

Domains of discourse

François Recanati

► **To cite this version:**

François Recanati. Domains of discourse. Linguistics and Philosophy, Springer Verlag, 1996, 19 (5), pp.445-475. ijn_00000178

HAL Id: ijn_00000178

https://jeannicod.ccsd.cnrs.fr/ijn_00000178

Submitted on 9 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.

Domains of discourse*

François Recanati
CREA (Ecole Polytechnique/CNRS)
1 rue Descartes, 75005 Paris, France

I. Austinian semantics

Natural language quantifiers often seem implicitly restricted. Thus when we say 'The burglar took everything', we have the feeling that 'everything' ranges over the domain of valuable objects in the house — not everything in the world. (In this case, it can be argued that there *has to be* some contextual restriction or other, for a totally unrestricted notion of 'everything in the world' hardly makes sense.) In the same way, someone who says 'Most students came to the party' is likely to have a particular group of students in mind, such that most students in that group came to the party.

Footnotes

* A decade ago I presented a first version of the theory of domains, inspired by Fauconnier's work, in a talk to the Aristotelian Society (Recanati 1987). Someone in the audience objected that the notion of a domain was unclear: were the domains supposed to be real situations or mental constructs? At the time I did not know what to say, and I left the topic. Recently, I came to realize that the domain of discourse approach is widely considered as refuted by a certain class of counter-examples originally due to Lewis and McCawley. So I thought I had to take up the issue again and I planned to write a defense of the domain framework. It soon occurred to me that the theory of reference expounded in Recanati (1993) provided not only an answer to the embarrassing question about the nature of domains, but also a way of articulating the very different perspectives offered by situation theory, Fauconnier's work on mental spaces, and Sperber and Wilson's theory of interpretation — three frameworks which had influenced me but which had seemed hard to conciliate.

A number of people commented on earlier versions of this paper. I am particularly grateful to Kent Bach, Roberto Casati, Mark Crimmins, Jerry Katz, Alberto Voltolini, and two referees for this journal. I am also grateful to Scott Soames, whose paper, 'Incomplete Definite Descriptions', provoked me and provided the impetus for writing this paper.

Similar phenomena can be found which do not involve overt quantifiers. If, when asked whether I have eaten, I reply 'I've had breakfast', I clearly mean that I've had breakfast on the day of utterance. In contrast, an utterance like 'I've been to Tibet' is likely to mean that the speaker has been to Tibet once — not (or not necessarily) that he has been to Tibet on the day of utterance (Sperber and Wilson 1986: 189-90). This difference can be accounted for in terms of a contextually variable domain of discourse (Recanati 1989b: 305-6). If we accept Davidson's analysis of action sentences, we can treat 'I've had breakfast' as meaning approximately: 'there is a past event which is my having breakfast'. Contextually the domain of quantification is restricted; the utterance 'concerns' only what happened so far on a certain day.

Among the phenomena which seem amenable to a treatment in terms of contextually restricted quantification we find incomplete descriptions. Often we use a description 'the F' even though Russell's uniqueness condition is not satisfied. Thus I say 'Close the door', knowing full well that there is more than one door in the universe. Does this object to Russell's analysis of descriptions? Not if we accept the notion of a contextually restricted domain of quantification. If we do, we can say that the uniqueness condition is satisfied — not in the world, but in a relevant portion of the world. The latter is the domain of discourse for incomplete descriptions. Such an analysis has been offered many times, e.g. in Kuroda (1982), Barwise & Perry (1983), Fauconnier (1985), Recanati (1987), Barwise & Etchemendy (1987).

We can go further and generalize the notion of a contextually variable domain of discourse. Following Barwise, Perry and their colleagues, we can view the domain of discourse as a 'situation' tacitly referred to in the discourse. Whenever there is quantification, it is relative to the situation tacitly referred to, but the generalized notion of domain of discourse as 'parameter situation' applies whether or not the utterance involves some form of quantification. Barwise and Etchemendy give the following example (Barwise and Etchemendy 1987:121-2). Looking at a poker game, I say 'Claire has a good hand'. I describe the situation I am witnessing as a situation in which Claire has a good hand. If I am mistaken and Claire is not a constituent of the situation (if she is not among the players of the game I am watching, contrary to what I believe), my utterance is not true — even if Claire is playing poker in some other part of the city and has a good hand there; the utterance is not true because the situation referred to is not of the type described. If we did not have partial situations as domains of evaluation for our statements, but only the total world, we would have to say that the utterance is true, provided Claire is playing poker in some other part of the city and has a good hand there. So we see that utterances, like quantifier phrases, are interpreted relative to some partial, contextually determined domain of discourse rather than relative to a fixed, total world.

John Austin's theory of truth explicitly relies on the notion of a parameter situation relative to which utterances are to be interpreted. An utterance is true, Austin says, if the 'historic' situation it refers to is of the type with which the sentence is associated by the descriptive conventions of the language (Austin 1970: 122). Contemporary situation theorists follow Austin in describing the complete content of an utterance as a pair $\langle \underline{s}, T \rangle$ consisting of a real situation and a type of situation (Barwise and Etchemendy 1987; Barwise 1989). The type of situation is very much like a 'proposition' in the traditional sense, but the utterance's semantic content (the 'Austinian proposition', as Barwise & Etchemendy say) is richer: it also involves the situation tacitly referred to by the utterance. The Austinian proposition $\langle \underline{s}, T \rangle$ is true just in case \underline{s} is of type T .

What are situations, and what are types of situation? Let us start with John Perry's notion of an 'issue' (Perry 1986). To get an issue, we need a \underline{n} -place relation and a sequence of \underline{n} objects; this generates the issue, whether or not those objects stand in that relation. There are two possible answers to an issue: Yes (+) and No (-). A putative 'fact' consists of an issue and an answer to that issue. If \underline{u} is the issue, whether John loves Mary, then the fact that John loves Mary can be represented as the pair: $\langle \underline{u}, + \rangle$. (From now on, I shall omit the sign '+' from the representation of positive facts.) The world, which provides answers to all issues, can be represented as the collection $@$ of all the facts, whether positive or negative. Situations — portions of the world — provide answers only to some issues; they can be represented as subsets of $@$, that is, as sets of facts.

The facts I have talked about so far are all atomic facts, i.e. triples consisting of a relation, a sequence of arguments, and a positive or negative answer. They are the only facts which situations contain. What about non-atomic facts, e.g. disjunctive facts (John loves Mary or Peter hates Paul), general facts (There is a man who loves Mary), and so forth? A situation \underline{s} cannot contain such facts, but it can support them, by virtue of the atomic facts it contains. The 'support' relation can be defined along the following lines:

- A situation \underline{s} supports an atomic fact σ (in symbols: $\underline{s} \models \sigma$) just in case $\sigma \in \underline{s}$.
- A situation \underline{s} supports a disjunctive fact $\sigma \vee \sigma'$ just in case $\underline{s} \models \sigma$ or $\underline{s} \models \sigma'$.
- A situation \underline{s} supports an existential fact $\exists x \sigma(x)$ iff, for some object \underline{a} , \underline{s} supports $\sigma(\underline{a})$.
- ...

A type of situation, T_σ , can now be represented as a set of situations, namely the set of all situations which support σ (whether σ is atomic or not). To say that a situation is of the type T_σ is to say that it belongs to $\{s \mid s \models \sigma\}$.

Let us go back to our example: 'Claire has a good hand'. The utterance 'states a fact', viz. the fact that Claire has a good hand: $\langle \text{GH}(x), \text{Claire} \rangle$. This determines a type of situation $T_{\langle \text{GH}(x), \text{Claire} \rangle}$. The utterance also refers to a situation \underline{pk} . The complete content of the utterance is:

$\langle \underline{pk}, T_{\langle \text{GH}(x), \text{Claire} \rangle} \rangle$

The utterance is true iff $\underline{pk} \in T_{\langle \text{GH}(x), \text{Claire} \rangle}$.

II. A survey of logical space

Before considering alleged objections to Austinian semantics, it will be useful to undertake a brief review of the theories available for dealing with incomplete quantifiers and similar phenomena. I can think of four different classes of theory, corresponding to the four levels which must be distinguished in the interpretation of an utterance: the sentence itself qua syntactic object, its linguistic meaning or 'character', its Kaplanian 'content', i.e. the proposition literally expressed, and finally what the utterance 'conveys' via Gricean mechanisms.

The standard alternative to the explanation in terms of domains of discourse is the ellipsis theory. It says that the sentence at issue (e.g. 'The burglar took everything') is elliptical for a longer sentence, e.g. 'The burglar took everything valuable in the house'. In the same way, 'Most students came to the party' is short for 'Most students in my class came to the party'; 'I've had breakfast' is short for 'I've had breakfast today'; 'Close the door' is short for 'Close the door of our bedroom'; 'Claire has a good hand' is short for 'In this poker game, Claire has a good hand', and so forth. The original, 'elliptical' sentence is not what undergoes semantic evaluation, on this theory; the sentence is first completed, then semantically evaluated. Thanks to this two-step procedure, it is possible to maintain that the domain of discourse is always the complete 'world'.

