what is said (the proposition expressed) by an utterance of the sentence. For example, the English sentence 'I am French' has a certain meaning which, *qua* meaning of a sentence-type, is not affected by changes in the context of utterance. This context-independent meaning contrasts with the context-dependent propositions which the sentence expresses with respect to particular contexts. Thus 'I am French', said by me, expresses the proposition that I am French; if you utter the sentence, it expresses a different proposition, even though its linguistic meaning remains the same across contexts of use.
- Second, we have the distinction between what is actually said and what is merely 'conveyed' by the utterance. My utterance of 'I am French' expresses the proposition that I am French, but there are contexts in which it conveys much more. Suppose that, having been asked whether I can cook, I reply: 'I am French'. Clearly my utterance (in this context) provides an affirmative answer to the question. The meaning of the utterance in such a case includes more than what is literally said; it also includes what the utterance 'implicates'.

'What is said' being a term common to both distinctions, we end up with a triad:

² The basic insight underlying that view can also be found in McCarthy 1993.

sentence meaning

vs.

what is said

vs.

what is implicated.

The distinguishing characteristic of sentence meaning (the linguistic meaning of the sentence-type) is that it is conventional and context-independent. Moreover, in general at least, it falls short of constituting a complete proposition, i.e. something truth-evaluable. In contrast, both 'what is said' and 'what is implicated' are context-dependent and propositional. The difference between 'what is said' and 'what is implicated' is that the former is constrained by sentence meaning in a way in which the implicatures aren't. What is said results from fleshing out the meaning of the sentence (which is like a semantic 'skeleton') so as to make it propositional. The propositions one can arrive at through this process of 'fleshing out' are constrained by the skeleton which serves as input to the process. Thus 'I am French' can express an indefinite number of propositions, but the propositions in question all have to be compatible with the semantic potential of the sentence; that is why the English sentence 'I am French' cannot express the proposition that kangaroos have tails. There is no such constraint on the propositions which an utterance of the sentence can communicate through the mechanism of implicature. Given enough background, an utterance of 'I am French' might implicate that kangaroos have tails. What's implicated is implicated by virtue of an inference, and the inference chain can (in principle) be as long and involve as many background assumptions as one wishes.

The basic triad can be mapped back onto the simple sentence meaning/speaker's meaning distinction by grouping together two of the three levels. There are two ways to do it, corresponding to two interpretations for the triad. The first interpretation stresses the close connection between sentence meaning and what is said; together, sentence meaning and what is said constitute the *literal meaning* of the utterance as opposed to what *the speaker* means:

literal meaning	}	sentence meaning
		what is said

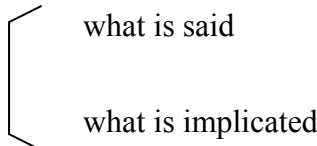
vs.

speaker's meaning

The other interpretation stresses the commonality between what is said and what is implicated, both of which are taken to be pragmatically determined:

sentence meaning

vs.

speaker's meaning 

Essential to this interpretation is the claim that 'what is said', though constrained by the meaning of the sentence, is not as tightly constrained as is traditionally thought.

The first interpretation corresponds to a widespread doctrine which I call pragmatic Minimalism. According to that doctrine, 'what is said' departs from the meaning of the sentence (and incorporates contextual elements) *only when this is necessary to 'complete' the meaning of the sentence and make it propositional*. In other words, the distance between sentence meaning and what is said is kept to a minimum (hence the name 'Minimalism'). Thus an indexical sentence such as 'He is tall' does not express a complete proposition unless a referent has been contextually assigned to the demonstrative pronoun 'he', which acts like a free variable in need of contextual instantiation. 'Saturation' (or, in Kent Bach's terminology, 'completion') is the contextual process whereby the meaning of such a sentence is completed and made propositional. Other contextual processes — e.g. the inference process generating implicatures — are semantically *optional* in the sense that the aspects of meaning they generate are dispensable: The utterance would still express a complete proposition without them. According to Minimalism, those extra constituents of meaning which are not necessary for propositionality are external to what is said. The only justification for including some pragmatically determined constituent of meaning into what is said (as opposed to what is merely conveyed) is the indispensability of such a constituent — the fact that the utterance would not express a complete proposition if the context did not provide such a constituent.

As an illustration, consider examples (1) and (2), respectively borrowed from Sperber and Wilson (1986) and Kent Bach (1994):

- (1) I've had breakfast
- (2) You are not going to die

Arguably, the first sentence, 'I've had breakfast', expresses the proposition that the speaker has had breakfast before the time of utterance. Strictly speaking this

proposition would be true if the speaker had had breakfast twenty years ago and never since. That is clearly not what the speaker means (when she answers the question 'Do you want something to eat' and replies 'I've had breakfast'); she means something much more specific, namely that she's had breakfast *this morning*. This aspect of speaker's meaning, however, has to be construed as external to what is said and as being merely conveyed, in the same way in which the utterer of 'I am French' implies, but does not say, that he is a good cook. That is so because the 'minimal' interpretation, to the effect that the speaker's life was not entirely breakfastless, is sufficient to make the utterance propositional without having to bring in the implicit reference to a particular time.

The same thing holds for the other example. Kent Bach, to whom it is due, imagines a child crying because of a minor cut and her mother uttering (2) in response. What is meant is: 'You're not going to die from that cut'. But literally the utterance expresses the proposition that the kid will not die *tout court* — as if he or she was immortal. The extra element contextually provided (the implicit reference to the cut) is not necessary for the utterance to express a complete proposition, hence it does not constitute a component of what is said in the minimalist sense.

Opposed to Minimalism is pragmatic *Maximalism* — the view I have defended and elaborated over the years (Recanati 1993, 1995). According to that view, the relevant distinction is not between mandatory and optional contextual processes, but between those that are 'primary' and those that are 'secondary'. Primary pragmatic processes help determine what is said; secondary pragmatic processes are inferential processes: they take 'what is said' as input and yield further propositions (the implicatures) as output. Now primary pragmatic processes include not only saturation, but also 'optional' processes such as free enrichment. That is so because, in general, the notion of 'what is said' we need to capture the input to secondary, inferential processes already incorporates contextual elements of the optional variety. In the examples above, the speaker implies various things by saying what she does: she implies that she is not hungry, or that the cut is not serious. Those implicatures can be worked out only if the speaker is recognized as expressing the proposition that she's had breakfast *this morning*, or that the child won't die *from that cut*.

Minimalism*Sentence meaning*

saturation

*what is said_{min}*

optional processes

*what is communicated***Maximalism***Sentence meaning*primary pragmatic processes
(saturation and optional processes
such as enrichment or transfer)*What is said_{max}*

secondary pragmatic processes

*what is communicated*


What is said in the maximalist sense corresponds to the intuitive truth-conditions of the utterance, that is, to the content of the statement as the participants in the conversation themselves would gloss it. In contrast, the literal truth-conditions posited as part of the Minimalist analysis turn out to be very different from the intuitive truth-conditions which untutored conversational participants would ascribe to the utterance. Minimalist theorists acknowledge (and sometimes even applaud) this divorce between what is said as a theoretical entity and our intuitions of what is said. Maximalists like myself find it unbearable: insofar as 'saying' is a particular form of meaning_{mn}, the statement which is made by uttering a sentence depends upon, and can hardly be severed from, the speaker's publicly recognizable intentions.

§1.3 Austinian propositions and minimal content

It has been suggested that the dispute may be verbal to some extent. Why should we not distinguish two notions of what is said: a purely semantic, minimalist notion, and a pragmatic notion ('what is stated' as opposed to what is implied)? If we accept this suggestion (voiced by Nathan Salmon, 1991 and Kent Bach, 1994), we end up with four levels instead of three:

literal meaning	{	<ul style="list-style-type: none"> • Sentence meaning (character) • what is said_{min}
vs.		
speaker's meaning	{	<ul style="list-style-type: none"> • what is said_{max} • what is implicated

Such a compromise view would seem to be acceptable to both the Minimalist and the Maximalist. The Minimalist wants to isolate a purely semantic notion of content, that is, a notion of the content of a sentence (with respect to a context) which is compositionally determined and takes pragmatic elements on board only when this is necessary. The Maximalist wants to capture the intuitive notion of 'what is said' (as opposed to what is implied) and stresses that what is said in that sense is, to a large extent, determined in a top down manner by the context. Both notions can be had simultaneously if one accepts to replace the traditional triad by a four-level picture.



Is the suggested compromise workable? Ultimately I don't think it is, but I agree that we should try to go as far as we can in that oecumenical direction. It is here that Austinian semantics can be most useful. Remember the main claim: An utterance is true iff the situation s it refers to supports the fact σ it expresses. The complete truth-conditional content of an utterance therefore is:

$$s \models \sigma$$

This complete content, or Austinian proposition, is distinct from the proposition σ on the right hand side of the support sign ' \models '. So there are two levels of content, two propositions which every utterance expresses: the fact σ which is stated, and the Austinian proposition to the effect that the situation of reference supports that fact.

There is much in common between the right hand side σ in the Austinian proposition (I call it the 'nucleus') and minimalist content. Consider the examples I gave to illustrate Minimalism:

- (1) I've had breakfast
- (2) You're not going to die

In the Austinian framework we can put the meaning constituent generated by free enrichment on the situational side, and keep the nucleus 'minimal' or nearly so. Thus in (1) the speaker refers to a temporally circumscribed situation (viz. the situation on the day of utterance) and characterizes *it* as a situation in which the speaker has had breakfast. In the same way, the mother who utters (2) refers to the specific situation brought about by the child's cut, and characterizes *it* as a situation in which the child does not die. The proposition on the right hand side of the Austinian formula is thus exactly what the sentence seems to express; there is no need to consider the sentence as somehow 'elliptical'. The same thing holds for the cases in which a quantifier is contextually restricted: the contextual domain of the quantifier is nothing other than the situation talked about, in the Austinian framework (Recanati 1996). It follows that the nucleus itself is unaffected by contextual restrictions on the domain of quantification. The Austinian framework, therefore, helps us to implement the compromise view imagined by Salmon and Bach. As in the latter we have four levels, with the notion of 'what is said' split in two:

- sentence meaning
- what is said₁: nucleus
- what is said₂: Austinian proposition
- what is implicated

Ultimately, I think the equation between nucleus and minimal content will have to be given up (see §5.2 below). Still, there is a rather obvious similarity between them, and interpreting the nucleus in this light — that is, as an approximation of the sort of 'literal content' the Minimalist has in mind — usefully constrains our understanding of situations and the general form which Austinian semantics must take.

II. THE 'SUPPORT' RELATION

§2.1 Situations, facts, and worlds

According to Austinian semantics, utterances refer to situations and what they say (their Kaplanian content) is predicated of those situations. But what are situations? When confronted with this question, we should be wary of conflating situations in the ordinary sense and situations in the technical sense. The technical notion of situation we find in situation theory has much wider application than the ordinary notion. Situations, in the ordinary sense, are 'eventualities' — perhaps a specific sub-class of 'stative' eventualities. Eventualities are a special type of entity, distinct from other types (e.g. individual objects, groups, places, or times). But every entity, of whatever sort, is a 'situation' in the technical sense, insofar as we can think of it and store information about it. At least this is how I will elaborate the notion of 'situation' which is at the center of Austinian semantics.

The reason why I keep using the possibly misleading term 'situation' is this. A situation in the ordinary sense supports facts and can be seen as a micro-universe (the world itself being 'the collection of all the facts', as Wittgenstein wrote). Take the situation here and now as I am writing this. It contains a number of facts: the fact that I am typing on my computer, the fact that my children are playing with a cat and keep interrupting me, etc. Situations in the ordinary sense contain entities having properties and standing in relation to one another. That feature is precisely what the technical notion of 'situation' is meant to capture (and to generalize). What characterizes situations in the technical sense is their double nature: They are entities and, *qua* entities, they have properties and bear relations to other entities; at the same time they are like a micro-universe *containing* entities having properties and bearing relations to other entities. We capture this double nature by associating each entity (each situation) with a set of facts (the *factual set* of the entity).

