

Context and Information in Dynamic Semantics

Jeroen A.G. Groenendijk* and Martin J.B. Stokhof*

1 Introduction

A popular view on meaning holds that a distinction needs to be made between the semantic content of an utterance and its pragmatic, communicative function, and that both have to be taken into account in order to get a full specification of meaning. This implies that 'semantics' cannot be identified with 'theory of meaning, the latter being a broader notion. Another widespread belief about meaning is that the meaning of an utterance consists in the way in which it changes the context. However, the notion of the context of an utterance is a notoriously vague one, that has been taken to comprise a wide variety of things. Here, a third popular view on meaning is of help, which is that meaning is intimately related to information and information exchange. From this perspective on communication, a context can be looked upon as consisting of the (partial) information of the speech participants, about the world, about the information of each other, and so on.

One way of combining the above views in an overall approach to meaning is to reconstruct the notion of a communicative function as one that operates on the semantic content of a sentence and an information state to produce a new information state. So, the meaning of an utterance with a particular communicative function and a particular semantic content, i.e. a particular communicative act, is characterized by the change in information state that it brings about.

This view we call 'the pragmatic interpretation approach'. It underlies the analysis that Bunt has proposed for the special type of conversations he calls 'information dialogues'. To be sure, Bunt has added a lot more details to this general scheme, but we will not go into those here. Our primary concern in this paper is an assessment of the relationship between this approach of pragmatic interpretation and the kind of 'dynamic semantics'

* ITLI/Department of Philosophy & Department of Computer Science, Faculty of Arts, University of Amsterdam, Grimborgwal 10, 1012 GA Amsterdam, The Netherlands.

that we have developed as a method for the partial analysis of semantic content. In his paper Bunt (*this volume*) proposes to make use of the system of dynamic predicate logic of Groenendijk and Stokhof (1987) in extending his original analysis of information dialogues, which was carried out along the lines of the former approach. In the course of doing so, he criticizes dynamic predicate logic at several points, and suggests some improvements. We will argue below that Bunt's criticisms and the alternatives he suggests do not fully appreciate the true difference between the pragmatic interpretation approach and that of dynamic semantics. Superficially, both deal with context, information, and updating of information, but a closer scrutiny of the principles underlying dynamic semantics will reveal that they do so in different ways, and with different goals.

For a proper understanding of these differences, it is important to get a clear picture of the various roles that the notion of context plays, or is taken to be playing, in the theory of meaning, more in particular of the relationship between context and information. Therefore we will start with a short, general discussion of these matters, and turn to an exposition of dynamic semantics only after that. We follow up with a discussion of Bunt's criticisms and alternatives, making some suggestions ourselves as to how to overcome the limitations of the particular system of dynamic semantics that Bunt signals. Some of these extensions, we will argue, will also solve the problems that probably prompted Bunt to take up the issue. The comparison with the pragmatic interpretation approach will be largely implicit, but we trust that the reader who has read Bunt's paper, will have no trouble in spotting the differences.

2 Context and information

In logical semantics, the recursively defined interpretation of expressions characteristically takes place with respect to one or more parameters. One parameter that will always be present (and for that reason is often suppressed) is a model.

The model contains all the elements that are necessary to assign meanings to all the expressions of the logical language in question. It presents the ingredients of an ontology. A domain of discourse is nearly always among these ingredients, and certain other sets of primitive objects, such as possible worlds, or moments of time, and relations on them, such as accessibility of possible worlds, an earlier-than relation between moments of time, may

inhabit the model as well. And there is a function which interprets the basic expressions of the language in terms of the ontology the system presents and which forms the starting point of the general recursive definition of the interpretation of all expressions of the language.

Against the background that the model provides, we often encounter one or more different parameters which can be said to provide a context of interpretation. These contextual parameters constitute a wide variety of elements. One that like a model almost goes unnoticed is the assignment of values to variables with respect to which we evaluate expressions in standard predicate logic. Other more conspicuous examples are the possible worlds which function as contextual parameters in modal logic, the moments of time in tense logic, speaker and addressee in a logic for indexical expressions, and so on. By the way, calling them all 'contextual parameters' should not obscure the differences that exist between them. We will come back to this later.

This notion of context as a (set of) parameter(s) is essentially a feature of logical interpretation. In a sense, it may be said to provide the necessary information for the interpretation process to be carried out. However, we must not take the phrase 'information' too seriously here. For as such, this notion of context has little to do with, e.g., the information available to speech participants, or with the information carried by an utterance. As we will see, it plays, or may be made to play, a role there, too, but it cannot be simply equated with information in this sense of the word.