On its standard interpretation, the ellipsis theory is what it seems to be: a theory at the syntactic level. It says that the sentence, e.g. 'The burglar took everything', is elliptical for some other sentence, as 'John went to the movie and Bob to the theater' is elliptical for 'John went to the movie and Bob went to the theater'. This theory (thus interpreted) raises three main objections:

(1) It multiplies ambiguities without necessity (Recanati 1989a: 229-31). All the examples given above turn out to be ambiguous. For example, 'The burglar took everything' can be construed in two ways: either as an instance of the complete sentence

'The burglar took everything', or as an (elliptical) instance of the sentence 'The burglar took everything valuable in the house'.

(2) For there to be ellipsis in the literal, syntactic sense, the deleted material must be recoverable on syntactic grounds (Bach 1994: 131-2), or at least there must be independent syntactic reasons for positing deleted elements in the sentence (Bach 1987: 73-4). This is the case with 'John went to the movie and Jack to the theater': we know, on purely syntactic grounds, that this is short for 'John went to the movie and Jack went to the theater'. This is not the case with the other examples we talked about.

(3) It is indeterminate which particular sentence (among several possible candidates) a given sentence is elliptical for. "When one says, e.g., 'The table is covered with books', the table the speaker has in mind can be more fully described in any number of ways, by the use of any number of non-synonymous, uniquely denoting descriptions (for example 'the table in room 209 of Camden Hall at t_1 ', 'the table at which the author of *The Persistence of Objects* is sitting at t_1 ', etc.)" (Wettstein 1981:246). Some contextual completion seems to be required, but which one? As Wettstein argued, there is no determinate answer to this question, even if we appeal to speaker's intentions.¹

¹ Many philosophers believe that Wettstein's indeterminacy objection to the ellipsis theory has been overcome. Thus it has been suggested (following Wettstein himself) that the completing material may be 'singular' in the sense of involving a particular object. Take, for example, an attributive use of the incomplete description 'the murderer'. The description must be completed for it is not uniquely identifying (there are many murderers in the world, and the sentence does not tell us which is the one the speaker is talking about). Let us suppose that the speaker is in fact talking about *Smith's* murderer. The description must be accordingly completed, but there are two ways to do it. If the description is completed through provision of further 'descriptive' material, we face the indeterminacy problem: there are many alternative 'descriptive' completions ('the murderer of our second floor neighbour', or 'the murderer of Jane's fiancé'...), and no principle of selection. But another sort of completion is possible: the description can be completed not through the provision of further conceptual material, but through provision of an individual, namely the murdered Smith himself (irrespective of how this individual is represented). By virtue of this objectual enrichment, the incomplete description contextually expresses a 'singular concept' (Soames 1986:352) which is complete (uniquely identifying) and therefore can be evaluated with respect to the total world, even if it falls short of determining a single 'descriptive' concept or mode of presentation. This is the ellipsis theorist's standard reply to Wettstein.

These objections undermine the ellipsis theory syntactically interpreted. But another interpretation is available. It can be argued that 'The burglar took everything' is short for 'The burglar took everything valuable in the house' not in the syntactic sense but in the 'Gricean' sense: the speaker literally says that the burglar took 'everything', but what she means is that he took everything valuable in the speaker's house. On this view, the proposition literally expressed by the utterance corresponds to the 'absolute' interpretation, to the effect that the only burglar there is in the world took everything there is in the world. This is very far from our intuitions concerning what the utterance means, but a Gricean mechanism is supposed to take us from this proposition literally expressed to some other proposition corresponding to what is actually meant.

The 'Gricean' theory raises the following objection. It implies that the participants in the speech situation are unaware of 'what is said' even though by normal

But the appeal to singular contents cannot save the ellipsis theory interpreted at face value. For even if we accept that the expressed content is singular, still it's totally indeterminate *which particular sentence expressing that content* the uttered sentence is elliptical for — it's indeterminate which particular 'relationally descriptive term' the incomplete description is elliptical for (on 'relationally descriptive' vs. 'thoroughly descriptive' terms, see Salmon 1981: 17-8). Is 'the murderer' short for 'the murderer of this man' or for 'Smith's murderer'? There is no determinate answer to this question. The natural conclusion is that the uttered sentence is not elliptical for some other sentence, whether the latter expresses a singular or a descriptive sort of content. (Of course, if the sentence is supposed to be completed 'descriptively' rather than 'singularly', the indeterminacy objection applies twice: once because of the indeterminacy of descriptive content, and a second time because of the indeterminacy of the linguistic material used in expressing that content.) It follows that Stephen Neale is mistaken when he writes: "There is simply no difference between saying that there is an "implicit reference" to the victim and saying that the incomplete description 'the murderer' is elliptical for a uniquely denoting description, such as 'the murderer of *that man*' (where 'that man' refers to the victim), or 'the murderer of *him*' (where 'him' refers to the victim), or '*his* murderer' (where 'his' refers to the victim), all of which contribute the same thing to the proposition expressed, viz. the descriptive condition *murderer-of-b* where *b* is the victim himself rather than some description of *b*." (Neale 1990:99). An attributively used incomplete description, in conjunction with some implicit (contextual) reference to an individual or a situation (Salmon 1981: 18n16, 1982:44), expresses the same content as a relational description explicitly referring to that individual or situation, but that does not entail that the incomplete description is *elliptical for* the relational description. (See Reimer 1992: 353 for a similar point.)

standards they fully understand the utterance. (The speaker who says 'The burglar took everything' is unaware of expressing the proposition that the only burglar there is in the world took everything there is in the world). Now this divorce between what-is-literally-said as a theoretical entity and our *intuitions* of what is said, though commonly accepted by semantic theorists, is highly problematic. In a nutshell: The notion of something which is literally said without being accessible and open to view rests on a false model of human communication, variously called 'the code model' (Sperber and Wilson 1986) or the 'message model' (Bach & Harnish 1979). According to this model, 'what is said' is determined, in a more or less complicated and context-dependent manner, by the conventional meaning of words independent of speaker's communicative intentions. Now it has been amply demonstrated in the literature that this model fails to account for context-sensitivity in natural language. 'What is said' is in fact inseparable from publicly recognizable speaker's intentions. Once we acknowledge this, we are no longer free to sever the link between our ordinary intuitions and 'what is literally said'.²

Having rejected the first and the fourth theory, we must consider those that remain. They are built around the following theses, intended to immunize them from the various objections above.

[1] The sentence, *qua* syntactic object, is not affected by the process responsible for the 'restricted' interpretations we are trying to explain.

[2] The propositional content of the utterance *is* affected by the process in question.

Thesis [1] is meant to protect us from the objections raised by the ellipsis theory syntactically interpreted. Thesis [2] follows from our rejection of the Gricean theory,

² The issue is complicated by the alleged ambiguity of the phrase 'what is said'. Several philosophers (e.g. Salmon 1982, Bach 1987, 1994) believe that there is a purely semantic notion ('what is said' is a property of the sentence, relative to a Kaplanian context) distinct from the pragmatic notion ('what is said' is what *the speaker* states — still distinct, however, from what she merely 'conveys'). I think it is hard to deny that the availability constraint applies to the pragmatic notion; this is why the Gricean theory (standardly interpreted) is unacceptable — for the distinction it uses is that between what the speaker literally states and what she merely conveys (Recanati 1989b). But what about the former, semantic notion of 'what is said'? As we shall see in the next footnote, there is a variant of the Gricean theory which appeals to the distinction between what is said (in the 'purely semantic' sense) and what the speaker states. This variant escapes the availability objection.

according to which the restricted interpretation is generated at the level of what is conveyed rather than at the level of what is said

From [1] and [2] it follows that the locus of restricted interpretations is either the linguistic meaning of the sentence (its Kaplanian 'character'), or the proposition contextually expressed (its Kaplanian 'content'). Now the former option is excluded; we cannot say that the two interpretations of an utterance like 'Most students came to the party', viz. the absolute interpretation (most students in the world) and the restricted interpretation (most students in my class), correspond to two different 'characters', for that too would multiply ambiguities without necessity. How can a given sentence be associated with two distinct characters, unless the sentence in question is ambiguous? We must therefore add a third thesis:

[3] The process responsible for restricted interpretations is 'pragmatic', in the sense that it leaves the linguistic meaning of the utterance unchanged.