Facts, or states of affairs, are triples consisting of an n-place relation, a sequence of n appropriate arguments, and a polarity which is + if the relation obtains between the arguments and - if it doesn't. (A fact is negative if the polarity is '-'. By default, a fact is considered as positive; thus I will often omit the '+' sign in what follows.) The items in the sequence of arguments are called the (material)³ constituents of the fact. Each situation determines a set of facts, namely the set of facts in which the situation itself is a constituent. Thus consider the situation $s =$ London in the 19th century. Suppose it's a

³ The *formal* constituents of a fact are the n-place relation and the polarity.

fact that: *in 19th century London, poor children used to work hard*. That fact concerns the situation s (it contains it as a constituent), hence it goes into the set of facts which the situation determines. In the same way, the fact that Freud was bald concerns Freud, hence it goes into the set of facts determined by the 'situation' Freud. The set of facts determined by a situation contains all the facts of which that situation is a constituent.

Suppose it was a contingent fact that Freud was bald. We want to capture this contingency. It will not be captured if the set of facts associated with a situation is considered as definitive of that situation. If a situation *is* the set of facts it determines, then it cannot be a contingent fact that a situation contains such and such a fact. That is why we must distinguish situations (entities) from the sets of facts they determine. Situations determine sets of facts only relative to a world.

What is a world? On the present view there are two components in a world w : a domain of entities, $\text{Dom}(w)$; and a function w from entities (situations) in that domain to sets of facts. The set of facts associated with a situation $s \in \text{Dom}(w)$ is the set $w(s)$. The contingency of Freud's baldness is thus captured: the fact that Freud was bald belongs to the factual set actually associated with Freud, but in other possible worlds the factual set associated with Freud would not have contained that fact. Different worlds can associate different sets of facts with the same situations (entities) if the situations in question are in the domains of those worlds. Thus Nixon has the property of resigning in the actual world, but not in other possible worlds in which he exists. This simply means that the fact that Nixon resigned belongs to $@(\text{Nixon})$, but not to $w(\text{Nixon})$, for some world $w \neq @$.

Let us now turn to Austinian propositions. An Austinian proposition says that a situation, s , supports a fact σ . The 'support' relation is relative to a world, hence instead of ' $s \models \sigma$ ', we should write:

$$s \models_w \sigma$$

In what follows, however, I will subscript the relation only when necessary (i.e. when distinct worlds are simultaneously at issue: see §3.3-4).

Knowing what situations and facts are, can we define the 'support' relation? The obvious candidate is:

$$s \models \sigma \text{ iff } \sigma \in w(s)$$

But that will not do. Consider, for example, the utterance 'Everyone is asleep'. We want σ to approximate the minimal content of the utterance. This means that the contextual restriction (everyone *in group G*) does not affect σ but only the situational component s : 'Everyone is asleep' thus states the same (nonpersistent) fact σ whatever situation

we are referring to. But if we accept the definition of 'support' above, we'll have to take σ to be the (persistent) fact that everyone is asleep *in situation s*. That is so because the facts in $w(s)$ contain s as a constituent. Owing to that feature, the definition of 'support' above entails that the contextual restriction of the quantifier affects the nucleus σ . If we want the latter to approximate the minimal content of the utterance, we must change the definition of 'support'. I suggest the following:

$$s \models \sigma \text{ iff } \sigma \in w(s)_r$$

The subscript 'r' indicates that the facts of $w(s)$ have been 'relativised', where relativisation is a form of 'backgrounding'. To that topic I now turn.

§2.2 Backgrounding and variable polyadicity

From a given relation R , further relations can be *generated* by increasing or decreasing the arity of R . Consider the relation between the seller and the buyer in a commercial transaction. There are a number of argument roles in that relation: not only *buyer* and *seller*, but also the *goods* to be transferred, the *money*, and so forth. As Fillmore pointed out, verbs like 'buy' and 'sell' cannot be understood without mastering the complete 'frame' which serves to define both (see e.g. Fillmore 1975, 1982, Fillmore and Atkins 1992). Still, each verb highlights different aspects of the frame, by making different sets of argument roles obligatory to fill. With 'buy', the buyer and the object bought must be specified (linguistically or contextually), but the other argument roles can be left unfilled; with 'sell', it is the seller and the goods which must be specified. The buyer is optional, as the seller in the case of 'buy'. This provides an argument in favour of the view that 'buy' and 'sell' express *distinct* relations, if we agree with Devlin that "part of the complex structure that constitutes a 'relation' is the collection of conditions that determine which particular groups of argument roles need to be filled" (Devlin 1991: 121). Let us therefore consider that the Fillmorean frame evoked by both 'sell' and 'buy' denotes a complex relation R^* , while each of the two verbs denotes a distinct relation R , such that the set of argument roles of R is included in the set of argument roles of R^* .

When an argument role is left unfilled, as the *seller* role in 'John bought a house', the existence of an object playing that role is implicit since it is part of the semantic frame evoked by the verb. In such a case, following Fillmore, I will say that the relevant argument role is *backgrounded*. The argument roles which are filled are *foregrounded* (or 'profiled', in the terminology of Langacker 1987). Leaving an argument role unfilled is equivalent to existentially quantifying over its values. Thus 'John bought a house' is equivalent to 'John bought a house from someone, who sold it'.

'John has eaten' is equivalent to 'John has eaten something edible.' But it does not follow that 'John has eaten something' and 'John has eaten' say the same thing and involve the same relations. On the view I am defending, 'John has eaten' involves a 1-place relation (a property), while 'John has eaten something' involves a two-place relation, one of the arguments of which is unspecified.

Grammatical constructions contribute a lot to foregrounding and backgrounding. Consider (1):

(1) The house was sold

The commercial transaction frame is evoked, and it is evoked through the verb 'sell' which lexically foregrounds the *seller* role. Yet the passive construction has the well-known effect of backgrounding the argument role associated with the subject term in the active form — here the *seller* role which the verb normally highlights.⁴ The grammatical construction can also foreground some aspect of the frame. Thus in (2), the ditransitive construction foregrounds the *buyer* role back despite its being backgrounded by the verb 'sell':

(2) John sold Bill the house

We cannot say that the relation here is a two-place relation between a buyer and an object bought, on the grounds that these two roles are the roles which the verb 'sell' renders obligatory to fill. The *buyer* role, though it might have been left unfilled without ungrammaticality (had another construction been used), is foregrounded in (2) through the ditransitive construction. In that construction, therefore, 'sell' expresses a three-place relation between the seller, the buyer, and the goods — the same type of three-place relation which the verb 'give' expresses, except that the verb 'give' always (or nearly always) expresses a three-place relation, while 'sell' makes only two argument roles obligatory to fill.

The ditransitive construction turns the two-place relation denoted by the verb 'sell' into a three-place relation. Similarly, adverbial adjuncts turn n-place relations (e.g. 'John studies geography') into n+1 place relations ('John studies geography at UCLA'). In this framework, the inferential pattern which motivated Davidson's analysis of action

⁴ See Tesnière 1969 for pioneering remarks on these matters. (The passive is a 'recessive construction', according to Tesnière, i.e. the opposite of causative constructions which increase the number of argument role.) See also Goldberg 1995: 57-58 for a distinction between two forms of constructional backgrounding: *shading* and *cutting*.

sentences can be accounted for by appealing to the (fairly general) principle that, when two relations are related to each other by arity-decreasing/increasing operations such as these, the relation with greater arity entails the other one: R_{n+1} entails R_n . Thus we have the following inference pattern:

John studies geography at UCLA --> John studies geography
 John sold Bill the house ---> John sold the house
 John walks the dog ---> The dog walks
 John eats a banana ---> John eats
 John buttered the toast in the kitchen ---> John buttered the toast

There are exceptions, though. For example, it is not obvious that 'It's raining in Nevada' entails 'It's raining', at least if 'It's raining' is interpreted in the sense of 'It's raining here and now'. It may be raining in Nevada without there being rain here and now. In general, the R_{n+1} ---> R_n inference pattern becomes problematic whenever the right hand side of the entailment relation is an instance of 'definite null instantiation' (DNI) in Fillmore's sense. DNI is a special case of backgrounding where the backgrounded argument role is contextually filled instead of being unfilled (Fillmore 1986, Fillmore and Kay 1994). I will return to that topic in the next section.

It should be noticed that the 'walk' alternation above is significantly different from the other examples. First, the arguments are linked to different grammatical functions on the two sides of the alternation; in contrast the other alternations do not affect linking. More important for our purposes, the 'walk' example does not support an argument pattern which the other examples illustrate. While we have:

John studies geography --> John studies geography somewhere
 John sold the house ---> John sold the house to someone,
 John eats ——> John eats something,

we don't have:

The dog walks --> Someone walks the dog.

This fact shows that the two-place relation 'walk' in 'John walks the dog' does not belong to the WALK frame. Rather, it belongs to a more complex frame containing the WALK frame as a part. A relation R belongs to a frame Φ iff it can be obtained from the complex relation R^* denoted by Φ by decreasing the arity of R^* . R then corresponds to a particular aspect of the frame which is highlighted or focussed on. When that is the case, the argument roles of R^* which are not profiled (replicated as

argument roles in R) are still present in the background via the evoked frame. As those background roles are not filled, the general effect is that of quantifying over the values of the role: that explains the inference pattern above. We don't find this inference pattern in the 'walk' case because the argument role corresponding to the person walking the walker is not part of the WALK frame; not being part of the evoked frame, there is no reason why it should be implicit when we say 'The dog walks'.

§2.3 Relativisation

I have just mentioned cases in which the backgrounded argument role is left unfilled. In other cases, however, the backgrounded argument role *is* filled. This corresponds to what Perry calls 'unarticulated constituents' (Perry 1986). An argument is unarticulated when it corresponds to an argument role which no linguistic expression is used to fill, but which is nevertheless *contextually* filled. Fillmore and Kay (1994) refer to this process as 'Definite Null Instantiation' (DNI), as opposed to 'Indefinite Null Instantiation' (INI). Compare 'I have eaten' (INI) and 'I have noticed' (DNI). In both cases the object (the thing eaten or the thing noticed) is left implicit. In the case of 'I've noticed', however, the object noticed must be contextually identifiable for the utterance to be felicitous. This is not the case for 'I've eaten', as we have seen: 'I have eaten' can be understood to mean that the speaker has eaten something or other. This sort of reading is not available for 'I've noticed'.⁵

In Indefinite Null Instantiation, the relevant argument role is left unfilled; in Definite Null Instantiation, it is contextually filled. Shall we say — as I did above — that it is backgrounded (because it is left implicit), or that it is foregrounded (because it is filled)? It does not really matter what we say, as long as we make a clear distinction between *three* types of case:

Unfilled argument roles

vs.

Contextually filled argument roles

vs.

Linguistically filled argument roles

The operation on facts which Barwise calls 'relativisation' corresponds to the second case (Barwise 1989: 253-4). A fact which contains a certain constituent can be turned into a fact not containing that constituent, by appropriately decreasing the relation. The resulting fact is said to be 'relativised to' the constituent in question. This

⁵ The same distinction can be found in Prandi 1992: 48-49fn.

operation makes sense whenever the missing constituent is contextually provided or 'presupposed'.

Relativisation is pervasive in natural language. More often than not, when we state a fact, some constituent of the fact is contextually given and somehow taken for granted. Instead of expressing the more complex fact containing that constituent, what we linguistically express is the simpler fact relativised to that constituent. Barwise gives perspectives as an example. When Holmes says 'The salt is left of the pepper', he expresses a relativised fact:

Left-of (salt, pepper)

This fact holds with respect to Holmes's perspective, but that perspective is not a constituent of the (relativised) fact. Rather, it is an aspect of the situation talked about.

Perry's distinction between 'concerning' and 'being about' (Perry 1986) is useful here. A fact is about something iff that thing is a constituent of the fact. The fact

Left-of (salt, pepper, Holmes's perspective)

is 'about' Holmes perspective in that sense. But Holmes's own words, 'The salt is left of the pepper', do not express that fact; they express the *relativised* fact that the salt is left of the pepper. The relativised fact 'concerns' Holmes's perspective, but it is not 'about' it.

In Barwise's example, the perspective need not be linguistically articulated because it is given in the context. Similarly in Perry's own example, 'It is raining'. The person who says 'It's raining' talks about a certain location and expresses a relativised fact concerning that location: the fact that it's raining. That fact is distinct from the (apparently) unrelativised fact that *it is raining at that location*. It is felicitous to express the latter type of fact only if the location in question is contrasted with other possible places where it might rain (or not rain). If no contrast is envisioned and only that location is concerned, there is no need to articulate it by expressing the more complex fact.