That this is so can perhaps be appreciated by the following observation. A characteristic feature of information in the second, more colloquial sense of the word is that more often than not it is only partial: complete information about the world, about the information of others, and so on, is what we strive after, but unfortunately never seem to be able to obtain. This is a definite characteristic that contextual parameters lack: assignments, possible worlds, moments of time, etcetera, are complete, total objects. E.g., an assignment in standard predicate logic is a total function from the variables of the language to objects in the domain of discourse. Possible worlds, too, are conceived of as total objects: they verify or falsify all formulas. And the same holds for moments of time, and so on.

If we want to represent information about the context, i.e. information about the values of the variables, information about the world, etc., the simplest way to go about is to take a subset of the entire set of contexts, and regard the elements of this subset as those possibilities that are still open according to our information. Adding information amounts to eliminating

certain possibilities. As our information grows, the information set shrinks. Having no information means leaving open all possibilities, i.e. having the entire set as our information set. Complete information will correspond to a unit set. And the empty set represents the deplorable situation where we have added too much to our information causing it to become inconsistent. Although this 'elimination approach' to the representation of information has some shortcomings, some of which are overcome by refinements which have been proposed, it is by far the most common; for our purposes in this paper we need not consider more complex theories.

So, if we want to think of a context in terms of available information about something (the values of variables, the world), this is not a context in the sense of a contextual parameter of our interpretation function, though it can be represented in terms of such. This leaves us with two different, though related notions of a context. It should also be clear that what kind of information we can take a context in the second sense to consist of, depends on the logical system we are discussing. In ordinary predicate logic for example, the only notion of information that is present is that of information about the values of variables. But if the interpretation function in some logical system has more parameters, then this system provides us with more than one notion of information. And it should be stressed that as the contextual parameters from which they are built may play different roles, these notions of information may be relevant at different levels of interpretation.

One might be inclined to object to the above argument against the equation of a context as parameter of interpretation with a context as information, by pointing out the existence of partial analogues of the total notions involved. Could we not conceive of partial assignment functions from variables to objects, and do we not have situations as a kind of 'pieces of the world'? This observation is correct, of course, but misses the point.

Just using partial assignment functions, or 'partial' worlds, instead of total ones, is not sufficient to get an account of the partiality of information. (We have seen above that it is not necessary either, since partiality can be represented by taking a set of alternative complete possibilities.) This can be argued as follows. For example, a partial assignment will leave the value of certain variables undefined. If the value of a certain variable is undefined, this can be viewed as a representation of a situation where we have no information at all about its value. If it does assign a certain value to it, this means that we are completely informed about the value of this particular variable. But what about the case where we have partial information about the value of a particular variable, for example that it is either this object

or that, but we do not know which? One single partial assignment cannot account for this situation. For any particular variable it only gives us the choice between total information or no information at all. Partial information can still be accounted for only by having a set of alternatives at hand. A similar reasoning will show that a situation, or other kind of 'part of the world' which evaluates only part of all formulas, by itself does not represent partiality of information about the world.

Of course, we do not want to claim that partializing contextual parameters is of no use in semantics, on the contrary. It plays an important role, e.g. in the semantic analysis of vague predicates, presuppositions, semantic paradoxes, and so on. With respect to the representation of information, we see (at least) two functions it may perform. First, a partialized parameter can be regarded as a convenient 'abbreviatory' representation for a situation of having no information about those objects for which it is undefined. Second, and more important, partialized parameters may be used to enrich the elimination perspective indicated above, by allowing it to make a distinction between a situation of having no information about an object, but being in a sense acquainted with it, and a situation of not being aware of the object's existence at all. However, it should be noticed, first, that the latter distinction just comes on top of the distinctions which the elimination perspective allows us to make, and, second, that a proper representation of this distinction, too, calls for the use of sets of partialized parameters.

Recapitulating our discussion so far, we have seen that (at least) two notions of context need to be distinguished: contexts as parameters of interpretation, and contexts as available information about such parameters. Examples are variable assignments, and sets of such, which represent information about the values of variables, or possible worlds, and sets of those (often also called 'propositions') which embody information about the (actual) world. According to popular belief, both notions of context play a role in interpretation, i.e. in establishing meaning. In the introduction we mentioned the view that meaning consists of semantic content and communicative function, and which reconstructs the latter as a function from information states to information states. The first notion of context comes into play in establishing semantic content, whereas the second one does its job by providing the means to represent information states.

As for the first, we have remarked above that various contextual parameters may differ in the precise role they play as arguments in the interpretation process. Let us limit ourselves to the distinction between assignments and possible worlds (truly 'indexical' parameters form a third category).