It may be felt that there is a tension between [2] and [3]. [2] rules out a Gricean approach locating restricted interpretations at the level of what is meant without being said, whereas [3] requires a pragmatic explanation. The two theses do not contradict each other, however. Not all pragmatic explanations are 'Gricean' in the sense of involving a 'secondary' pragmatic process, that is, a process generating a further layer of meaning over and above what is literally said. Secondary pragmatic processes, like that which generates conversational implicatures, take the Kaplanian content of the utterance as input (among other things) and yield the implicatures as output; but there are other pragmatic processes which are 'primary' in the sense of constitutive of the Kaplanian content of the utterance (Carston 1988, Recanati 1989b, 1993). The paradigm example is 'saturation', the contextual process of assigning semantic values to indexicals and other 'free variables' in the utterance.

The Austinian theory precisely is an account in terms of 'saturation'. Austinian semantics adds a further contextual parameter to those associated with standard indexicals: the situation talked about. The absolute and the restricted interpretations result from different ways of filling out the situational slot in semantic structure. On this theory the linguistic meaning of the sentence is fixed; what varies is the value of a contextual parameter.

This review was not meant to be exhaustive. Other theories than Austinian semantics appeal to primary pragmatic processes and incorporate the three theses above.³ My aim in this section was merely to show that Austinian semantics, with its

³ In Recanati 1993 two types of primary pragmatic process are distinguished: *saturation*, the process of assigning values to contextual 'variables' in the utterance, is

generalized appeal to contextually variable domains of discourse, satisfies the desiderata which any plausible theory must satisfy, and which the standard alternatives

semantically mandatory; other primary processes, like *enrichment* and *transfer*, are optional and responsive only to pragmatic pressures. Now Kent Bach has offered an account of restricted interpretations in terms of enrichment (Bach 1987, 1994). As I understand it Bach's picture involves two steps. An utterance like 'Most students came to the party' is first associated with a mental sentence *Most students came to the party* whose truth-conditions correspond to the 'absolute' interpretation (most students in the world); as this interpretation turns out to be contextually inadequate, a pragmatic-psychological process of 'expansion' occurs which takes the mental sentence as input and yields as output another mental sentence, viz. *Most students *in my class* came to the party*. In this theory, the natural language sentence 'Most students came to the party' is *not* elliptical for 'Most students in my class came to the party'. It's not ambiguous either: it has only one literal interpretation, and it is the 'absolute' one (*Most students came to the party*). What happens is that a pragmatic process of 'expansion' takes us from this literal interpretation of the sentence to some non-literal interpretation (*Most students in my class came to the party*). Being a form of enrichment, hence a primary pragmatic process in my sense, the process of expansion appealed to by Bach has a serious advantage over the secondary (postpropositional) process appealed to by the Gricean explanation: *qua* primary pragmatic process, expansion occurs at a 'sub-doxastic level', prior to that of publicly accessible content, hence it is not subject to the objection leveled against the Gricean explanation (Recanati 1989b: 314n). In other words, Bach can assign the utterance a counter-intuitive content at the sub-doxastic level, without falling prey to the objection based on the publicity of communicative intentions.

One difference between saturation-based and enrichment-based accounts is the dissymmetry which the latter posits between the absolute and the restricted interpretation. While, in Austinian semantics, the two interpretations are on a par (they are both generated by the normal semantic processes associated with the utterance by virtue of its linguistic meaning), on Bach's theory the absolute interpretation is taken to be 'literal' and primary while the restricted interpretation is pragmatically 'derived'.

Bach's theory is worth exploring, but the notion of expansion is mysterious. It's a syntactic notion: what's enriched is not the informational content (as in Sperber's and Wilson's notion of enrichment) but the sentence which is enriched with a few extra 'words'. What sort of words are these? This theory seems to presuppose a very literal interpretation of the language of thought hypothesis and is hard to evaluate until we are also given some details about the overall framework.

(the ellipsis theory and the Gricean theory) precisely do not satisfy. I now turn to specific objections allegedly raised by the Austinian account.

III. Local domains

A class of counter-examples, originally due to Jim McCawley and David Lewis, is often quoted as having refuted the 'domain of discourse' approach to incomplete quantifiers. (L) below is an instance of this class:

(L) The dog got in a fight with another dog

In order to deal with a true instance of (L), it seems that the Austinian semanticist has to posit a situation in which (i) there is one and only one dog (in order to satisfy the uniqueness requirement conveyed by the description 'the dog') and (ii) there are two dogs (in order to make the sentence true). Austinian semantics is taken to be refuted by those cases in which certain constituents of the sentence 'place demands on situations that conflict with those arising from other constituents in the sentence' (Soames 1986: 357).

Is Austinian semantics really refuted by (L)? I do not think so. The attempted refutation relies on the crucial assumption that there is at most one situation tacitly referred to by a given utterance. But there is no reason why such an assumption should be maintained. The situation talked about is a feature of the 'context'. Now it is well-known that the context changes as the utterance takes place, to a large extent as a result of the utterance itself. Thus the earlier part of an utterance often affects the context in which the later part is interpreted. This phenomenon of intra-sentential variability has been noted many times in the literature on context-change, in connection with conjunctions, conditionals, and the defeasibility of presuppositions (see e.g. Stalnaker 1974, 1978). It is therefore to be expected that the domain of discourse itself can change in mid-utterance. This means that there can be more than one domain, more than one 'situation', corresponding to a given utterance. All the authors I mentioned earlier as having appealed to contextually restricted domains of discourse have made this point explicitly. Now, if there are more than one domain, more than one situation, then an utterance like 'The dog got in a fight with another dog' no longer raises an insuperable problem: it can be accounted for by positing two situations, one in which there is one and only one dog, and another one in which there are two dogs.

To account for the 'shifty' nature of the contextual domain of discourse Barwise and Perry (1983) introduce 'resource situations' over and above the situation the global utterance is about. As they point out, many such situations can be exploited in interpreting a given utterance: 'there is no reason to suppose that there is at most one

resource situation per utterance any more than there should be only one thing around referred to by IT in a given utterance' (Barwise and Perry 1983:153). Indeed, there can be as many domains as there are constituents in the sentence. Nowhere has this point been made more explicitly than in Kuroda's paper on 'indexed predicate calculus' (Kuroda 1982). In Kuroda's framework, sentential constituents are indexed and the indices refer to the domains (or 'mini-worlds') with respect to which those constituents are to be interpreted. The constituents within a single sentence can but need not be co-indexed. Thus Kuroda gives the following example (among many others):

(1) Since it was so stuffy in the house, Mary went up to the attic and opened the window

An utterance of (1) talks about a situation involving Mary and a particular house (whose identifiability is presupposed by the description 'the house'). Let us call this situation 'i'. 'The house' is interpreted with respect to i: the house is the house of the situation. 'The attic' is interpreted with respect to a different domain j, namely the house that is a constituent of situation i: the attic is the attic of the house. Here we have a case of domain focalization: the domain is restricted to a given portion of the initial domain. This sub-domain is easy to focus on, for it is made salient by the description 'the house'. A further, and more significant, shift of domain occurs when the third description, 'the window', is tokened. Under pragmatic pressure, a new domain k is selected: the introduction of the attic as a new topic makes it possible to interpret 'the window' as the window of the attic, rather than (implausibly) as the window of the house (Kuroda 1982:46). This is a standard instance of domain shift. It is driven by the process of 'accommodation' described by Lewis in the same paper in which the dog counterexample is presented (Lewis 1979). Using indices to represent domains of discourse (though not quite in Kuroda's manner), (1) can be represented as follows:

(1*) (Since it was so stuffy in (the house)_i, Mary went up to (the attic)_j and opened (the window)_k)_i

i = the situation the global utterance talks about

j = the house

k = the attic

This example shows that the 'domains of discourse' posited by the Austinian semanticist can be associated with expressions at all levels, from constituents of sentences to extended pieces of discourse. The Austinian semanticist is therefore happy to grant that 'contextual supplementation [i.e. the provision of a contextual domain] works at the level of constituents of sentences or utterances' (Soames 1986: 357). It is

true that Austinian semantics locates the tacit reference to a situation at the global, utterance level. But nothing prevents the pragmatic process of domain provision from working both at the level of constituents and at the level of global utterances. Thus in (1*), there is a 'main' domain, namely *i*, corresponding to the global utterance. The main domain is the situation the utterance as a whole talks about (the 'historic' situation referred to, in Austin's terminology). Local domains, such as *j* and *k* in (1*), are not alternatives but additions to that main domain.⁴

To illustrate the dialectic between local domains and main domains, consider the following variant of the Barwise-Etchemendy example: 'Claire has a better hand than the bearded man'. Let us assume that the global utterance is interpreted with respect to a particular situation, namely the poker game, as in the original example. Now the description 'the bearded man' must be interpreted with respect to a partial situation in which there is one and only one bearded man. This opens up two possibilities: either the two situations are the same, or they are not. The case is simpler if the bearded man is the bearded man of the poker game (i.e. the one and only male participant in the poker game to have a beard), but this need not be the case. It is easy to imagine a context such that the bearded man referred to in the utterance 'Claire has a better hand than the bearded man' is no longer bearded in the poker situation referred to, but was bearded in

⁴ It could be argued (indeed, Soames has argued) that once we have 'local' domains, we no longer need the main domain: we can drop the Austinian proposition and equate the content of the utterance with the 'complete' proposition resulting from local domain provision. On this view, situations are needed only locally, to help contextually enriching the content which is then evaluated in the standard way (with respect to the total world).