It is in terms of this notion of relativised fact that I propose that we define the 'support' relation. I said that all the facts in the factual set of a situation contain that situation as a constituent:

$$(\sigma) (s) (\sigma \in w(s) \supset s \in \sigma)$$

Now a fact σ belonging to a situation s can be relativised to that situation. The relativised fact σ_r holds with respect to s — it 'concerns' s — even if it is not part of

$w(s)$ (it does not contain s as a constituent and thus is not 'about' it). I interpret Austinian semantics as the claim that natural language utterances express relativised facts of this sort. That is why I build relativisation into my definition of the 'support' relation. Corresponding to $w(s)$, the factual set of the situation s , there is $w(s)_r$, the set of all the facts of $w(s)$ relativised to s . A situation s supports a fact σ iff $\sigma \in w(s)_r$.

III. δ -STRUCTURES

§3.1 Mentioning situations

In §2.3 I claimed that the fact linguistically expressed by an utterance is relativised to the situation talked about. But that does not seem always to be the case. Instead of merely saying 'It is raining', I can say: 'It is raining here' — thereby making the location I am talking about explicit. Similarly, Holmes could say: 'The salt is left of the pepper *from my perspective*', thereby making the perspective explicit. As the situation talked about is made explicit in the utterance itself, it becomes a constituent of the fact linguistically expressed. It follows that the fact in question cannot be relativised to that situation, for a fact cannot be relativised to α if it contains α as a constituent.⁶ By building relativisation into our definition of ' \models ', do we not run the risk of depriving ourselves of the means of accounting for such cases?

We do not. We have no trouble accounting for 'explicit' cases in the framework set up so far, in which only relativised facts are taken to be linguistically expressed. In explicit cases like 'It's raining here' or 'The salt is left of the pepper from Holmes's perspective', the fact is not relativised *to the situation which is explicitly mentioned in the utterance*, but it is relativised nonetheless — to a different situation.

Let us contrast Holmes's utterance with Watson's:

Holmes: - The salt is left of the pepper

Watson: - The salt is left of the pepper from Holmes's perspective

The first utterance presupposes Holmes's perspective and expresses a fact relativised to that perspective. The second utterance makes that perspective explicit by integrating it as a constituent of the fact linguistically expressed. The latter is not relativised to Holmes's perspective, but it is not unrelativised either: it is relativised to a distinct situation, namely the situation *Watson* is talking about. The situation Watson talks

⁶ See §5.1 for a qualification.

about contains Holmes, the salt and the pepper. That is not the same situation as the situation Holmes was talking about in the first place, for the latter contained only the pepper and the salt (it did not contain Holmes).⁷ In general, when we make some aspect of the earlier situation explicit by incorporating it into the fact linguistically expressed, the utterance demands a broader situation of reference than the implicit utterance did: it demands (= is felicitous only with respect to) a situation of reference containing the new constituent in its domain.⁸

On the view I am presenting, the fact linguistically expressed by an utterance is relativised to the situation talked about, and there is a principled difference between the situation talked about and the 'mentioned situation' which is a constituent of the linguistically expressed fact. The situation talked about must conform to a specific constraint concerning the size of its domain. The domain of a situation contains the constituents of all the facts in its relativised factual set. Now to assert something is to claim that the fact the utterance expresses belongs to the relativised factual set of the situation talked about. It follows that, when a fact is linguistically expressed, the situation it concerns — that which the speaker is talking about — must contain all the constituents of the expressed fact in its domain, on pains of pragmatic contradiction. Since the mentioned situation is a constituent of the linguistically expressed fact, the situation of reference s' must contain the mentioned situation s in its domain. The mentioned situation is subject to no such constraint: it need not contain itself in its domain.

Consider the 'It is raining' example again. Suppose John is in Condé sur Huisne, and he says: 'It is raining'. His utterance concerns Condé sur Huisne and is true if and only if it's raining in Condé sur Huisne. As Perry emphasized (Perry 1986:216), when we say or think 'It's raining', we need not think *reflectively* about the place we're in, even though that place enters the truth-conditions of what we say. We let the place we are in complete the content of our thought, instead of representing it explicitly. The complete content of John's utterance can therefore be represented as:

Condé sur Huisne $\models \langle\langle \text{It's raining} \rangle\rangle$

where Condé sur Huisne is on the situational side, rather than on the side of linguistically expressed content.

⁷ Well, perhaps it did contain Holmes after all. The situation is complicated by the existence of reflexive situations. See §5.1 below and Recanati 1997, where I suggest that perspectives are reflexive situations.

⁸ Situations, like worlds, have domains. The domain of a situation s (with respect to a world w) contains all the entities (situations) which are constituents of the facts in $w(s)_r$.

If John reflectively thinks of the location he is in, instead of letting it be provided by the context, he will say something different, namely 'It is raining in Condé sur Huisne'. The utterance is now *about* Condé sur Huisne. One difference with the previous utterance is that the situation talked about *must* contain Condé sur Huisne as a constituent, while this was not necessary for the simpler utterance 'It's raining'. Anticipating somewhat, we can represent the content of 'It's raining in Condé sur Huisne' as:

$$s' \models \langle\langle \text{Condé sur Huisne} \models \langle\langle \text{It's raining} \rangle\rangle \rangle\rangle$$

where the situation s' must contain Condé sur Huisne in its domain. For example, we can imagine that John is talking about the meteorological situation in various places, such as Condé sur Huisne and Cajarc. He says: 'In Cajarc, the sun is shining, but it's raining in Condé sur Huisne'. Here the situation s' John is talking about contains both Cajarc and Condé sur Huisne in its domain. It is *that* situation which is said to support 'It's raining in Condé sur Huisne'.

At this point, it is worth introducing a technical term instead of using a number of equivalent expressions such as 'the situation talked about', 'the situation of reference', or 'the situation with respect to which the utterance is interpreted'. From now on I will use the technical term 'exercised situation' instead of 'situation talked about' and its cognates. In formulas, I will enclose the expression standing for the exercised situation in square brackets. Thus the Austinian proposition expressed by an utterance will be represented as:

$$[s'] \models \sigma$$

The special case in which a situation is linguistically mentioned and is a constituent of the expressed fact will be represented by iterating the Austinian formula:

$$[s'] \models \langle\langle s \models \sigma \rangle\rangle$$

Such an iterated Austinian proposition I call a δ -structure. It contains two situations: s' , the exercised situation, and s , the mentioned situation; two 'support' relations: one between the exercised situation and the linguistically expressed fact, the other internal to the linguistically expressed fact; and two facts: the fact that $s \models \sigma$ and the fact that σ .

§3.2 δ -structures in English

In my version of Austinian semantics, the content of most utterances in ordinary English has the structure:

$$[s'] \models \langle\langle s \models \sigma \rangle\rangle$$

That is so because utterances are tensed and the tenses I treat as 'mentioning' temporal situations (much as temporal adverbials do). Thus utterances of the form

(At t) it will be the case that p

and

(At t) it was the case that p

are instances of the schema:

$$[s'] \models \langle\langle l_t \models \sigma \rangle\rangle$$

where s' is the exercised situation, and l_t , the mentioned situation, is a (past or future) temporal location. Concrete examples are:

Yesterday Paul attended the meeting

I saw the accident on TV

When he arrived, Peter was upset

You will break you neck!

Jane is feeding the rabbits

Similar examples involving spatial locations are easy to come by:

In Spain, the government is appointed by the King

Three miles from here, there is an architectural treasure

Those examples I analyse as

$$[s'] \models \langle\langle l_s \models \sigma \rangle\rangle$$

where I_s is the location denoted by 'in Spain' or 'Three miles from here'. (In the first case, the location is geopolitical rather than spatial in a strict sense.) It is an interesting fact that the reference to temporal locations is built into the grammar of the language via the tense feature, while spatial locations are in most cases introduced into the linguistically expressed fact via adverbial adjuncts. But the expressed fact has, in both cases, the same structure, and that is all that matters here.

Beside locations, any type of situation (in the technical sense, that is, any type of entity) can be mentioned in a δ -structure. That is true in particular of individual objects, insofar as they can be treated as situations in the technical sense (§2.1). In the present framework, the content of a simple subject-predicate sentence such as 'John is British' is also a δ -structure:

$$[s'] \models \langle \langle \text{John}^0, \text{British}^0, + \rangle \rangle$$

The subject term contributes the mentioned situation (viz. John), and the predicate contributes a 'thetic fact' which that situation is said to support (viz. the fact of being British). Let me explain.

The *factual set* of an individual object is the set of facts involving that object as constituent, while the *relativised factual set* consists of all the facts in the factual set relativised to the object itself. This means that, whatever relation R^n is a constituent of a fact in the factual set of a , the corresponding fact in its relativised factual set will be built around the relation R^{n-1} , with a itself disappearing from the sequence of arguments. The relativised factual set is the set of all such relativised facts. Now consider an individual object, John, and the property of being British. This gives us a fact: $\langle \langle \text{British}^1, \langle \text{John} \rangle, + \rangle \rangle$, a fact which belongs to the factual set of John. The monadic relation British^1 is a constituent of that fact, together with the individual John. If we relativise that fact to John, we obtain a fact with an 0-adic relation and no argument: the fact $\langle \langle \text{British}^0, + \rangle \rangle$. The relativised factual set of John thus contains a number of such 0-adic facts, facts relativised to John himself, in the same way in which the relativised factual set of a spatio-temporal location can contain the 0-adic fact $\langle \langle \text{Rain}^0, + \rangle \rangle$. (On the 'rain' predicate as 0-adic, see Barwise and Perry 1983 pp. 50-51, and Fillmore and Kay 1994 §4.3).⁹

0-adic facts can sometimes be expressed directly, as when we say: 'Superb!' about a performance we are watching. An Austinian proposition is expressed, in which the performance is the exercised situation and the nucleus is the 0-adic fact $\langle \langle \text{Superb}^0, + \rangle \rangle$.

⁹ Of course, the relativised factual set of John will also contain monadic facts such as $\langle \langle \text{Married-to}^1, \langle \text{Josephine} \rangle, + \rangle \rangle$, and other n-adic facts derivable through relativisation from n+1-adic facts in the factual set.

+>>. This corresponds to what grammarians in the Brentanian tradition called a thetic judgement, and to what others in the Saussurean tradition called a 'monoreme'. Thetic judgements were said to be simple, in contrast to categoric judgements such as that expressed by 'John is British' (or 'This performance is superb'). Categoric judgements are 'double' in that they involve two ancillary acts: (i) the identification of some entity (or entities), and (ii) the predication of some property or relation holding of that entity (or sequence of entities). Since the second ancillary act is common to the thetic judgement and the categoric judgement, we can analyse categoric judgements as denoting complex facts with two components: a thetic fact, plus a situation which supports it.

$$\begin{aligned} &<<\text{British}^1, \langle \text{John} \rangle, +>> \\ &= \\ &<<\text{John} \models \langle \langle \text{British}^0, +>> \rangle \rangle \end{aligned}$$

On this analysis there is a semantic contrast between the thetic utterance 'Superb!', *concerning* the performance, and the categoric utterance 'This performance is superb', in which the performance is *mentioned* in the utterance itself:

$$\begin{aligned} \text{Thetic:} & \quad [\text{the performance}] \models \langle \langle \text{Superb}^0, +>> \\ \text{Categoric:} & \quad [s'] \models \langle \langle \text{the performance} \models \langle \langle \text{Superb}^0, +>> \rangle \rangle \end{aligned}$$

This contrast is the same as that between 'It's raining' and 'It's raining here' (§3.1).