If we are only interested in sentences, i.e. closed formulas (a reasonable limitation from a natural language point of view), the role of assignments is only local: during the interpretation of a formula they allow us to keep track of the values that have been assigned to the various variables that we have encountered. In the end the 'information they contain', to use that dangerous metaphor, is discarded; it is of no further use: a sentence is true with respect to all assignments, or with respect to none. The assignments we use are like the books we keep until our accounts are definitely settled and we throw them away. Possible worlds, however, perform a different function: they parametrize the interpretation of the descriptive constants of the language, allowing us to define a notion of 'intension of an expression' that captures semantic content at the object language level. Hence, its role is global, rather than local: the interpretation of a fully interpreted sentence makes essential reference to it. As for the second notion of context, that of information state, we may observe that the difference in function of contextual parameters just observed, carries over. Information about the values of variables plays a local, more or less auxiliary role, whereas information about the world is global: the latter is what we strive for, the former is subsidiary to that end.

The relationship between the roles that the notions of context play in an account of meaning is hierarchical: contextual parameters are used in determining semantic content, and do not refer to context as information states. The latter give body to a notion of 'updating' of information states, which is defined in terms of semantic content.

We think that a third notion of context in the process of establishing meaning needs to be distinguished. And this is where dynamic semantics comes in. This third notion is a notion of update that plays a role within the determination of semantic content. It differs from context as an information state about the world, and it also differs from the kind of contextual parameters of the interpretation function that we have been concerned with above, though it bears resemblances to both. It plays a role in building the semantic content of discourses or texts in a dynamic way. If we want to interpret a text or discourse, and update our initial information with the information it carries, we need to establish its semantic content. Although, in the end, that is when the entire text or discourse is finished, its semantic content can be represented by a simple proposition. This essentially 'static' notion of semantic content cannot be used while the interpretation process is still going on, i.e. while information is still coming in. Here, we need an essentially 'dynamic' notion of content.

One reason for this, and the only one that will be discussed in this paper, is the existence of cross-sentential anaphoric links. In order to account for such links, we need to be able to keep track of the values of variables and the properties assigned to them, established by the interpretation of subsentential expressions, such as noun phrases. We need such information to be able to determine antecedents for anaphora in sentences to come. And in fact, it can be argued that certain phenomena concerning anaphora already call for this dynamic notion of content at the sentential level. Some of the relevant phenomena will be discussed in the next section. Here, it suffices to notice that this notion of a context is a separate one.

3 Anaphora and compositionality

If we use standard first-order predicate logic (henceforth, PL) in translating a sentence or discourse in natural language, anaphoric pronouns will turn up as bound variables. In many cases this means that in order to arrive at formulas which express the right meaning (i.e. which are good translations) we have to be pretty inventive, and should not pay too much attention to the way in which the natural language sentence or discourse is built up.

Let us illustrate this with two simple examples, which nevertheless are representative for the kind of problems we meet:

- (1) A man walks in the park. He whistles
- (2) Every farmer who owns a donkey, beats it

In order for the *he* in the second sentence of (1) to be anaphorically linked to *a man* in the first sentence, we have to give an existential quantifier wide scope over the conjunction of the two sentences involved. Doing so, we arrive at (1a):

$$(1a) \exists x[\text{man}(x) \wedge \text{walk_in_the_park}(x) \wedge \text{whistle}(x)]$$

Now notice that the translation of the first sentence in (1), which would be $\exists x[\text{man}(x) \wedge \text{walk_in_the_park}(x)]$, does not occur as a subformula in (1a). Apparently, we do not get from (1) to (1a) in a step-by-step, i.e. compositional way. If we did, we would translate (1) as (1b):

$$(1b) \exists x[\text{man}(x) \wedge \text{walk_in_the_park}(x)] \wedge \text{whistle}(x)$$

But this is not a proper translation of (1), at least not in standard predicate logic, since in (1*b*) the last occurrence of the variable x is not bound by the existential quantifier, and hence the anaphoric link in (1) is not accounted for.

However, suppose we could interpret (1*b*) in such a way that it would be equivalent with (1). Evidently, (1*b*) would be preferred to (1*a*) as a translation of (1) since it would be the result of a compositional, 'on-line' procedure.

Turning to example (2), we observe that its proper translation into predicate logic is (2*a*):

$$(2a) \forall x \forall y [[\text{farmer}(x) \wedge \text{donkey}(y) \wedge \text{own}(x, y)] \longrightarrow \text{beat}(x, y)]$$

This case is more dramatic than the previous one. Although (2) contains an indefinite term, which normally translates as an existentially quantified phrase, we need universal quantification to account for its meaning in this kind of example. And notice moreover that the corresponding universal quantifier $\forall y$ has to be given wide scope over the entire formula, whereas the indefinite term *a donkey* in (2) to which it corresponds, appears way inside the relative clause attached to the subject term *every farmer*. If we use standard logic as our means to represent semantic content, this kind of example prevents us from uniformly translating indefinite terms as existentially quantified phrases. Again, this constitutes a breach of the principle of compositionality, which is not only intuitively appealing, but also theoretically parsimonious and computationally plausible. From a compositional view, a translation like (2*b*) is to be preferred:

$$(2b) \forall x [\text{farmer}(x) \wedge \exists y [\text{donkey}(y) \wedge \text{own}(x, y)]] \longrightarrow \text{beat}(x, y)]$$

But then again, (2*b*) does not have the proper meaning, since the variable y in the consequent is not bound by the existential quantifier in the antecedent. Hence, (2*b*) is not equivalent with (2*a*), at least not in PL.