But it is not true that only local domains are needed. In the Barwise-Etchemendy example, the proposition that Claire has a good hand is already complete in the sense of truth-evaluable, yet the utterance provides a situation with respect to which it is to be evaluated. If this is treated as a case of domain provision (rather than as a case of ellipsis), then, I think, the domain must be considered as global, because there is no sub-sentential constituent which is affected by the provision of the domain. Moreover, the global domain is truth-conditionally relevant in this example: the intuitive truth-conditions co-vary with the situation the speaker refers to. If the speaker refers to the poker game pk₁ which takes place in front of him, the utterance is true iff Claire has a good hand in pk₁; if the speaker refers to another poker game pk₂ taking place somewhere else, the utterance is true iff Claire has a good hand in pk₂. It is therefore not possible to argue that only local domains can affect truth-conditions. (I apologize for the sketchiness of those indications; I am conscious that those issues deserve a much fuller discussion than I can offer here.)

a prior situation of which both the speaker and the listener are aware and which serves as 'resource situation' for the interpretation of the description. In such a case the description is interpreted with respect to a local domain distinct from the main domain (the poker situation):

(2*) (Claire has a better hand than (the bearded man)_i)_j

i = local domain (the previous encounter with the 'bearded man')

j = main domain (the poker situation)

Ruth Kempson gives another example in which the shift of domain is of a temporal nature. Think of the contrast between (3) and (4):

(3) The hostages were systematically ill-treated

(4) The hostages were welcomed home by the President

In the domain framework, the contrast can be represented as follows:

(3*) ((the hostages)_i were systematically ill-treated)_i

(4*) ((the hostages)_i were welcomed home by the President)_j

i = the 'hostage' situation

j = the 'welcome' situation

As Kempson says,

The interpretation of (3) implies that throughout the time described as one in the past with respect to the time of the utterance, a set of people who were being held hostages were systematically ill-treated. In (4), however, at the point in time in the past at which a set of people is welcomed home by their president, they are transparently no longer hostages. But the assertion made by (4) may yet be true. The resource situation, s_0 , against which a definite NP is evaluated demonstrably need not be the same as the described situation s' (Kempson forthcoming, p.15).

In other words, (4) involves a temporal shift of domain. Such shifts are pervasive in everyday discourse.

A third example, due to Andrea Bonomi (1992: 100-101), is this:

Imagine this situation. I live in the countryside. My family includes five nice dogs. Unfortunately, some of them are extremely lively and, when they run off, the cats in the neighborhood are in serious danger. All of a sudden I hear some dog bark in the

distance. I look anxiously to my wife, who has just inspected the garden, but she reassures me. 'All the dogs are sleeping' she says.

Clearly, a particular set of dogs is selected in the situation, and all the dogs in that set are sleeping; yet, as Bonomi points out,

the situation [referred to] is large enough to involve dogs that are not sleeping. (Am I not hearing some dog bark?) But this characteristic does not lead me to consider the statement false. The point is that, besides this large segment of the world with respect to which the statement is evaluated, there is another one — used as a context — that I have in mind when I must select the relevant universe of quantification. And this is a smaller segment. As a result, it turns out to be true, with respect to the more inclusive situation, that all the dogs individuated by the smaller situation are sleeping.

Bonomi's example, therefore, can be represented as follows:

(5) ((All the dogs)_j are sleeping)_i

j = the smaller situation

i = the larger situation

Further examples of domain shift can be found in Cooper 1993 and in Stanley and Williamson 1995.

Now we see how the Lewis-McCawley counterexample can be disposed of. It is easy to imagine true instances of (L) analysable as involving a shift of domain, for example a temporal shift. Suppose we hear a dog barking at some distance. 'It's a dog', I say. But after some time, still concentrating on what I hear, I say 'The dog got in a fight with another dog'. Here again, time is the dimension along which the domain of discourse shifts. The situation we start with has one and only one dog, barking; then the situation changes. Even if the situation attended to does not really change (there were two dogs from the start, but we paid attention to only one of them) at least we re-analyse it and our representation changes. (The change occurs in 'processing time', in Langacker's terminology.) Be that as it may, the initial situation, with only one dog, gives way to a different situation with two dogs. The latter situation is what the global utterance talks about, but the description 'the dog' is still interpreted with respect to the earlier situation (described by the previous utterance, 'It's a dog'). (L) can therefore be analysed as follows:

(L*) ((The dog)_j got in a fight with another dog)_j

i = the earlier situation (described by 'It's a dog')

j = the new situation (described by (L))

Such an utterance would be true, if the dog which is a constituent of the initial situation *i* had got in a fight with another dog in the main situation *j*.

IV. Attributive uses of incomplete descriptions

According to Austinian semantics, when an incomplete description 'the F' is used, a situation is referred to which includes one and only one F. In the situation-theoretic framework set up in section 1, this means that there is one and only one object \underline{x} such that the situation contains the atomic fact that \underline{x} is F. Consider Donnellan's classic example: 'The murderer is insane', said in presence of Smith's dead body. This is an incomplete description, for there are many murderers in the world. In interpreting this description, we are supposed to focus on a restricted portion of the world, namely a situation including Smith's murderer and no other murderer: a situation \underline{s} such that, for some \underline{x} , \underline{s} contains the fact that \underline{x} killed Smith, and for all \underline{y} , if, for some \underline{z} , \underline{s} contains the fact that \underline{y} killed \underline{z} , then $\underline{y} = \underline{x}$.

As Soames 1986 pointed out, a problem for this view arises from the possibility of using incomplete descriptions attributively ('The murderer, whoever he is, is insane'). Let us assume that it is Robinson who actually killed Smith. According to Austinian semantics, the utterance refers to a situation \underline{s} which includes the fact that Robinson killed Smith. Had someone else killed Smith, the utterance would have referred to a different situation \underline{s}' including the fact that, say, Jones killed Smith. Since the situation referred to is part of the Austinian proposition expressed by the utterance, which Austinian proposition the utterance expresses depends upon who the referent is. But attributive uses of definite descriptions are distinguished from referential uses precisely by the fact that the identity of the referent does not affect what is said: the speaker merely says that whoever killed Smith is insane, and what is thus said is, or should be, unaffected by who the actual murderer is (Soames 1986: 356).

The difficulty can be summed up as follows. Real situations are individuated by the atomic facts they contain, e.g. the fact that Robinson killed Smith. But if the situation referred to, which is an aspect of the (Austinian) content of the utterance, includes the referent of the description (Jones or Robinson), we can no longer capture what distinguishes attributive uses from referential uses of descriptions, namely the fact that the actual referent is irrelevant to the content of what is said. To account for attributive uses of 'the murderer', we would have to include in the situation referred to the general fact that someone killed Smith, rather than the atomic fact that Robinson (or Jones) killed Smith.

To deal with this difficulty, can we replace real situations by types of situation, in the Austinian framework? I don't think we can. The gist of the Austinian account is the idea that an utterance refers to a situation s , and describes or (better) depicts a 'type' of situation T which serves to classify the situation referred to. There are two correlative distinctions here: between referring and depicting, and between real situations and types of situation. It seems that s , the situation referred to, can only be a real situation: in contrast to depiction, reference can only be to something which actually exists — some part of reality (e.g. some real situation). There is a tension between Austinian semantics thus understood and the version we would end up with if we replaced real situations by types of situation in order to account for the referential/attributive distinction.

The difficulty is more apparent than real, however. Let us assume that the situation referred to is indeed a real situation, individuated by the atomic facts it contains. Still, we can account for the referential/attributive distinction, for there is no reference without a mode of presentation of the reference. Once we construe the real situation as referred to, we must bring a mode of presentation of it into the picture, and this gives us all we need; for the mode of presentation under which a real situation is referred to determines a type of situation, viz. the type of situation the situation referred to is believed to instantiate. In other words, we need not replace real situations by types of situation; for whenever a real situation is referred to, a type of situation is eo ipso provided qua mode of presentation.