Traditionally, existential sentences ('There are spies') and weather sentences ('It's raining') were considered as expressing thetic judgments. Since those sentences are tensed, the claim that they express thetic judgements is somewhat controversial. A better example is provided by children's one-word utterances, e.g. 'Rain!'. In this type of utterance, some contextually provided topic is globally characterized by means of a 0-adic predicate (Lyons 1975, Sechehaye 1926). Except for this type of case, English sentences express δ -structures of various degrees of complexity — i.e. not only simple δ -structures, but also *iterated δ -structures* conforming to the following pattern:

$$[s_n] \models \langle \langle s_{n-1} \models \langle \langle s_{n-2} \models \dots \models \sigma \dots \rangle \rangle \rangle$$

Iterated δ -structures abound in ordinary discourse. A simple example in which the only mentioned situations are locations is:

Two miles from here, a huge building will be erected

The structure of this representation is

$$[s_3] \models \ll s_2 \models \ll s_1 \models \sigma \gg \gg$$

where σ is the fact that a huge building is erected, s_1 is the temporal situation indicated by the future tense, s_2 is the spatial situation denoted by the phrase 'two miles from here', and s_3 is the exercised situation. If what I said above is correct, a grammatically simpler sentence such as 'Napoleon was ambitious' can be construed as expressing an iterated δ -structure of the same degree of complexity:

$$[s] \models \ll I_t \models \ll \text{Napoleon} \models \ll \text{Ambitious}^0, + \gg \gg \gg$$

More complex δ -structures are far from uncommon. A good example (from Recanati 1997) is the English sentence 'The landlord thought that, in 1996, Peter would be penniless', which involves five distinct situations.

δ -structures can also be quantified. In a *quantified δ -structure*, no specific situation is mentioned, but a general fact is linguistically expressed. We have:

$$[s'] \models \ll Qs: s \models \sigma \gg$$

where Q can be any (possibly restricted) quantifier. For example the following utterances express universally quantified δ -structures:

- (1) Each time he comes, John asks for a break
- (2) John always succeeds
- (3) Everywhere in France, there are communal swimming pools

(1) and (2) quantify over temporal situations, (3) over spatial situations. Except in (2), the quantifiers are explicitly restricted: (3) mentions all spatial locations within the French territory, and (1) all temporal locations supporting: $\ll \text{John comes} \gg$.

Finally, we must pay attention to an important difference between the exercised situation and the mentioned situation. The exercised situation is referred to, hence it must exist in the actual world.¹⁰ But the mentioned situation need not be actual; it can be 'virtual', in a sense shortly to be explained. When the mentioned situation is virtual, the δ -structure is an *intentional δ -structure*.

¹⁰ The cases in which a non-actual situation is used as exercised situation are cases of simulation, in which we do *as if* the situation was actual.

§3.3 Intentional δ -structures

The most obvious example of intentional δ -structure is provided by conditionals. Conditionals mention *hypothetical* situations and state facts concerning them. The content of an utterance like

(4) If he opens the fridge, John will be scared to death

can be analysed as:

$$[s'] \models \langle\langle s \models \langle\langle \text{John is scared to death} \rangle\rangle \rangle\rangle$$

where s' is the actual situation in which the conditional holds, and s is an hypothetical situation in which John opens the fridge. By analysing conditionals as δ -structures we capture the commonality between (4) and (5):

(5) When he opens the fridge, John will be scared to death

(5) is a temporal δ -structure similar to those mentioned above. The difference between (4) and (5) is that s , the mentioned situation, is a virtual situation in (4) and an actual temporal situation in (5).¹¹

In earlier writings (Recanati 1996, 1997) I analysed meta-representations — representations about representations — on the same pattern. For example, 'In the film, Robin Hood meets Frankenstein' was analysed as mentioning a *fictional* situation (that depicted by the film) and saying that it supports $\langle\langle \text{Robin Hood meets Frankenstein} \rangle\rangle$. In the same way, I analysed 'John believes that p ' or 'According to John, p ' as a δ -structure in which the mentioned situation is John's belief-world (the complex situation corresponding to John's beliefs, that is, the world as it is according to John):

$$[s'] \models \langle\langle \text{Bel}_{\text{John}} \models \sigma \rangle\rangle$$

('Bel_{John}' is my name for John's belief-world.) In general, I analysed the content of a meta-representation as a δ -structure $[s'] \models \langle\langle s \models \sigma \rangle\rangle$, where s , the mentioned situation, is the situation described by some representation given in the exercised situation. The representation can be a film, a book, a picture, an utterance, or a person's mental states

¹¹ I will characterize 'virtual situations' properly in the next section. For the time being I rely on an intuitive understanding of what it is for a situation to be virtual rather than actual.

(to mention only the most salient types of case). In the examples above, the relevant representations are John's belief state (which represents the world as being thus and so) and the film (which represents a certain complex situation). Other examples include:

In the picture, John is smiling
 John said that Peter would come

I now see several problems with my earlier analysis of meta-representations. First, it is pretty clear that the representation (the picture, the film, John's mental states, etc.) is actually mentioned in the meta-representation; it cannot be confined to the exercised situation which the meta-representation concerns. We must therefore introduce the representation itself into the fact denoted by the meta-representation. The general structure becomes:

$$(6) \quad [s'] \models \langle \langle R \models s \models \sigma \rangle \rangle$$

where s' is the exercised situation, R is the representation, and s is the situation the representation describes.

Once we have taken this step we face the Compositionality problem: how is it possible that a representation such as 'In the picture, p ' expresses a structure like (6), in which two situations are mentioned (the picture R , and the situation s it represents), while a very similar representation, 'In the kitchen, p ' expresses a much simpler structure (in which only one situation — the kitchen — is mentioned)? I think this particular problem could be solved by appealing to Nunberg's theory of meaning transfer (Nunberg 1995).¹² However there is another problem which cannot be solved so easily.

On my earlier analysis, as well as on the revised version corresponding to formula (6), a meta-representation mentions the virtual situation described by the representation R : it is that situation which is said to support a particular fact σ . Thus 'John believes that p ' mentions John's belief world and says that it supports the fact that p ; 'in the film, p ' mentions the fictional situation described by the film and says that it

¹² Briefly: While 'my car is parked out back' predicates a certain property of the car (the property of being parked out back), 'I am parked out back' predicates a different (and more complex) property of a different subject: it ascribes to the speaker the property of 'having a car which is parked out back'. The predicative expression 'is parked out back' thus undergoes a transfer of meaning (Nunberg 1995). In virtue of the same type of process, 'In the picture John is smiling' could be considered as ascribing to the picture a more complex property than 'In the kitchen, John is smiling' ascribes to the kitchen.

supports the fact that p ; etc. The major problem raised by this theory is this: what is 'the situation described by the representation'? Is it really a situation (e.g. something to which a world associates a factual set) or is it a complete world? It seems clear that John's 'belief world', for example, *is* a world. Similarly, a work of fiction describes a world, rather than merely some particular situation. There is a clear contrast with conditionals in this respect. While the antecedent of a conditional serves to identify the mentioned situation, in meta-representational sentences the sole function of the prefix arguably is to shift the world, without indicating any *particular* situation as supporting the fact expressed by the accompanying sentence. When we say 'John believes that kangaroos have tails', we present the fact that kangaroos have tails as holding in John's belief-world, but no particular situation is singled out, in contrast to what happens with conditionals.

Faced with this particular difficulty, we can do several things. We can treat meta-representational sentences such as 'John believes that kangaroos have tails' as expressing *existentially quantified* intentional δ -structures:

$$[s'] \models_{@} \langle\langle \exists s s \models_{\text{bel}_{\text{John}}} \langle\langle \text{kangaroos have tails} \rangle\rangle \rangle\rangle$$

On that view John's belief-world is a world, not a situation; it indexes the 'support' relation between the mentioned situation and the fact that kangaroos have tails. The mentioned situation in question is left indefinite instead of being specified — thus we account for the above-noted contrast with conditionals.

Alternatively, we could define, for each world, a 'maximal situation' (that very world, construed as a situation), and say that it is that maximal situation which the meta-representational prefix denotes. That would make the general picture much more complex than it is; for we would have to say either that a world can assign a factual set to another world, or that worlds are self-interpreting, in contrast to partial situations.

There is a third option, which enables us to solve several of our problems at once. We can *give up* the view that the meta-representational prefix (e.g. 'In the film', 'John believes that', etc.) mentions a virtual situation, and construe it instead as doing two things: (i) it mentions a *real* situation — the representation R itself (the film, or John's belief state); and (ii) it shifts the world with respect to which what follows the prefix is interpreted. It is this (very tentative) suggestion which I will pursue in what follows. Appearances notwithstanding, this analysis is compatible with the claim that meta-representations are intentional δ -structures, in which a virtual situation is mentioned. What we give up is only the claim that the virtual situation in question is contributed by the prefix itself.

§3.4 World-shifting

When I say that hypothetical or fictional situations are 'virtual', do I mean that they are not in the domain of @? Certainly not. Suppose I utter (7):

(7) If John went to Lyons, he took the 1:30 train

I am considering a hypothetical situation, viz. a situation in which John went to Lyons. Suppose that John indeed went to Lyons. Then the hypothetical situation turns out to be in the domain of @ (in the sense that @ associates a factual set with that situation). Yet it still is a hypothetical (hence a virtual) situation. So we cannot define a virtual situation as a situation which does not belong to the domain of @, the actual world.

Situations can be in the domain of several worlds; a situation does not essentially belong to the domain of this or that world. But when we mention a situation in a δ -structure, we mention it as supporting a certain fact. Now the 'support' relation is relative to a world: a situation s supports a fact σ iff that fact belongs to $w(s)_r$, the relativised factual set of the situation *as determined by a particular world function* w . I therefore suggest the following (tentative) definition:

(B)

A situation s mentioned in a δ -structure is virtual (hence the δ -structure itself is intentional) iff s is presented as supporting a certain fact *with respect to a world w different from @* (or at least, different from the world with respect to which the global δ -structure is evaluated¹³). That fact itself I call an 'intentional fact'.

At the end of §3.1 I said that a δ -structure involves two situations, two facts, and two 'support' relations. According to (B), the two facts which we find in an *intentional* δ -structure are not only presented as supported by different situations (the exercised situation and the mentioned situation), that is, as belonging to their respective relativised factual sets, but the relativised factual sets in question are themselves determined by different world functions:

$$[s'] \models_{@} \langle\langle s \models_w \sigma \rangle\rangle$$

In this formula the two 'support' relations bear different indices. This means that the relativised factual sets they implicitly refer to (via the equivalence between $s \models_w \sigma$ and $\sigma \in w(s)_r$) are determined by different world functions, @ and w . In non-intentional δ -

¹³ This qualification is needed to deal with the cases of 'projection' in which some previously mentioned situation is used as exercised situation (Recanati 1997).

structures the same world-function is appealed to throughout and need not be explicitly represented.

This analysis works well with conditionals. In (7), the situation s which is mentioned is a hypothetical situation in which John went to Lyons. To say that it is hypothetical is not to say that the actual world associates no factual set with that situation; as we have seen, it may be a fact that John went to Lyons (indicative conditionals are used precisely when the hypothetical situation's being actual is a live option). What makes the situation hypothetical is the world with respect to which that situation is said to support the fact that John took the 1:30 train: that world is not the actual world, but an hypothetical world *hyp* which may be more or less similar to $@$. The utterance can therefore be analysed as:

$$[s'] \models_{@} \ll s \models_{\text{hyp}} \ll \text{John took the 1:30 train} \gg \gg$$

where s , the mentioned situation, is characterized as temporally past and as supporting the fact that John went to Lyons. In a conditional sentence the consequent expresses the fact which is said to hold in the hypothetical situation, while the antecedent characterizes that situation as supporting a 'restricting' fact, a fact which serves to identify the hypothetical situation in question (in our example, the restricting fact is the fact that John went to Lyons).

The word 'if' in a conditional I take to be a *world-shifter*: it indicates that the facts in its scope are evaluated with respect to an hypothetical world distinct from the world with respect to which the global δ -structure is evaluated. The well-known context-sensitivity of conditionals comes from the fact that the world in question can be fleshed out differently in different contexts.

Meta-representational prefixes such as 'In John's mind', 'in the picture', 'in the film', 'according to Paul', 'John believes that', 'John said that', and so forth, are intuitively world-shifters, like the word 'if'. Consider, again, the Robin Hood example:

In the film, Robin Hood meets Frankenstein

The meta-representational prefix 'In the film' indicates that the fact expressed by the accompanying sentence ('Robin Hood meets Frankenstein') holds in a fictional universe, in contrast to the fact expressed by the global meta-representation (i.e. the fact that *in the film* Robin Hood meets Frankenstein): the latter holds in the actual world.

On this analysis, meta-representational prefixes such as 'In John's mind', 'in the picture', 'in the film', 'according to Paul', 'John believes that', 'John said that', and so forth, mention a real situation (Paul's mental states, the picture, the film, etc.) as supporting a certain fact, but the fact in question is such that *whichever 'support'*

relation it internally involves is indexed to the world of the representation. For example, the utterance

John believes that, in the eighteenth century, kangaroos had tails

would be analysed as denoting the following fact:

$$\ll \text{John's mental state} \models_{@} \ll 18^{\text{th}} \models_{\text{bel}_{\text{John}}} \ll \text{kangaroos have tails} \gg \gg \gg$$

where '18th' is the temporal location denoted by the phrase 'in the eighteenth century' (together with the past tense). Now that situation is 'virtual' because it is in some alternative world, namely John's belief-world, that it supports the fact that kangaroos have tails. So the meta-representational prefix, *qua* world-shifter, 'virtualises' the situations which the accompanying sentence mentions.

In that framework the complete content of a meta-representation such as 'In the film, *p*' or 'John believes that *p*' has the following structure:

$$[s'] \models_{w_1} \ll R \models_{w_1} \ll s \models_{w_2} \sigma \gg \gg$$

where *s*' is the exercised situation, *R* is the mentioned representation (e.g. the film, or John's mental state), and $\ll s \models_{w_2} \sigma \gg$ is the fact denoted by the sentence *p* which the meta-representational prefix introduces (e.g. 'Robin Hood meets Frankenstein', or 'In the eighteenth century kangaroos had tails'). In contrast, the complete content of a conditional retains the simpler structure:

$$[s'] \models_{w_1} \ll s \models_{w_2} \sigma \gg$$

where *s*' is the exercised situation, *s* is the situation mentioned by the antecedent, and σ is the fact denoted by the consequent. On that view, although the content of both conditionals and meta-representations is an intentional δ -structure, characterized by a world-shift, there is a significant difference between them. In meta-representations, the world-shift itself is represented, that is, it is internal to the fact denoted by the meta-representation; while in conditionals the world-shift takes place without being represented. The fact denoted by the conditional is the fact that $s \models_{w_2} \sigma$, while the fact denoted by a meta-representation is the fact that $R \models_{w_1} \ll s \models_{w_2} \sigma \gg$.

That consequence of the analysis is desirable. For, as I said earlier, it is clear that the representation (the picture, the film, John's mental states, etc.) is actually mentioned in the meta-representation; it cannot be confined to the exercised situation

which the meta-representation concerns. Yet the analysis of meta-representations which has just been sketched has another, less desirable consequence. A meta-representation is said to mention a real situation (viz. a representation: film, picture, etc.) as supporting an *intentional fact*: a fact such that the 'support' relation it internally involves is indexed to the world of the representation rather than to the world with respect to which the global δ -structure (the meta-representation) is evaluated. But aren't there simple facts which internally involve no 'support' relation? Arguably there are; they are the 'thetic facts' mentioned in §3.2 (e.g. the fact that it is raining). Assuming that there are such facts, it follows from the above analysis that they cannot be intentional: *they cannot be what the complement sentence in a meta-representation expresses*.

That consequence is surprising, but empirically correct as far as I can tell. If I say 'John believes that it is raining', I ascribe to John the belief that it is raining *in a particular place* (tacitly referred to). So the fact expressed by the complement sentence is not simple (thetic), appearances notwithstanding. This is less clear in an example like 'In the picture, it is raining', but we could perhaps treat that as an instance of quantified δ -structure.¹⁴ More generally, we could admit that the meta-representational prefix is a world-shifting operator requiring a complex fact to operate on, and handle the counterexamples by arguing as follows. Whenever the fact expressed by the accompanying sentence is not (already) complex, it is made so: thus if a sentence S expresses a simple (thetic) fact σ , embedding it within a meta-representational frame has the result that it expresses the complex fact that: $\exists s s \models \sigma$. The world-shifting operator can then operate on *that* fact, indexing the support relation it internally involves to the world of the representation.

IV. COMPLEX AND SCHEMATIC FACTS

§4.1 Complex facts

In §2.1 I represented facts as triples consisting of an n -place relation R^n , a sequence of n appropriate arguments, and a polarity. It is in terms of such facts that I defined the 'support' relation: a situation supports a fact iff that fact belongs to the relativised factual set of the situation. But what about *complex* facts in which simple facts enter as constituents? What is it for a situation to support a complex fact? This is the issue I will address in this section.

¹⁴ Alternatively, we could appeal to what I call the 'Meinongian pretense' to account for that example: see Recanati forthcoming, §5.5.

The first thing we must do is distinguish various forms of complexity. One form of complexity I call δ -complexity. A fact is δ -complex iff its constituents are a situation, the 'support' relation, and a fact which itself can be δ -simple or δ -complex. Adverbial adjuncts are δ -complexifiers in that they mention a situation in which the fact expressed by the remainder of the sentence is said to hold. In §3.2, I argued that tenses themselves are δ -complexifiers. It may be possible to treat moods also as δ -complexifiers, but I will not address this question here.¹⁵ More important for our purposes is the claim I made in the same section, to the effect that simple subject-predicate sentences themselves can be construed as expressing δ -complex facts. Categorical judgements, I said, can be represented as *δ -complexifications* ofthetic judgements:

Thetic: $[s] \models \langle\langle R^0, + \rangle\rangle$
 Categorical: $[s'] \models \langle\langle s \models \langle\langle R^0, + \rangle\rangle \rangle$

On this analysis a 'thetic fact' is expressed both in thetic and categorical judgements. In categorical judgements the fact in question is not expressed directly by the utterance, but indirectly, as a constituent of a δ -complex fact. In thetic judgements it is directly expressed.

What is it for a situation to support a δ -complex fact? There is an easy answer to that question. A situation s' supports a δ -complex fact $\langle\langle s \models \langle\langle R^n, \langle a_1, \dots, a_n \rangle, i \rangle\rangle \rangle$ if and only if it supports the corresponding 'unrelativised' fact $\langle\langle R^{n+1}, \langle a_1, \dots, a_n, s \rangle, i \rangle\rangle$. There is a complication if the supported fact is intentional, for that forces us to keep track of the world-indices. Still, the equivalence holds.

There is another form of complexity, which I call it ' \oplus -complexity'. ' \oplus ' stands for any connective such as \vee or $\&$ (but not for negation which affects the polarity of the fact, not its complexity). What is it for a situation to support e.g. a conjunctive fact or a disjunctive fact? In situation theory (e.g. Barwise 1989, Devlin 1991) \oplus -complexity is accounted for by extending the definition of 'support' as follows:

$$s \models \sigma \oplus \sigma' \text{ iff } s \models \sigma \oplus s \models \sigma'$$

For example:

¹⁵ Arguably, moods indicate the world which indexes the 'support' relation in the Austinian proposition. Thus 'Close the door' says that whatever situation is at issue supports the fact of your closing the door, where the 'support' relation is relative to your wish-world. On this analysis moods are *not* δ -complexifiers.

$$\begin{aligned}
s &\models \sigma \ \& \ \sigma' \text{ iff } s \models \sigma \ \& \ s \models \sigma' \\
s &\models \sigma \vee \sigma' \text{ iff } s \models \sigma \vee s \models \sigma' \\
s &\models \sigma \supset \sigma' \text{ iff } s \models \sigma \supset s \models \sigma'
\end{aligned}$$

A third form of complexity involves quantification. What is it for a situation to support a quantificational fact such as the fact that all men are mortal? There are two approaches to this problem in the present framework.

1. Quantificational facts can be treated as having the same global structure as simple (non quantificational) facts: $\langle\langle R^n, \langle a_1, \dots, a_n \rangle, i \rangle\rangle$. The difference between quantificational facts and simple facts is simply this: In quantificational facts the relation is a *second-order* relation and its arguments are themselves first-order relations. If such second-order facts are allowed into the relativised factual sets of situations, we do not have to extend the definition of 'support' to account for quantification.

2. Alternatively, we can appeal to the same method we use for \oplus -complexity, and extend the definition of 'support' as follows:

$$(N) \quad s \models \langle\langle Qx \sigma(x) \rangle\rangle \text{ iff } (Qx) (s \models \sigma(x))$$

where 'Q' is an arbitrary quantifier, ' $\sigma()$ ' is a schematic fact, and ' $Qx \sigma(x)$ ' is a quantificational fact. It is this approach which I will pursue in the next section.

§4.2 Schematic and quantificational facts

Schematisation is a process much like relativisation. *Relativising* a fact $\sigma = \langle\langle R^n, \langle a_1, \dots, a_n \rangle, i \rangle\rangle$ to one of its constituent arguments consists in backgrounding the argument role corresponding to that argument and generating a fact $\sigma' = \langle\langle R^{n-1}, \langle a_1, \dots, a_{n-1} \rangle, i \rangle\rangle$. *Schematizing* a fact σ with respect to one of its constituent arguments consists in 'parametrizing' the argument role corresponding to that argument: The arity of the relation is preserved (in contrast to what happens with relativisation), but the relevant argument role is filled only by a place-holder (a 'parameter'). While a relativised fact is a complete fact of arity $n-1$, a schematic fact is an *incomplete* fact of arity n , with one of the argument roles filled by a place-holder.

Being incomplete a schematic fact is not a fact (hence it does not belong to the relativised factual set of a situation), but a function from appropriate values for the argument role to non-parametrized facts. While $\sigma()$ is a schematic fact (hence not a

genuine fact), $\sigma(a)$, the value of the schematic fact $\sigma()$ for argument a , is a genuine fact. We are therefore in a position to understand the right hand side in formula (N): ' $(Qx) (s \models \sigma(x))$ ', where Q is an arbitrary quantifier. It says that for Q entity x , the situation s supports the fact which is the value of $\sigma()$ for argument x . Various types of *quantificational facts* can thus be defined through the support-conditions enunciated in formula (N). The following definitions are particular instances of (N):

$$s \models \langle\langle \exists x \sigma(x) \rangle\rangle \text{ iff } (\exists x) (s \models \sigma(x))$$

$$s \models \langle\langle \forall x \sigma(x) \rangle\rangle \text{ iff } (\forall x) (s \models \sigma(x))$$

I will henceforth represent a quantificational fact as $\sigma(Q)$, where $\sigma()$ is a schematic fact and Q is a quantifier ranging over the values of the parametrized argument role in $\sigma()$.

Schematic facts themselves can be schematized. When that is so, the resulting fact is *doubly schematic*. I represent doubly schematic facts thus:

$$(1) \quad \sigma() ()$$

This means that two of the argument roles of the relation $R^{n \geq 2}$ around which σ is built have been parametrized.

Quantificational facts themselves can be schematized:

$$(2) \quad \sigma(Q) ()$$

A schematized quantificational fact like (2) is a doubly schematic fact like (1) in which one of the two parametrized argument roles is 'bound' by the quantifier, the other argument role remaining 'free'. The free argument role in a schematized quantificational fact like (2) can itself be bound by a new quantifier; for example we can have:

$$(\exists x) (\sigma(Q) (x))$$

or any instance of

$$(Q'x) (\sigma(Q) (x))$$

where Q' is an arbitrary quantifier. Such a *doubly quantified fact* can be represented as follows:

$\sigma(Q) (Q')$

Of course, the process can be repeated. A doubly quantified fact can be schematized by parametrizing a third argument role (if the relation around which the fact is built has arity $n \geq 3$):

$\sigma(Q) (Q') ()$

This, in turn, can yield an even more complex quantificational fact:

$\sigma(Q) (Q') (Q'')$

A complex quantificational fact involves (i) a sequence $\langle Q_1, \dots, Q_n \rangle$ of quantifiers, and (ii) a schematic fact with n parametrized argument roles. I have represented the binding relation between quantifiers and argument roles by putting the quantifier in the relevant argument place. In English that relation is established through the lexical phenomenon of 'linking'. In the semantics of relational expressions (e.g. verbs) different argument roles for the denoted relation are linked to different grammatical functions such as 'subject', 'object' and 'indirect object'. Now quantificational expressions, that is, words or phrases denoting quantifiers, are themselves associated with specific grammatical functions, viz. the functions they fulfil or occupy within the sentence in which they occur. Hence a simple solution to the problem of determining which quantifier binds which argument role:

BTL ('binding through linking') Principle:

The quantifier Q contributed by a token θ of a quantificational expression *Quant* binds the argument role linked to the grammatical function occupied by θ .