What I have just said about reference and modes of presentation may sound controversial. Is it not the case that, sometimes, we refer to objects 'directly'? But in the sense in which I am taking 'mode of presentation', it is, I think, uncontroversial that there is no reference without a mode of presentation. This claim is perfectly consistent with the theory of direct reference, for the modes of presentation under which the objects of (direct) reference are thought of are truth-conditionally irrelevant. 'Cicero is a great writer' and 'Tully is a great writer' are both true iff Cicero (= Tully) is a great writer, but there is a difference in cognitive significance, though no truth-conditional difference, between the two utterances. This theory, argued for at length elsewhere (Recanati 1993, 1995), reconciles the theory of direct reference with Fregean insights about cognitive significance and modes of presentation. On this view, the complete content of an utterance like 'Cicero is a great writer' is not a singular proposition but a quasi-singular proposition. Quasi-singular propositions are like singular propositions except that their first constituent is not an object or sequence of objects but an ordered pair (or a sequence of ordered pairs) consisting of an object and a truth-conditionally irrelevant mode of presentation. Thus 'Cicero is a great writer' expresses the following (quasi-singular) proposition:

<<Cicero, '**Cicero**'>, the property of being a great writer>

(As usual, I ignore modes of presentation of properties and relations.) The expression in bold type stands for a truth-conditionally irrelevant mode of presentation of Cicero. Such modes of presentation are dossiers of information, and they contain 'facts' (or pseudo-facts) about the referent, like the fact that Cicero was a Roman orator or the fact that Cicero is called 'Cicero'.⁵

The reference to situations in discourse can be construed in exactly the same fashion. The (real) situation referred to must be paired with a mode of presentation, in such a way that Austinian propositions themselves are quasi-singular propositions. Instead of:

<s, T>

where 's' is the situation referred to and 'T' the type of situation, an Austinian proposition has the following structure:

<<s, **m_s**>, T>

where 's' is a real situation, '**m_s**' is a mode of presentation of that situation, and 'T' the type of situation it is said to instantiate. The domain of discourse turns out to be a structured entity consisting of two components: an internal component (the mode of presentation) and an external component (the real situation).

In this framework we have three 'situations' to distinguish, corresponding to the referent, its mode of presentation, and what is said about it:

- s, the real situation referred to;
- T¹, the type of situation s is believed to instantiate, i.e. the way s is represented;
- T², the type of situation s is said to instantiate

It is the intermediate situation T¹, not the real situation s, which one characterizes when one says that (in attributive uses) the situation with respect to which 'the murderer' is interpreted includes the general fact that Smith was murdered, rather than the atomic fact that Robinson (or Jones) killed Smith. The real situation s referred to is

⁵ The referent is not determined 'satisfactorily' as whatever satisfies the content of the dossier, but 'relationally' as that from which most information (or misinformation) in the dossier derives. (The relational/satisfactorial terminology is borrowed from Bach 1987).

that of Smith's murder, and it does include the atomic fact that Robinson (or Jones) killed Smith; but it is mentally represented as a situation T^1 in which Smith was murdered by someone or other.

The difference between attributive and referential 'readings' turns out to be a cognitive difference at the mode of presentation level. In both cases, the speaker refers to a real situation \underline{s} (Smith's murder) and classifies this situation as a situation in which the murderer is insane. But the fact that Robinson killed Smith is an aspect of the situation of reference of which the speaker and hearer need not be aware. When the description is used attributively, the situation is represented merely as a situation in which someone or other killed Smith. When the description is used referentially, the murderer is identified: the situation of reference is represented as a situation in which Robinson killed Smith.

V. Domains and the mental encyclopedia

I have just analysed 'domains of discourse' as complex entities consisting of a real situation and a mode of presentation of that situation. In this section I want to say more about the 'internal' component of the domain of discourse.

Let me first introduce the notion of a 'mental encyclopedia'. By this I mean the individual's total belief-system conceived as a rich conceptual network. Each node in the network, each 'concept', can be seen as a labelled dossier of information; thus the concept of cat consists of a certain label ('cat') and a certain content (what the individual knows about cats). For the sake of simplicity, let us assume that the informational content of a concept is a set of propositions. Sperber and Wilson suggest that propositions themselves consist of concepts (i.e. labels) put together in a certain way. Thus the proposition that the cat is on the mat includes the concept 'cat', and the concept 'cat' itself includes in its content propositions about cats: that cats purr, that the neighbour's black cat is called 'Tobby', etc. This dual role of labels, which serve both as constituents of propositions and as 'addresses' of dossiers containing other propositions (Sperber and Wilson 1986: 86), determines accessibility relations within the network. Thus if I am currently entertaining the proposition 'The cat is on the mat', the concept of 'cat' thereby activated provides access to a number of propositions about cats, stored under the concept 'cat' in encyclopedic memory. From these propositions, in turn, one can go to other propositions by the same method (for example the proposition that the neighbour's black cat is called 'Tobby' will give access to propositions about the neighbour). The lesser the number of steps necessary to access a given representation from a given starting point, the more accessible that representation will be.

The 'modes of presentation' which I talked about in the previous section are themselves dossiers of information. Like my concept of cat, my concept of Cicero is a

certain dossier of information. When I think of Cicero, I think of him under a certain mode of presentation, that is, through a certain 'dossier' which is part and parcel of my mental encyclopedia. The modes of presentation of situations also are 'dossiers', like the modes of presentation for individuals. World War II is a certain situation or event, and my mental file about it is a typical dossier; it contains everything I know or believe about World War II.

Sperber and Wilson describe the process of interpreting an utterance as involving, as a preliminary step, a search through the mental encyclopedia for an appropriate 'context' of interpretation. The context, they say, is not 'given' but 'chosen' (Sperber and Wilson 1986: 132-50; see also Lewis 1973: 115 and Simon 1994: 9-11). On this view, interpreting an utterance involves selecting the particular regions of encyclopedic space which are relevant to the interpretation of the utterance. Dossiers, i.e. sets of propositions stored together in memory⁶, are typical regions in this sense.

In this framework, the particular region of encyclopedic space relative to which an utterance is interpreted corresponds to the internal component of the domain of discourse. The real situation referred to (external component) can itself be construed as the worldly correlate of the relevant region. The notion of a 'worldly correlate' ('w-correlate' for short) is not new; it is the familiar notion of 'reference', as applied to regions. Take the simplest case: a dossier about a particular individual, say my dossier about Cicero. This is a region of the encyclopedia. It corresponds to something in the world, namely to the referent of the dossier (Cicero himself). In the same way, I have a situational dossier about World War II; this dossier corresponds to a real situation in the world, namely World War II. Now suppose someone utters 'So many men died!', in a piece of conversation about World War II. This utterance will be interpreted with respect to a domain consisting of (i) an internal component: the speaker's or hearer's file about World War II, and (ii) an external component: World War II itself. The former is cognitively relevant; the latter is truth-conditionally relevant.

Situational dossiers are dossiers about situations, in the ordinary sense of the term; individual dossiers correspond to another, more typical sort of object, namely individuals (like Cicero). So far I have assumed that only situational dossiers can fill the mode of presentation role in a domain of discourse. This is close to the letter of Austinian semantics (both Austin and latter-day situation semanticists insist on the centrality of situations). But is there a deep reason for this restriction? Why must domains be situations as opposed to e.g. individual objects?

⁶ In Recanati 1993 I distinguished two sorts of dossiers: stable dossiers ('encyclopedia entries') and temporary files ('egocentric concepts'). Both are part of the 'mental encyclopedia' in the sense in which I am using this phrase here.

Situations are natural domains because situations are fact-supporting. What is going on at a certain spatio-temporal location is a situation in the ordinary sense. What happens to a particular individual at a particular time is also a situation in the ordinary sense. For every such situation (e.g. the situation of John at t), there is a set of facts, namely the facts which obtain in the situation (e.g. the fact that John is sleeping at t). Since every situation in the ordinary sense determines a situation in the technical sense — a set of facts — there has been a tendency to equate situations in the technical sense (domains) with situations in the ordinary sense. But if we consider that fact-supportiveness is criterial, we must recognize that individuals, too, are fact-supporting. An individual object determines a set of facts, as much as a situation does; and a piece of information can be about an individual object as well as about a situation in the ordinary sense.

Let us analyse the Austinian proposition expressed by an utterance like 'Mary is in love with John'. This utterance states a fact, viz. the fact that Mary is in love with John. Now what is the domain of discourse, in this example? It can certainly be a situation in the ordinary sense, e.g. the current situation in the philosophy class. The fact that Mary loves John can be fed into the dossier 'current situation in the class', which dossier is a situational dossier. This would be the case in the following dialogue:

- What's new in the philosophy class?
- Mary is in love with John!

But it seems to me that it is also possible to feed the fact that Mary loves John into the dossier for Mary; the fact may be a relevant fact about Mary, not (only) in the sense that it involves Mary as a constituent, but in the sense that the dossier 'Mary' is the most relevant one for interpreting the utterance. Thus we could have the following dialogue:

- How is Mary?
- She is in love with John!