For example, consider the sentence 'John sold something to Bill'. The verb in this sentence expresses a three-place relation *sell*³ whose argument roles are borrowed from the commercial transaction frame. The three foregrounded argument roles in this construction are: *seller*, *buyer*, and *goods*. In the semantics of the construction the *seller* argument role is linked to the grammatical function 'subject', the *goods* argument role is linked to the grammatical function 'object', and the *buyer* argument role is linked to the grammatical function 'indirect object'. In the sentence, the proper nouns 'John' and 'Bill' respectively occupy the functions 'subject' and 'indirect object', whereas the quantificational expression 'something' occupies the object position. It follows that (i) the referents contributed by the proper nouns 'John' and 'Bill' fill the *seller* and the

buyer argument roles respectively; and (ii) by virtue of BTL, the quantifier contributed by 'something' binds the *goods* argument role.

§4.3 Questions

I have analysed quantificational facts as involving a schematic fact $\sigma()$, and a quantifier Q which binds the (or a) parametrized argument role in the schematic fact. Though they involve schematic facts, quantificational facts themselves are not schematic (unless they are 'schematized quantificational facts'): They are not schematic because the parameter they involve is not 'free'.

Can schematic facts be expressed directly by natural language sentences, without the parametrized argument role being quantified over? I think so. I take the semantic content of *questions* to be such a schematic fact.

There are various sorts of questions: *yes-no* questions ('Is John home?'), *wh*-questions ('Where is John?'), alternative questions ('Will John go or will you come?', 'Will John go or not?'), etc. So-called indirect questions, such as 'Paul wonders *where John is*', are not questions in the pragmatic sense, and I will not be concerned with them in what follows. I will consider only the simplest cases: *yes-no* questions and *wh*-questions. I start with the latter.

In 'Who came?', the argument role linked to the subject position is parametrized. No complete fact is therefore expressed, but only a function from parameter values to facts. The role of question-words like 'who', 'what' or 'where' is triple:

- (i) Like quantificational expressions, they occupy a particular grammatical function (subject, object, etc.) and thereby indicate which argument role the question concerns: the question concerns that argument role of the expressed relation which is linked to the grammatical function occupied by the question-word.
- (ii) They provide an indication concerning the sortal nature of the parametrized argument: person (who), thing (what), place (where), time (when), etc.
- (iii) Together with word-order and intonation, they indicate that the utterance is a question, that is, an utterance with such and such *felicity conditions*. A question expressing a schematic fact $\sigma()$ is felicitous only in a context in which it is followed by an utterance providing a value for the parameter. The pair consisting of the question followed by the answer expresses a fact, but the question in isolation expresses only a schematic fact.

In order to stress what is common to questions and quantificational utterances, I represent questions as follows:

$\sigma(?)$

The content of both a quantificational utterance $\sigma(Q)$ and a question $\sigma(?)$ involves a schematic fact $\sigma()$. The difference between quantificational utterances and questions is this: In quantificational utterances the schematic fact is turned into a complete (quantificational) fact. Not so with questions: questions express schematic rather than complete facts.

At this point a question arises. I said that sentences containing quantificational expressions do *not* express schematic facts, but complete facts (namely quantificational facts). That is so, I said, because the parametrized argument role is not free but 'bound' by the quantifier. But question-words also bind an argument role of the relation expressed by the verb: In 'Who came?', the question-word binds the argument role linked to the subject position, exactly as 'Someone came' binds the argument role linked to the subject position. Still, I maintained, questions express schematic facts. How can that be? How can binding affect schematicity in one case but not in the other?

In fact there are two forms of binding, and two senses in which a quantificational expression *Quant* (e.g. 'someone' or 'everything') 'binds' a particular argument role. *Quant* does two things:

(i) It has the effect of parametrizing a particular argument role, namely that which is linked to the grammatical position occupied by *Quant* in the sentence. In this respect there is no difference between a quantificational expression *Quant* and a question-word *Wh*. They both parametrize a particular argument role through the BTL principle. The latter must be reformulated as a general principle concerning parametrization: Both quantificational expressions and question-words parametrize a particular argument role in the relation expressed by the verb, namely that argument role which is linked to the grammatical position occupied by the expression. Expressions which have this property I call 'parametrizers'. A parametrizer, whether a quantificational expression or a question-word, 'binds' a particular argument role in the sense of parametrizing it. But quantificational expressions do something else as well.

(ii) *Quant* denotes a quantifier. That quantifier 'binds' the parametrized argument role in the sense that it quantifies over its values. Binding in this sense has the effect of completing the schematic fact into a quantificational fact. A question-word *Wh* does not 'bind' the parameter in this sense, for it does not denote a quantifier — indeed it does not denote anything. The meaning of a question-word is pragmatic: question-words

signal that the utterance is a question, that is, an utterance with certain felicity conditions, and presuppose that those conditions are satisfied. Question-words therefore constrain the context in which the sentence containing them occurs, but make no contribution to the content expressed by such sentences (except insofar as they parametrize the argument role linked to the position they occupy).

So much for *wh*-questions. *Yes-no* questions are a more delicate matter, because they do not appear to express a schematic fact. According to the traditional speech-act analysis, a *yes-no* question such as 'Is John home?' expresses a complete proposition, rather than a propositional function. In situation theory, however, a fact consists of a relation, a sequence of arguments, *and a polarity*. So far I have been concerned only with one form of schematicity: that which results from parametrizing an argument role. But it is also possible for the formal constituents of the fact, viz. the relation or the polarity, to be parametrized. *Yes-no* questions can thus be construed as a special case in which parametrization concerns the polarity of the expressed fact. If we use question marks to represent parametrized constituents, the difference between *wh*-questions like 'Who is bald?' and *yes-no* question like 'Is Paul bald?' comes out as follows:

Who is bald? = $\langle\langle \text{Bald}^1, ?, + \rangle\rangle$
 Is Paul bald? = $\langle\langle \text{Bald}^1, \text{Paul}, ? \rangle\rangle$

This analysis, however, raises a serious objection. As Cornulier pointed out in his insightful paper on the semantics of questions (Cornulier 1982), a *yes-no* question such as 'Is Paul bald?' is not equivalent to its negation: 'Is not Paul bald?' The first question asks whether a certain state of affairs (Paul's being bald) is actual, while the other concerns a different state of affairs (Paul's not being bald). If the polarity of the expressed fact was parametrized in *yes-no* questions, there could be no such difference. 'Is Paul bald' and 'Is not Paul bald' would both have the schematic content mentioned above: $\langle\langle \text{Bald}^1, \text{Paul}, ? \rangle\rangle$.

To deal with negative questions I think we must complicate the analysis and remember the generalization I made earlier: unless it is a mere interjection, an English sentence always expresses a δ -complex fact $\langle\langle s \models \sigma \rangle\rangle$. As such a complex fact contains the simpler fact σ as a constituent, two polarities are involved, corresponding to the δ -complex fact $\langle\langle s \models \sigma \rangle\rangle$ and to the internal fact σ respectively. This comes out clearly if we represent δ -complex facts thus:

$\langle\langle \models, \langle s, \langle\langle \mathbb{R}^n, \langle a_1, \dots, a_n \rangle, i \rangle \rangle, i \rangle \rangle$

It is therefore possible to parametrize the polarity of the complex fact $\langle\langle s \models \sigma \rangle\rangle$, without parametrizing the polarity of the internal fact σ — the latter can be negative as well as positive. This accounts for the non-equivalence of questions and their negations: Both express a δ -complex schematic fact whose polarity (but *not* that of the internal fact) has been parametrized:

Semantic content of positive questions:

$$\langle\langle \models, \langle s, \langle\langle R^n, \langle a_1, \dots, a_n \rangle, + \rangle \rangle \rangle, ? \rangle\rangle$$

Semantic content of negative questions:

$$\langle\langle \models, \langle s, \langle\langle R^n, \langle a_1, \dots, a_n \rangle, - \rangle \rangle \rangle, ? \rangle\rangle$$

V. INFORMATIONAL STRUCTURE

§5.1 Exercised situations as 'topics'

In the course of this paper I made several assumptions which it is time to scrutinize. Two assumptions are especially important. First, I suggested that we equate the exercised situation and the 'topic' of the utterance in the traditional sense: that which the speaker is talking about. Correspondingly, the nucleus (the right hand side in the Austinian proposition) can be equated with the 'comment': what the speaker says concerning the topic. If this equation is correct, Austinian semantics should connect up with the vast literature on the topic/comment distinction and accommodate its findings (see Lambrecht 1994 for a recent survey).

Second, I implied (and in Recanati 1997 I said) that the exercised situation is necessarily external to the nucleus. If the exercised situation is mentioned and becomes a constituent of the nucleus, then it ceases being the exercised situation and a new exercised situation emerges (§3.1). Let me state this assumption a bit more explicitly:

Principle of Non-Redundancy:

The exercised situation must be distinct from any situation mentioned in the utterance.

The Principle of Non-Redundancy can be abridged as: $[s'] \neq s$, where $[s']$ is the exercised situation and s is the mentioned situation. It can be read as saying that the topic cannot be part of the comment.

The Principle of Non-Redundancy raises two problems. The first problem concerns the cases in which a situation is reflexively about itself. In 'Situations, Sets and the Axiom of Foundation', Barwise gives seven examples of such reflexive situations, including Gricean intentions, the Cartesian *cogito*, self-referential and 'specular' representations, and common knowledge (Barwise 1989: 177-200). I myself analysed *perspectives* as reflexive situations (Recanati 1997: 70n). But if there are reflexive situations, why is it not possible for a situation to be both exercised and mentioned?

As I pointed out in §3.1 the reason why it seems that a situation cannot be both exercised and mentioned is this. A situation is exercised when it is presented as supporting the fact which the utterance expresses. Now a situation cannot support a fact which contains that situation itself as constituent, because a fact which is relativised to an entity α cannot contain α as a constituent. Hence it seems that we cannot have it both ways: Either we give up our definition of the support relation in terms of relativisation, or we stick to the Non-Redundancy Principle and give up the hope of accounting for reflexive situations.

There is a way out of this dilemma. We can deny that a fact cannot be relativised to an entity α if it contains α as a constituent. For a fact can contain the same entity twice, if the entity in question fills two distinct argument roles in the fact. Let us imagine that that is the case. We can relativise the fact in question to α , by suppressing one of the two argument roles filled by α , without automatically suppressing the other. In such a situation, the fact is relativised to α yet it (still) contains α as a constituent. That is precisely what happens when a situation s is reflexive. The unrelativised fact involving s is the fact that: *in situation s*, $F(s)$. The situation s occurs in that fact twice. Through relativisation, we obtain the simpler fact which the situation supports — the fact that: $F(s)$.

Let us assume that reflexive situations can indeed be handled in that manner. It follows that the Principle of Non-Redundancy is not unrestrictedly valid, yet we need not give it up. Reflexive situations are rather special cases; and non-reflexive situations are sufficiently typical for a generalization such as the Principle of Non-Redundancy to be worth making. If we want a more universal principle, we can modify the Principle of Non-Redundancy as follows:

Principle of Non-Redundancy (modified)

Unless it is reflexive, the exercised situation must be distinct from any situation mentioned in the utterance.

From now on I shall ignore reflexive situations, hence I will stick to the unmodified version of the Principle of Non-Redundancy.

The Principle of Non-Redundancy raises a second problem, one which arises even if we disregard reflexive situations. So far, I have assumed that a situation is mentioned just in case it is linguistically *articulated* in the sentence. As we have seen, a situation can be articulated in many different ways: through tenses, adverbial expressions, attitudinals ('John believes (that)', etc.), or singular terms in subject position. Let us make this auxiliary assumption explicit:

(M1)

A situation *s* is mentioned in an utterance *u* if and only if there is something in the uttered sentence which denotes *s* (possibly in a context-dependent manner).

If we interpret the Principle of Non-Redundancy in the light of (M1), it says that the exercised situation cannot be articulated, that is, denoted by something in the sentence, without ceasing to be the exercised situation. In conjunction with the equation of topics and exercised situations, this entails that *topics cannot be articulated*. Now that consequence conflicts with well-documented findings in the theory of informational structure.

In the informational structure literature, examples are given and analysed in which some constituent of the sentence identifies the 'topic' while other constituents contribute to the 'comment'. Think of the following contrast (from Recanati 1996: 463-464):

- (1) - What's new in the class?
- Mary is in love with John
- (2) - How is Mary?
- She is in love with John.
- (3) - How is John?
- Mary is in love with him

In the first case the topic is the situation in the class; that topic is mentioned in the question, but not in the answer. As the topic of the answer is not mentioned in that answer itself (but only in the question), the Principle of Non-Redundancy is respected. In the other two examples, however, the topic — Mary and John respectively — are articulated in the answer, through the pronoun.

In order to identify the topic, a sentential constituent must be unstressed. Thus in

(4)

(4) *Mary* is in love with John

Mary cannot be the topic. Whatever is stressed is 'in focus' and cannot be (part of) the topic. Other ways of putting in focus are word order, cleft-constructions, and so forth. Now pronouns are always unstressed, and this suggests that the referents of pronouns are somehow constitutive of the exercised situation, of the 'topic'. Whatever we think of this suggestion, which will be pursued below (§5.2), it is pretty clear that, in examples such as (2) and (3), the topic is articulated in the utterance. Hence we cannot maintain the Principle of Non-Redundancy (interpreted in the light of M1) if we equate the exercised situation and the topic in the traditional sense.

To solve that difficulty, I suggest that we replace (M1) by a weaker interpretation of what it means for a situation to be mentioned:

(M2)

A situation is mentioned in an utterance *u* if and only if it is a constituent of the nucleus of *u* (i.e. of the right hand side of the Austinian proposition expressed by *u*).

Interpreted in the light of (M1), the Principle of Non-Redundancy is unacceptable because it conflicts with well-established findings in the theory of topics. Interpreted in the light of (M2), however, the Principle of difference is fine. It says — and says only — that the exercised situation must be *external to the nucleus*. If a situation is a constituent of the nucleus, it cannot be the exercised situation (the topic); if an expression identifies or articulates the topic, it cannot contribute to the nucleus.

Interpreted in the light of (M2), the Principle of Non-Redundancy prevents the exercised situation from being 'mentioned', but it does not prevent it from being 'articulated'. If the topic is articulated, however, the expression which articulates it cannot at the same time contribute to the nucleus. (This is reminiscent of Austin's analysis of performative utterances. According to Austin the performative prefix does not contribute to the propositional content of the speech act, but only to its force.)

The effect of the Principle of Non-Redundancy can be described in the framework set up previously. If the situation denoted by a sentential constituent is elected as 'topic' (exercised situation), it is automatically banned from the nucleus. This is an instance of 'backgrounding'. The mechanism through which the topic is backgrounded from the nucleus is that which I discussed in §2.3: The nucleus is *relativised to* the topic. Thus if we consider the 'Mary is in love with John' example, we see that the nucleus covaries with the topic: In the first case ('What's new in the class?') the nucleus is the fact that Mary is in love with John; in the second case ('How is Mary?') the nucleus is the property of being in love with John; in the third case ('How is

John?') the nucleus is the property of being loved by Mary. The Austinian propositions respectively expressed are:

[The situation in the class] \models $\langle\langle$ Love², \langle Mary, John \rangle , $\rangle\rangle$

[Mary] \models $\langle\langle$ Love¹, John, $\rangle\rangle$

[John] \models $\langle\langle$ Loved-by¹, Mary, $\rangle\rangle$

§5.2 Nucleus vs. minimal content

On the view I have just sketched the nucleus can no longer be equated with the utterance's 'minimal content'. The nucleus results from a process of *relativisation to the topic* ('t-relativisation') which can only take place when the topic has been contextually identified. As the 'Mary is in love with John' example shows, that process is non-minimalist (optional), for the sentence expresses a complete proposition (the proposition that Mary is in love with John) before that process occurs. We must therefore give up the picture we started with. Instead of




figure 1

we now have the following picture:




figure 2

Once the nucleus is seen as resulting from the operation of a non-minimalist (optional) pragmatic process, namely t-relativisation, we have no reason to consider it as unaffected by other optional processes like enrichment or transfer. The non-minimal character of the nucleus must be fully acknowledged (figure 3):




figure 3

I conclude that the minimalist construal of the nucleus, which was adopted at the beginning of this paper as a working hypothesis, must be rejected if we want to accommodate well-known observations concerning informational structure. This, in itself, does not settle the debate over what is said: it is still possible to accept or

alternatively to reject Minimalism as a general doctrine. Minimalism posits an intermediary level of 'minimal' truth-evaluable content between the meaning of the sentence-type and what is said in the intuitive sense. In Figure 3 there still is such a level of minimal content, but it is no longer equated with the nucleus. Instead of three level of meaning (sentence meaning, what is said_{min}, what is said_{max}) there are four:

- Sentence meaning
 - Minimal content (what is said_{min})
 - Nucleus
 - Austinian proposition
- } what is said_{max}

As far as I am concerned, I see no reason to maintain such a complicated picture. We are much better off if, giving up Minimalism, we get rid of the intermediary level of minimal content. If we do, we are left with the following picture:




figure 4

The only levels of meaning thus posited are: the meaning of the sentence-type; the nucleus; and the complete Austinian proposition. I think all important semantic and pragmatic phenomena can be accounted for in a framework with only these three levels.

The picture can still be greatly simplified. In figure 4 we see that the nucleus is affected by the context in two ways: (i) directly, through the usual primary pragmatic processes (saturation, enrichment, transfer); (ii) indirectly, through the topic which itself affects the nucleus via t-relativisation. The simplification I have in mind is made possible by the following principle:

CTT ('context through topic') Principle:

The context acts on the nucleus only through the topic (exercised situation)

This is a bold conjecture which I cannot argue for here (no more than I can argue in favour of Maximalism). Figure 5 displays the picture which would result from accepting CTT:



Figure 5

According to CTT, all primary pragmatic processes, whether mandatory (saturation) or optional (enrichment, transfer, t-relativisation), proceed from the exercised situation. This is clear at least in the case of t-relativisation, which is obviously topic-dependent. In the case of enrichment too, the topic (the situation talked about) appears to be the controlling factor. If we are talking about a policeman regulating the traffic, the sentence 'He stopped three cars this morning' will be understood as entailing that the cars were stopped by means of an appropriate signal to the driver, rather than, say, by pressing the brakes (Rumelhart 1979). In this typical instance of top-down enrichment, it is the knowledge of the situation talked about which enables an interpreter to correctly recover the intended meaning.¹⁶

¹⁶ In the example I have just mentioned a minimalist would argue that the entailment at issue is not 'semantic': the proposition literally expressed is the proposition before enrichment, namely the minimal proposition that the policeman stopped the cars in some way or other. Maybe so. But there are cases in which the proposition literally expressed is so vanishingly abstract as to hardly make sense. Thus in *Desire*, a film by Frank Borsage (1936), the following dialogue takes place:

- Pedro!
- Yes sir.
- Take the plate to the kitchen and disarm the fricassee.

What does the complex phrase 'disarm the fricassee' literally mean? It is hard to tell. To make sense of that phrase, we must know the situation talked about. In the film, the exercised situation contains the following facts: (i) Gary Cooper (the speaker) is handing a fricassee plate to the waiter (Pedro); (ii) the fricassee plate contains a gun;

That enrichment is topic-dependent has been noted many times in connection with so-called 'incomplete quantifiers' (see e.g. Recanati 1986: 60-61, Neale 1990: 115-116). So we can admit the CTT as far as enrichment is concerned. But what about saturation? Saturation is supposed to be rule-dependent rather than topic-dependent. Let me summarize the traditional view concerning the saturation of indexicals:

Indexicals and saturation: the traditional view

It is a semantic rule (e.g. the rule that 'I' refers to the speaker) which determines the contextual value of an indexical. The situation talked about is irrelevant here.

Different indexicals depend upon different aspects of the context of utterance, but that *context* is distinct from the *situation talked about*. The context includes the speaker, the hearer, the time of utterance and such things. It can contain the situation talked about as a particular parameter, but it is distinct from it.

This view is known to be defective or at least exaggerated. Saturation is much more topic-dependent and much less rule-dependent than it suggests. Genitives provide a classical example of the need for contextual saturation: an utterance including the phrase 'John's book' does not express a complete proposition unless a particular relation has been identified as holding between the book and John. Now which particular relation is contextually singled out clearly depends on the situation talked about. The same thing holds for demonstratives: what a demonstrative like 'he' or 'this' refers to is highly topic-dependent. In general, semantic underdetermination can be overcome only by the participants' knowledge of the exercised situation. (See e.g. Clark 1992 on 'contextual expressions'.) It is only 'pure indexicals' — admittedly a very limited category — which are immune to that sort of topic-dependence.

The traditional view can easily accommodate the topic-dependence of (impure) indexicals, however. The only thing one has to do is to take the topic on board as one contextual parameter among others. Thus we could rephrase Austinian semantics within a Kaplanian sort of framework, by treating the meaning of a sentence as a function from extended Kaplanian contexts to Austinian propositions (where an 'extended Kaplanian

(iii) that gun has just fallen from the hands of someone during a brief fight around the dinner table. With respect to that situation, the phrase 'disarm the fricassee' makes sense: it means that the waiter must remove the gun from the plate. This is, again, a typical instance of top-down enrichment; and it clearly involves the situation talked about. (Other examples in the same vein can be found in Searle 1978 and 1980; see also Travis 1975, 1981). Whether a minimal proposition can still be isolated in such examples is a central issue for the debate over what is said, but not one I can deal with in the limits of this paper.

context' is a Kaplanian context containing the exercised situation along with the standard parameters).

The CTT represents a totally different option: All contextual features relevant to the determination of what is said are construed as aspects of the exercised situation. Exercised situations thus *replace* Kaplanian contexts. This view is perfectly compatible with the notion that indexicals are rule-governed. The 'rules' governing indexicals are essentially *constraints on their use*; and these constraints, in turn, can be described as *constraints on the exercised situation*. Thus a referential expression (be it an indexical or a proper name) can be construed as demanding that the exercised situation contain an appropriate referent. Here 'demands that' can be replaced by: 'is felicitously used only if'. The Familiarity Theory of Definiteness (Hawkins 1978, Heim 1988) can be reformulated within that sort of framework, but I will not go into this matter here.

§5.3 Topics and truth-conditions

According to Peter Strawson (1977: 93), the failure of reference entails a 'truth-value gap' only when the referring expression contributes the 'topic' of the utterance.

Consider, for example, utterances (5) and (6):

(5) The King of France visited the Exhibition yesterday

(6) The Exhibition was visited by the King of France yesterday

As there is no King of France, both (5) and (6) suffer from reference failure; but in (6) the topic is the Exhibition, while in (5), presumably, the topic is the King of France.¹⁷ Hence the failure of reference is felt as more damaging in (5) than it is in (6). As the object which the speaker attempts to talk about and characterize in (5), namely the King of France, does not exist, the characterization of that object as having visited the Exhibition yesterday is 'neither correct nor incorrect'. In (6), however, the object talked about and putatively characterized, namely the Exhibition, undoubtedly exists; the failure of reference only affects the characterization which is offered of that object.

I think Strawson is right in his observation that our intuitions concerning falsity and truth-value gaps in cases of reference failure depend a lot on the topic of discourse.¹⁸ Apparently Stephen Neale thinks this is a reason for disregarding those intuitions, for, he says,

¹⁷ Normally (though not invariably) the subject term identifies the topic.

¹⁸ Strawson's claim about topics and truth-value gaps should be carefully separated from his claim concerning the semantic role of definite descriptions. If we are uncertain

Surely the truth-value of what one says depends upon whether the world is as one has said it is; to let the decision as to whether one has said something false or said nothing at all depend upon such things as what is the primary or overriding focus of the discourse at any given moment — to the extent that such a notion is even theoretically manageable — is to give up this idea. Indeed it is to give up doing serious semantical work altogether... (Neale 1990: 28)

I think Neale is moved by the putative fact that (5) and (6), which differ only by their respective 'topic', express the same proposition.¹⁹ How then could they fail to have the same truth-value? There is the world on the one hand, and a single proposition — that which both (5) and (6) express — describing the world as being thus and so on the other hand. If the world is as described by that proposition, (5) and (6) are both true; if not, they are both false. How could extraneous matters such as the 'point' of the statement affect the truth-value of the utterance?