As I have just said, the fact that Mary is a constituent of the proposition expressed is irrelevant. Mary could be the domain without being directly involved in the fact stated by the utterance. For example, suppose the conversation is about Mary qua patient in a hospital. The doctor has just examined her, and the following dialogue takes place:

- How is Mary?
- The liver is in good condition.

Such examples show that domains of discourse need not be situations in the ordinary sense, even if they always are 'situations' in the technical sense. The individual object Mary determines a set of facts, viz. the set of all the facts concerning Mary; hence it qualifies as domain. When it is interpreted with respect to that domain, the utterance 'Mary is in love with John' says that the set of facts concerning Mary includes the fact that she is in love with John.

I will henceforth assume that any kind of entity which we are capable of thinking about (i.e. any kind of entity such that we have the resources for 'filing' information concerning it) can serve as 'domain'. This notion of a domain goes well beyond the ordinary notion of a situation; it is about as inclusive as that of a 'region', and it closely corresponds to the traditional notion of 'topic' — that which an utterance talks about. (More precisely, it is the notion of a global domain which corresponds to that of topic.) Thus 'Mary is in love with John' can be uttered as a comment on John, on Mary, on love, etc. It is one of the ambitions of the theory of domains to account for the phenomena (topic, focus, rhematic structure, theticity and categoricity, etc.) documented in the vast literature on informational structure.

VI. Windows

A particularly interesting case of region is what I will call an 'intensional window', or 'window' for short. Think of Gone with the Wind. This is a book and we have an encyclopedia entry (a dossier) about it; it includes e.g. the information that the book was written by Margaret Mitchell and similar things. But a certain part of the entry concerns the content of the book, e.g. the story told by the author. This part of the content of the dossier constitutes a sub-structure, a 'window', in which we find propositions like: Rhett Butler loves Scarlett O'Hara, etc.

Windows are dossiers, but they are distinguished by the regional character of the propositions they contain. It is true in the window (i.e. in the novel) that Rhett Butler loves Scarlett O'Hara, but this is no longer true if we step out of the window by moving upward to the level where we have the proposition that Margaret Mitchell wrote the book; it is not true that Rhett Butler loves Scarlett O'Hara in the sense in which it is true that Margaret Mitchell wrote the book (see Lewis 1978).

Let us consider the issue more carefully. Dossiers often contain other dossiers. Thus the dossier for World War II contains a sub-dossier for the Stalingrad battle; the dossier 'recent murders in London' contains the sub-dossier 'Smith's murder', etc. Similarly, the dossier 'Gone with the wind' contains the sub-dossier 'content of the book', which includes propositions true in the book. Now let me define a proposition P as 'regional', with respect to a pair of dossiers <D, D'>, when the following conditions are satisfied:

(Regionality)

- (1) $w(D) \models P$ (P holds in the w-correlate of D)
 (2) $D < D'$ (D' contains D as a 'sub-dossier')
 (3) $w(D') \not\models P$ (P does not hold in D')

Regionality is related, but not identical, to the familiar notion of 'non-persistence' (Barwise and Cooper 1981). A non-persistent statement is often characterized as a statement which holds in a situation \underline{s} but possibly not in a larger situation \underline{s}' extending \underline{s} .⁷ For example, 'Everyone is asleep' is non-persistent, because it may be true in a local situation in which everyone is asleep, without being true in a larger situation including the former situation but also including a bunch of alert people. Similarly, 'The murderer is insane' is non-persistent because the unicity condition associated with the definite description may be satisfied in a local situation \underline{s} without being satisfied in a larger situation \underline{s}' extending \underline{s} . To see the connection between regionality and non-persistence, suppose 'The murderer is insane' belongs to the dossier 'Smith's murder' and holds in the w-correlate of that dossier; suppose further that (as I said earlier) the dossier 'Smith's murder' is included in the dossier 'recent murders in London'. As the w-correlate of the latter dossier is a larger situation than the w-correlate of the former, the non-persistence of 'The murder is insane' entails its regionality: the proposition 'The murderer is insane' does not hold in the (larger) situation which is the w-correlate of the dossier 'recent murders in London' — unless it turns out that there has been only one (recent) murderer in London.

Non-persistence is not the only source of regionality, however. Consider the proposition: 'There is a man who can fly'. It is persistent because the existential quantifier is persistent (if it is true in a situation \underline{s} that there is a man who can fly, then this is true in every situation \underline{s}' extending \underline{s}). Still, this proposition is regional with respect to the pair consisting of the sub-dossier 'content of the Superman story' and the overall dossier 'Superman story'. The latter contains all sorts of information about the story, concerning its author, the various versions that exist, and so forth — not to mention the fact that the story commonly serves as example in philosophical discussions of propositional attitudes. Information about the content of the story belongs to the dossier via the sub-dossier 'content of the Superman story', where we find the proposition 'There is a man who can fly' (together with a host of other propositions such as 'Lois Lane is Clark Kent's colleague', 'Clark Kent is mild-mannered', etc.). In this case also, though for totally different reasons, the three conditions for regionality are satisfied: (1) 'There is a man who can fly' belongs to the

⁷ There is a problem with this characterization, but I cannot go into this issue here.

sub-dossier, and holds in the situation the latter describes; (2) the sub-dossier is included in the overall dossier 'Superman story'; (3) still, the proposition 'There is a man who can fly' does not hold in the w-correlate of the latter, for it is not true that a man can fly in the sense in which it is true that such and such a book was published on such and such a year. The w-correlate of the overall dossier 'Superman story' is a portion of the real world where we find authors creating fictions, actors interpreting them, people enjoying them, and so forth. No flying man is to be found in this portion of the world (nor in any other portion of the real world) but only representations of a flying man. Regionality, here, arises not from the fact that the w-correlate of the overall dossier is a 'larger' situation than the w-correlate of the sub-dossier, but from the fact that the w-correlate of the sub-dossier belongs in a different world than the w-correlate of the overall dossier.

Windows are an invaluable tool for storing and manipulating information about representations. Representations have two aspects: they are objects like tables and chairs, and as such they belong to the real world; but they also have a content by virtue of which they represent the world as being a certain way, possibly distinct from the way it actually is. In order to deal with representations (in order to represent them) we need two levels, corresponding to the overall dossier and the window it contains. At the first level we store information concerning the representation as worldly object; at the second level we store information about the content of the representation, i.e. the world as represented.

Once windows are recognized as a particular case of region, it is tempting to enroll them into 'domains of discourse' and to say that the domain with respect to which a fictional statement such as 'Sherlock Holmes lives at 221B Baker Street' is evaluated is not (a portion of) the actual world but, rather, the world of Conan Doyle's novels. This would explain why this true statement seems 'cut off from its consequences', as Lewis says: we cannot infer that Holmes lives in a bank even if we know that the only building at 221B Baker Street is a bank. The explanation for this is the fact that one premiss ('Sherlock Holmes lives at 221B Baker Street') occurs in a window, and the other ('The only building at 221B Baker Street is a bank') outside it; the regionality of information in windows prevents it from mixing with information outside the window in the manner required for deriving consequences.

Natural though it is, this extension of the notion of domain of discourse to cases involving worlds other than the actual world (intensional domains, as I will call them) raises an obvious problem. I said that domains of discourse consist of situations we refer to (together with modes of presentation of the situations in question). I have so far assumed that, in order to be referred to, something must be real — it must exist in the actual world. (This follows from the theory of Direct Reference.) Extending the

framework to cases in which the domains are intensional and concern fictional or other non-actual worlds raises the problem of reference to the non-existent.

The notion of pretended reference provides a way out of the dilemma. When we 'refer' to something non-actual — something fictional, say — we pretend to refer. As the 'referent' does not exist, it cannot be really referred to; but it can be fictionally referred to (Lewis 1978, Walton 1990, Currie 1990). What is true of individuals is true of situations: in the same way as we can fictionally refer to Sherlock Holmes, we can fictionally refer to the situation described by the Superman story.

Pretended reference is not genuine reference; it is only 'quasi'-reference (Geach 1972: 161sq). Yet, as cognitive semanticists (e.g. Fauconnier 1985) keep stressing, the difference does not matter from a strictly linguistic point of view. For we use the same linguistic material, with the same linguistic meaning, whether we genuinely refer or only pretend to refer. This follows from the very notion of 'pretense'. Hence, linguistically, there is no difference between 'Holmes lives at 221B, Baker Street' and 'Recanati lives in Paris'. Nor is there any significant difference, from a linguistic point of view, between intensional domains and 'genuine' domains involving real situations. It is therefore legitimate, for linguistic purposes, to use a notion of 'domain of discourse' which encompasses both.