Neale's remark concerning the alleged truth-conditional irrelevance of topics objects not only to Strawson's treatment of truth-value gaps but also to my version of

that definite descriptions are singular terms, we should use examples involving names rather than descriptions. Suppose we have heard of a city in Spain, called Torpedo, with certain properties. Unbeknownst to us, the city does not exist (the person who mentioned it was joking, but we did not realize). John, who has also heard of the city, expressed the desire to visit it while in Spain. A couple of days after his departure, I ask my friend: 'Where do you think John is presently?' She answers: 'John is either in Valencia or in Torpedo, depending on how much time he spent in France'. Suppose that in fact John is still in France. Then we are prone to consider my friend's utterance as false, even if we know that Torpedo does not exist. For we are talking about John's present location, and what is said about his location ('either in Valencia or in Torpedo') is clearly incorrect. The non-existence of Torpedo does not make it any less incorrect as a characterization of John's current location. Contrast this case with the case in which we are talking, not about John, but about Torpedo. I say to my friend: 'Torpedo is very small'. Is this true or false? According to Strawson, 'if we know of the reference-failure, we know that the statement cannot really have the topic it is intended to have and hence cannot be assessed as putative information about that topic' (Strawson 1977: 93). The speaker makes neither a correct nor an incorrect statement about Torpedo, for there is no such thing as Torpedo.

¹⁹ (5) and (6) also differ at the level of grammatical form, but this is irrelevant. A difference in context is sufficient to induce the relevant difference in topic, without any accompanying grammatical difference. It is that sort of case which Neale has in mind.

Austinian semantics. If Neale is right, the view I have sketched is badly mistaken: it presents itself as a theory of truth and content, while in fact it's only a theory about 'topics' and related pragmatic issues.

Neale is not alone in thinking that topics have no (direct) relevance to truth-conditions; that claim is almost universally accepted among philosophers. This is because we have rather strong intuitions, for example the intuition that an active sentence and its passive counterpart (e.g. (5) and (6) above, or 'John kissed Mary' and 'Mary was kissed by John') say the same thing and are true in the same circumstances, despite their difference in informational structure. In the same way, the sentence 'Mary is in love with John' was presented above in three different contexts characterized by a variation in the topic of the utterance:

- (1) [- What's new in the class?]
- Mary is in love with John
- (2) [- How is Mary?]
- She is in love with John.
- (3) [- How is John?]
- Mary is in love with him

Intuitively, it seems that 'Mary is in love with John' in (1), 'She is in love with John' in (2), and 'Mary is in love with him' in (3) all have the same truth-conditions, despite the difference in their respective topics. The fairly strong intuitions we have concerning those examples suggest a general principle:

Invariance Principle

The truth-conditions of an utterance are invariant under topic change.

Evidently, this principle conflicts with the theory I have sketched.

Despite our intuitions to the contrary, I think the Invariance Principle is misguided. There are cases in which a change in topic affects the intuitive truth-conditions of the utterance: All the cases in which the topic is unarticulated ('extrasentential topic cases', as I call them) fall in that category. It is when the topic is articulated in the sentence ('sentential topic cases') that we have strong intuitions in support of the Invariance principle; but I will show that those intuitions themselves can easily be accounted for without accepting the Principle.

Let us start with extrasentential topic cases. Consider the Barwise-Etchemendy example I mentioned earlier in this paper (§1.1). Commenting on the poker game I am

watching, I say 'Claire has the ace of hearts'. Let us call the poker game in question 'pk1'. Now suppose that, unbeknown to me, another poker game is taking place elsewhere; let us call it 'pk2'. According to Barwise and Etchemendy, my utterance is true iff pk1 supports $\langle\langle$ Claire has the ace of hearts $\rangle\rangle$. Thus if I am mistaken and Claire is not a participant in the game I am watching (pk1), my utterance is false, even if Claire is a participant in the game pk2 and has the ace of hearts there. The truth-conditions of my utterance are therefore different from what they would have been had pk2 been the topic of my utterance instead of pk1. My actual utterance is about pk1 and expresses the Austinian proposition (7); the counterfactual utterance would have been about pk2 and would have expressed the Austinian proposition (8):

(7) [pk1] $\models \langle\langle$ Has², <Claire, the ace of hearts>, + $\rangle\rangle$

(8) [pk2] $\models \langle\langle$ Has², <Claire, the ace of hearts>, + $\rangle\rangle$

Our intuitions in cases like this go against the Invariance Principle, to a certain extent at least. We *do* have the intuition that the actual utterance is not straightforwardly true if Claire is not a participant in pk1. (In contrast, the counterfactual utterance would be straightforwardly true in such circumstances.) To be sure, we do also have the intuition that, in some sense, what is said is true even in such a case, provided Claire has the ace of hearts somewhere. Traditional semantics handles those conflicting intuitions by saying that the first one concerns what is 'conveyed' by the utterance, while the second one — the only one that matters to semantics — is about what is strictly and literally said. Austinian semantics handles those intuitions quite differently, by distinguishing two sets of truth-conditions, corresponding to the two levels of semantic evaluation (the nucleus and the complete Austinian proposition). At the nuclear level, the utterance is true in a situation *s* iff Claire has the ace of hearts in *s*. Those nuclear truth-conditions are the same in (7) and (8), because the nucleus is the same: this captures the invariance intuition. But the Austinian truth-conditions are not invariant under topic change: (7) is true iff Claire has the ace of hearts *in pk1*, while (8) is true iff Claire has the ace of hearts *in pk2*.

Let us now consider cases in which the topic *is* articulated in the sentence itself. When that is so, which constituent is topicalized does not seem to affect the truth-conditions even at the 'Austinian' level: they remain invariant, in accordance with the Principle. Thus (5) and (6) have the same overall truth conditions, (2) and (3) have the same overall truth-conditions, etc. But our invariance intuitions can be accounted for by saying that the Austinian propositions respectively expressed are *equivalent*. This equivalence is indeed guaranteed by the fact that (owing to the Principle of Non-Redundancy) the nucleus is systematically relativised to the topic. Thus the Austinian

propositions expressed by 'She is in love with John' (where Mary — the referent of the pronoun — is the topic) and 'Mary is in love with him' (where John is the topic) are equivalent, even though structurally they are quite different:

[Mary] \models $\langle\langle\text{Love}^1, \text{John}, +\rangle\rangle$

[John] \models $\langle\langle\text{Loved-by}^1, \text{Mary}, +\rangle\rangle$

The first Austinian proposition says of Mary that she has the property of loving John; the second Austinian proposition says of John that he has the property of being loved by Mary. Though equivalent, those propositions are different, as their constituents (the exercised situation and the nucleus) are different. Exactly the same thing can be said of the Strawsonian examples (5) and (6).

I conclude that topics are truth-conditionally relevant, appearances to the contrary notwithstanding. In extrasentential topic cases, the nuclear truth-conditions are invariant under topic change, but the Austinian truth-conditions systematically vary. In sentential topic cases, the Austinian truth-conditions are invariant (because the Austinian propositions are equivalent) but, owing to the primary pragmatic process of relativisation, the nuclear truth-conditions co-vary with the topic.

References

- Austin, J. L., 1971: *Philosophical Papers*. 2nd ed., Oxford: Clarendon Press.
- Bach, K., 1994: Semantic Slack: what is said and more. In S. Tsohatzidis (ed.), *Foundations of Speech Act Theory*, London: Routledge, pp. 267-291.
- Barwise, J., 1989: *The Situation in Logic*. Stanford: CSLI.
- Barwise, J. and J. Etchemendy, 1987: *The Liar: An Essay on Truth and Circularity*. New York: Oxford University Press.
- Barwise, J. and J. Perry, 1983: *Situations and Attitudes*. Cambridge, Mass: MIT Press/Bradford Books.
- Beeson, F., 1979: J. L. Austin on Descriptive and Demonstrative Conventions. *Studies in Language* 3: 1-36.
- Clark, H. 1992: *Arenas of Language Use*. Chicago: The University of Chicago Press & CSLI.
- Cornulier, B. de, 1982: Sur le sens des questions totales et alternatives. In F. Récanati (ed.), *Langages 67: La Signalisation du Discours*, Paris: Larousse, pp. 55-109.
- Davidson, D., 1967: The Logical Form of Action Sentences. Reprinted in his *Essays on Actions and Events*, Oxford: Clarendon Press, 1980, pp. 105-122.
- Devlin, K., 1991: *Logic and Information*. Cambridge: Cambridge University Press.
- Fillmore, C., 1975: An Alternative to Checklist Theories of Meaning. *BLS* 1: 123-131.
- Fillmore, C., 1982: Frame Semantics. In Linguistic Society of Korea (ed.), *Linguistics in the Morning Calm*, Seoul: Hanshin, pp. 11-138.
- Fillmore, C., 1986: Pragmatically Controlled Zero Anaphora. *BLS* 12: 95-107.
- Fillmore, C. and B. Atkins, 1992: Toward a Frame-Based Lexicon: the Semantics of RISK and its neighbors. In A. Lehrer and E. Kittay (eds.), *Fields, Frames and Contrasts: New Essays on Semantic and Lexical Organization*, Hillsdale, N.J: Lawrence Erlbaum, pp. 75-102.
- Fillmore, C. and P. Kay, 1994: *Construction Grammar Coursebook*. University of California, Berkeley.
- Goldberg, A., 1995: *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago: The University of Chicago Press.
- Hawkins, J., 1978: *Definiteness and Indefiniteness*. London: Croom Helm.
- Heim, I., 1988: *The Semantics of Definite and Indefinite Noun Phrases*. New York: Garland.
- Lambrecht, K., 1994: *Information Structure and Sentence Form*. Cambridge: Cambridge University Press.
- Langacker, R., 1987: *Foundations of Cognitive Grammar*, vol. 1. Stanford: Stanford University Press.

- Lyons, J., 1975: Deixis as the Source of Reference. In E. Keenan (ed.), *Formal Semantics of Natural Language*, Cambridge: Cambridge University Press, pp. 61-83.
- McCarthy, J., 1993: Notes on Formalizing Context. *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, San Mateo: Morgan Kaufmann Publisher, vol. 1, pp. 555-560.
- Neale, S., 1990: *Descriptions*. Cambridge, Mass: MIT Press/Bradford Books.
- Nunberg, G., 1995: Transfers of Meaning. *Journal of Semantics* 12: 109-132.
- Perry, J., 1986: Thought without Representation. Reprinted in *The Problem of the Essential Indexical and Other Essays*, New York: Oxford University Press, 1992, pp. 205-225.
- Prandi, M., 1992: *Grammaire philosophique des tropes*. Paris: Minuit.
- Recanati, F., 1986: Contextual Dependence and Definite Descriptions. *Proceedings of the Aristotelian Society* 87: 57-73.
- Recanati, F., 1993: *Direct Reference: From Language to Thought*. Oxford: Blackwell.
- Recanati, F., 1995: The Alleged Priority of Literal Interpretation. *Cognitive Science* 19: 207-232.
- Recanati, F., 1996: Domains of Discourse. *Linguistics and Philosophy* 19: 445-475.
- Recanati, F., 1997: The Dynamics of Situations. *European Review of Philosophy* 2: 41-75.
- Recanati, F., forthcoming: Context-Shifting in Metarepresentations.
- Rumelhart, D., 1979: Some Problems with the Notion of Literal Meaning. In A. Ortony (ed.), *Metaphor and Thought*, 2nd ed., Cambridge: Cambridge University Press, 1993, pp. 71-82.
- Salmon, N., 1991: The Pragmatic Fallacy. *Philosophical Studies* 63: 83-97.
- Searle, J., 1978: Literal Meaning. *Erkenntnis*, 13: 207-24.
- Searle, J., 1980: The Background of Meaning. In J. Searle, F. Kiefer and M. Bierwisch (eds.), *Speech Act Theory and Pragmatics*, Dordrecht: Reidel, pp. 221-32.
- Sechehaye, A., 1926: *Essai sur la structure logique de la phrase*. Paris: Champion.
- Sperber D. & D. Wilson, 1986: *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Strawson, P.F., 1964: Identifying Reference and Truth-Value. Reprinted in Strawson 1977, pp. 75-95.
- Strawson, P.F., 1977: *Logico-Linguistic Papers*. London: Methuen.
- Tesnière, L., 1969: *Eléments de Syntaxe Structurale*. 2nd ed., Paris: Klincksieck.
- Travis, C., 1975: *Saying and Understanding*. Oxford: Blackwell.
- Travis, C., 1981: *The True and the False: the Domain of the Pragmatic*. Amsterdam: Benjamins.