Intensional domains of discourse provides the key to the proper analysis of many puzzling phenomena such as:

- (i) irony, free indirect speech, semi-quotations and other 'interpretive' uses of language;
- (ii) alleged 'ambiguities' in intensional contexts (locutionary and attitude reports, descriptions of images and pictorial representations, etc.);
- (iii) Conditionals and suppositions, including counterfactual suppositions;
- (iv) fictional discourse;
- (v) meta-representations in general.

In what follows, I will give a few examples of the role intensional domains can play in semantic analysis. The examples I will give belong to the first two areas mentioned above: interpretive uses of language, and intensional 'ambiguities'. The first group of examples is intended to illustrate the naturalness and intuitive appeal of the theory; they show that sometimes what one says must be interpreted from someone else's point of view.⁸ The second group of examples is intended to show how useful intensional domains can be in dealing with some of the hard problems of semantics.

⁸ This phenomenon can be compared to the shift of deictic point of view described by e.g. Fillmore (1982). See also Bühler (1934) on 'displacement' and 'imagination-oriented deixis'.

VII. Intensional domains

Let us start with a simple example of 'free indirect speech' (from Recanati 1987: 62).
Suppose someone says:

(6) John is completely paranoid. Everybody wants to kill him, including his own mother!

Here, arguably, the sentence 'Everybody wants to kill him' is to be evaluated with respect to John's belief-world. The speaker somehow projects herself into John's mind. She does not really assert that everybody wants to kill John; it is his thought, not her own, that she expresses. The utterance is presented as 'true with respect to John's belief-world' — not as true with respect to the actual world:

(6*) (John is completely paranoid.)_i (Everybody wants to kill him.)_j
i = the actual world;
j = John's belief world

Counterfactual situations, belief worlds, fictional worlds and the like can play the role of the contextual domain of discourse and affect the interpretation of quantifiers. Consider the following example. Suppose Peter wrongly believes that Ann is your sister; Ann is coming over and I say to you, ironically:

(7) Hey, 'your sister' is coming over.

Although the utterance as a whole is about the actual world, the description 'your sister' is to be interpreted with respect to Peter's belief-world. In this context, the description refers to Ann, whereas in another context it would refer to Nicole, your actual sister. Here the domain of discourse shifts within the utterance, as in the examples above.

(7*) ((Your sister)_j is coming over)_i
i = the actual world;
j = Peter's belief world

There are examples of the same phenomenon occurring outside ironical speech (Recanati 1994:158-9). Imagine the following situation: John wrongly believes that Bush is the President of the USA. Knowing that Bush is in the next room, I say:

(8) If he goes in the next room, John will be surprised to meet the President of the USA.

Here the description 'The President of the USA' refers to Bush rather than to the actual President (Clinton) because it is intended to be interpreted with respect to a world, namely John's belief-world, in which Bush is the President of the USA. Thus the consequent of (8) can be analysed as follows:

(8*) (John will be surprised to meet (the President of the USA)_k)_j.

j = the hypothetical situation in which John goes in the next room

k = John's belief world

The consequent of (8) is presented as true in the hypothetical situation in which John goes in the next room; but the description 'the President of the USA' is evaluated in John's belief world rather than in the hypothetical situation. It is Bush (the person whom John thinks is President) whom John is surprised to meet, in the hypothetical situation in which John goes in the next room.

A fuller analysis of (8) would have to take into account the fact that, by describing a hypothetical situation *j*, a hypothetical statement like (8) describes the actual world *i* as a world in which that situation is possible. The analysis therefore becomes:

(8**) (If he goes in the next room, (John will be surprised to meet (the President of the USA)_k)_j)_i.

i = the actual world

j = the hypothetical situation in which John goes in the next room

k = John's belief world

A further complication stems from the fact that (8) involves an attitude verb: 'to be surprised that...' Now attitude reports in general can be analysed as involving a mandatory domain shift. It is part of the semantics of 'believe' and other attitude verbs that their complement (e.g. the embedded sentence) is to be interpreted with respect to the attitude-world of the ascriber. An utterance like 'John believes that grass is green' therefore involves two domains, one (the 'parent space', in Fauconnier's terminology) for the interpretation of the global utterance, and another one (the ascriber's belief world) for the interpretation of the complement clause:

(John believes that (grass is green))_i

i = parent space (e.g. the actual world).

j = belief world

Applied to (8), this yields:

(8***) (If he goes in the next room, (John will be surprised that (John meets the President of the USA)_k)_j)_i.

i = the actual world

j = the hypothetical situation in which John goes in the next room

k = John's belief world

This is not an alternative to (8**); for the description itself must be indexed, as it can be interpreted either with respect to *k* ('opaque' reading, with the description denoting Bush) or with respect to *j* ('transparent' reading, with the description denoting Clinton). The two readings can be represented as follows:

(8***a) (If he goes in the next room, (John will be surprised that (John meets (the President of the USA)_k)_k)_j)_i.

(8***b) (If he goes in the next room, (John will be surprised that (John meets (the President of the USA)_j)_k)_j)_i.

i = the actual world

j = the hypothetical situation in which John goes in the next room

k = John's belief world

As Fauconnier (1985) has shown, it is one of the main advantages of this type of framework that it permits to handle intensional ambiguities (such as transparent/opaque or specific/nonspecific) easily and straightforwardly. This is particularly clear when we consider phenomena which raise difficulties for classical analyses in terms of scope. Consider (9):

(9) Teri believes that a spy is following her

This example exhibits the specific/non-specific 'ambiguity'. According to the traditional scope analysis, (9) can be analyzed either as:

(9') $(\exists x: \text{spy } x) (\text{Teri believes } (x \text{ is following Teri}))$

or as:

(9") Teri believes $((\exists x: \text{spy } x) (x \text{ is following Teri}))$.

In the domain framework another analysis of the ambiguity is available — one that does not appeal to scope. In 'Teri believes that a spy is following her', there are two worlds at stake, the actual world (where we find Teri and her beliefs) and the world of her beliefs (which the complement clause describes). The embedded sentence is evaluated with respect to the belief world: in Teri's belief world j , a spy is following her.

(9*) (Teri believes that (a spy is following her)) $_j$) $_i$.

The ambiguity of the indefinite description 'a spy' stems from the fact that it can be interpreted with respect to either of the two domains at stake, the belief world j and the parent space i . The two readings are:

(9**) (Teri believes that ((a spy) $_i$ is following her)) $_j$) $_i$.

(9***) (Teri believes that ((a spy) $_j$ is following her)) $_j$) $_i$.

In (9***) there is a double 'domain shift'. A first domain shift is linguistically encoded; as I said earlier, it is part of the meaning of 'believe' that the embedded sentence must be interpreted with respect to a local domain (the belief world) distinct from the domain with respect to which the global utterance is evaluated (e.g. the actual world). But the description, in the sentence, can be evaluated with respect to i , despite the fact that the sentence where it occurs is evaluated with respect to j . This is another instance of domain shift. On this analysis, the 'ambiguity' of the description comes from the fact that two domains are contextually available for its evaluation. It's not really an ambiguity, in the technical sense, but rather a particular form of semantic indeterminacy. While it is semantically mandatory for the embedded sentence to be evaluated with respect to the ascriber's belief world, the description's domain of evaluation is not semantically constrained in this manner.

So we have two ways of analysing the specific/nonspecific ambiguity: the scope analysis and the domain analysis. At this stage they need not be considered as alternative analyses; perhaps it is the scope of the description which determines which domain is relevant. However, the scope analysis is much more restricted in its application than the domain analysis. For the specific/non-specific ambiguity can be found also in sentences without sentential complement, like:

(10) Teri painted a mountain.

On one reading, there is a particular mountain (say, Mount Everest) which Teri represented in her painting; on the other reading there is no particular mountain — Teri's painting merely happens to be a 'mountain-representation' (as Goodman would say). Now the scope analysis does not apply to examples like (10). We can represent the specific reading as

(10') $(\exists x: \text{mountain } x) (\text{Teri painted } x)$

but how shall we represent the non-specific reading? We can do so only by doing 'some violence to both logic and grammar', as Quine admits (Quine 1956: 186). Thus we have to equate the painting of an object (the mountain) with the painting of a state of affairs (viz. the fact that there is a mountain):

(10'') Teri painted $(\exists x: \text{mountain } x)$

The same problem arises for Quine's example, 'I want a sloop' (Quine 1956). In spite of Quine's best efforts, the scope analysis is clumsy and implausible when applied to such examples (see Kaplan 1986: 266-8).

This problem does not exist for the analysis in terms of contextual domains. In 'Teri painted a mountain', 'a mountain' can be interpreted with respect to the actual world (there is a real mountain, which Teri represented in her painting) or with respect to the world of the painting (in the painting, there is a mountain):

(10*) $(\text{Teri painted } (a \text{ mountain})_i)_i$

(10**) $(\text{Teri painted } (a \text{ mountain})_j)_i$

The second phenomenon which the domain framework permits to handle is that of 'intermediate' readings. Consider (11), from Fauconnier 1985:

(11) Tom is looking for a fish.

On the standard 'specific' reading (11) means: there is a fish Tom is looking for. On the 'non-specific' reading, Tom is looking for some fish or other. But there is an intermediate reading: Tom is looking for the fish Max claimed he saw (but which does not exist in reality). Even if the first two readings could be handled in terms of scope, à la Quine, the third reading would still raise a problem — a problem which vanishes in the domain framework (Fauconnier 1985). What characterizes the intermediate reading

in this case is the fact that the indefinite description is interpreted with respect to Max's belief-world.

(11*) (Tom is looking for (a fish)_i)_i.

(11**) (Tom is looking for (a fish)_j)_i.

(11***) (Tom is looking for (a fish)_k)_i.

i = the actual world;

*j = Tom's target situation (= the situation he is trying to bring about:
a situation where he finds a fish);*

k = Max's belief world

I refer the reader to Fauconnier 1985 for further details and examples.

VIII. Conclusion

In the first part of this paper I have presented a defense of the Austinian semantic approach to incomplete quantifiers and similar phenomena (section 2-4). It is part of my defence of Austinian semantics that it incorporates a cognitive dimension (section 4). This cognitive dimension makes it possible to connect Austinian semantics to various cognitive theories of discourse interpretation. In the second part of the paper (sections 5-7), I have established connections between Austinian semantics and four particular theories:

- the theory of reference and modes of presentation in terms of information files (see e.g. Perry 1992),
- the theory of discourse interpretation as involving a process of context selection (see Sperber and Wilson 1986),
- the theory of informational structure (for a survey, see Lambrecht 1994),
- the theory of mental spaces (Fauconnier 1985).

The structural similarity between the theory of mental spaces and Austinian semantics, as elaborated by situation theorists (Barwise and Perry 1983, Barwise and Etchemendy 1987, Barwise 1989), is particularly striking. The main difference is that mental space theorists, and cognitive semanticists more generally, weaken the notion of reference in such a way that mental entities such as Sherlock Holmes can be 'referred to' (see Bencivenga 1983 on cognitive spaces and the 'epistemic' theory of reference). This weakening is unwelcome from the point of view of direct reference theory: the latter is a relational theory, which entails that the objects of reference must belong to @ (insofar as the users of the language themselves belong to @). This is why situation theory and

the theory of mental spaces have remained entirely separated, without any attempt at integration, despite their striking similarity. As I have shown, this apparent incompatibility can be overcome by appealing to the notion of 'pretense', which has become central in various areas of philosophy.

Beside the cognitive emphasis, there is another side to my defense of Austinian semantics: it pertains to the distinction between local domains and global domains. I think the theory of reference is entirely subsumed under the theory of local domains, and the theory of topics under the theory of global domains. I could not go into those developments in this paper, but I have provided many examples of utterances which can be usefully analysed by appealing to the distinction between local domains and main domains.

References

- Austin, J. L. (1970), *Philosophical Papers*. Second edition, Oxford: Clarendon Press.
- Bach, K. (1987), *Thought and Reference*. Oxford: Clarendon Press.
- Bach, K. (1994), Conversational Implicature. *Mind and Language* 9:124-62.
- Bach, K. & Harnish, M. (1979), *Linguistic Communication and Speech Acts*. Cambridge, Mass.: MIT Press.
- Barwise, J. (1989), *The Situation in Logic*, CSLI Lecture Notes 17
- Barwise, J. and Cooper, R. (1981), Generalized Quantifiers and Natural Language. *Linguistics and Philosophy* 4: 159-219.
- Barwise, J. and Etchemendy, J. (1987), *The Liar*. New York: Oxford University Press.
- Barwise J. and Perry, J. (1983), *Situations and Attitudes*. Cambridge Mass: MIT Press.
- Bencivenga, E. (1983), An Epistemic Theory of Reference. *Journal of Philosophy* 80: 785-805.
- Bonomi, A. (1992), Persistent Truth. *Intellectica* 13-14: 79-103.
- Bühler, K. (1934), *Sprachtheorie*. Jena: Gustav Fischer Verlag. [Translated in English as *Theory of Language*, Amsterdam: John Benjamins, 1990.]
- Carston, R. (1988), Implicature, Explicature, and Truth-Theoretic Semantics. In R. Kempson (ed.), *Mental Representations: The Interface Between Language and Reality*, Cambridge: Cambridge University Press, p.155-81.
- Cooper, R. (1993), Generalized Quantifiers and Resource Situations. In P. Aczel et al. (eds.), *Situation Theory and Its Applications* 3, Stanford: CSLI.
- Currie, G. (1990), *The Nature of Fiction*. Cambridge: Cambridge University Press.
- Fauconnier, G. (1985), *Mental Spaces: Aspects of Meaning Construction in Natural Language*. Cambridge, Mass.:MIT Press/Bradford Books.

Fillmore, C. (1982), Towards a Descriptive Framework for Spatial Deixis. In R. Jarvella and W. Klein (ed.), *Speech, Place and Action*, John Wiley and Sons, pp. 31-59.

Geach, P. (1972), *Logic Matters*. Oxford: Basil Blackwell.

Kaplan, D. (1986), Opacity. In L. Hahn and P. A. Schilpp (eds.), *The Philosophy of W. V. Quine*, La Salle, Illinois: Open Court, pp. 229-89.

Kempson, R. (forthcoming), *Language and Cognition: A Licensing Grammar*.

Kuroda, S.Y. (1982), Indexed Predicate Calculus. *Journal of Semantics* 1: 43-59.

Lambrecht, K. (1994), *Information Structure and Sentence Form*. Cambridge: Cambridge University Press.

Lewis, D. (1973), *Counterfactuals*. Oxford: Basil Blackwell.

Lewis, D. (1978), Truth in Fiction. *American Philosophical Quarterly* 15: 37-46.

Lewis, D. (1979), Scorekeeping in a Language Game. Reprinted in D. Lewis, *Philosophical Papers*, vol. 1, New York: Oxford University Press, 1983, pp. 233-49.

Neale, S. (1990), *Descriptions*. Cambridge, Mass: MIT Press/Bradford Books.

Perry, J. (1986), From Worlds to Situations. *Journal of Philosophical Logic* 15:83-107.

Perry, J. (1992), *The Problem of the Essential Indexical and Other Essays*, New York Oxford University Press.

Quine, W. v. O. (1956), Quantifiers and Propositional Attitudes. In W. v. O. Quine, *The Ways of Paradox and Other Essays*, revised edition, Cambridge, Mass: Harvard University Press, 1976, pp. 185-196.

Recanati, F. (1987), Contextual Dependence and Definite Descriptions. *Proceedings of the Aristotelian Society* 87:57-73.

Recanati, F. (1989a), Referential/Attributive: A Contextualist Proposal. *Philosophical Studies* 56: 217-49.

Recanati, F. (1989b), The Pragmatics of What is Said. *Mind and Language* 4: 295-329.

Recanati, F. (1993), *Direct Reference: From Language to Thought*. Oxford: Basil Blackwell.

Recanati, F. (1994), Contextualism and Anti-Contextualism in the Philosophy of Language. In S. Tsohatzidis (ed.), *Foundations of Speech Act Theory*, London: Routledge, pp. 156-66.

Recanati, F. (1995), Quasi-singular Propositions: The Semantics of Belief Reports. *Proceedings of the Aristotelian Society*, Supplementary Volume 69, pp. 00-00.

Reimer, M. (1992), Incomplete Descriptions. *Erkenntnis* 37: 347-63.

Salmon, N. (1981), *Reference and Essence*. Princeton: Princeton University Press.

Salmon, N. (1982), Assertion and Incomplete Definite Descriptions. *Philosophical Studies* 42:37-45.

Simon, H. (1994), Literary Criticism: A Cognitive Approach. *Stanford Humanities Review Supplement* 4(1): 1-26.

Soames, S. (1986), Incomplete Definite Descriptions. *Notre Dame Journal of Formal Logic* 27: 349-75.

Sperber D. and Wilson, D. (1986), *Relevance: Communication and Cognition*. Oxford: Blackwell.

Stalnaker, R. (1974), Pragmatic Presuppositions. Reprinted in S. Davis (ed.), *Pragmatics: a Reader*, New York: Oxford University Press, pp. 471-82.

Stalnaker, R. (1978), Assertion. *Syntax and Semantics* 9: 315-22.

Stanley, J. and Williamson, T. (1995), Quantifiers and Context-Dependence. *Analyst* 12.1, preprint 10.

Walton, K. (1990), *Mimesis as Make-Believe*. Cambridge, Mass: Harvard University Press.

Wettstein, H. (1981), Demonstrative Reference and Definite Descriptions.
Philosophical Studies 40: 241-57.