Examples like (1) and (2) have been treated successfully in Discourse Representation Theory (see Kamp, 1981), but at a cost: the problem of providing a compositional translation is not really solved, and DRT uses a rather non-orthodox logical language. In DRT (1) would be represented as (1*c*), and (2) as (2*c*):

- (1c) $[x][\text{man}(x), \text{walk_in_the_park}(x), \text{whistle}(x)]$
 (2c) $[] [x, y][\text{farmer}(x), \text{donkey}(y), \text{own}(x, y)] \longrightarrow [] [\text{beat}(x, y)]$

We will not go into the semantics of these ‘discourse representation structures’ here, but just notice that (1c) and (2c) have essentially the same truth conditions as (1a) and (2a) respectively. It should be noted, however, that the interpretation of these structures do not really present a compositional interpretation of the corresponding sentences. The structure of (1c) is essentially that of (1a) and not that of (1b). And in (2c) there is no separate representation of the relative clause *who owns a donkey* in (2). These issues, and the exact role of compositionality, will be discussed extensively in Groenendijk and Stokhof (*forthcoming*; see also Groenendijk and Stokhof, 1987).

In the latter paper we have given an alternative account of the phenomena indicated above, by replacing the standard semantics of the language of first order predicate logic by a dynamic semantics, inspired by systems of dynamic logic, used in the semantics of programming language (see Harel, 1984). The resulting system (henceforth referred to as DPL) constitutes an improvement over DRT in the following sense: it gives a more compositional treatment of the relevant phenomena, even though the syntax of the language used, being that of standard predicate logic, is an orthodox one. For the sentences (1) and (2), DPL gives (1b) and (2b) as representations, which, as we already remarked above, are to be preferred from a compositional and computational point of view. The new dynamic semantics ensures that (1b) and (2b) come out with the same truth conditions as (1a) and (2a) have in PL.

In the next section, we will give a sketch of what the dynamic semantics of DPL amounts to. We will do so by informally discussing the interpretation procedure that is involved, from the perspective of context, information and update, and by comparing it with the familiar semantics of PL. In doing so, we will also give some body to the third, ‘dynamic’ notion of context, announced in the previous section.

4 Context and information in DPL

Whether we are talking about PL or DPL, both being extensional logical systems, the notion of context is restricted to that of a (complete) assignment of values to variables, and hence, as we have argued in section 2, information

at any stage in the interpretation process is information about the values variables can have at that stage.

In PL, the semantic content of a formula in a model can be identified with a set of assignments, those assignments which satisfy the formula. For any closed formula this is either the set of all assignments, in case the formula is true, or the empty set, in case the formula is false. Viewed from the perspective of information, the set of all assignments represents the situation of having no information at all (all possibilities are still open) and the empty set with that of having inconsistent information (all possibilities are excluded). Non-empty, real subsets of the set of all assignments are encountered only as interpretations of formulas with free occurrences of variables. In the interpretation of sentences (closed formulas) such sets play a role only at intermediary stages. Once a closed formula is interpreted entirely, information about the values of variables has become irrelevant. At that final stage we know what the semantic content of the formula in the model is, viz. we know whether it is true or false.

Let us now turn to update of information in PL. The notion of update is closely connected with that of conjunction, so we first take a look at that. Let $|\phi|$ denote the set of assignments which satisfy a formula ϕ . The set of assignments which satisfy a conjunction $\phi \wedge \psi$ is the intersection $|\phi| \cap |\psi|$. From an update perspective we get the following. Let I be the set of assignments still possible according to our information at some particular stage of the interpretation process. If we update our information I with ϕ , then our new information is $I \cap |\phi|$. In this update procedure, the information which is already available, does not play a role in the interpretation of ϕ . Updating is not a local process, incorporated in the recursive notion of interpretation itself, but comes on top of it in a global way. We could define a wider notion of meaning of ϕ , denoted by $\|\phi\|$, which captures the updating aspect. Meaning in this sense would be a function from information to information, defined as $\|\phi\| = \lambda X : X \cap |\phi|$. If we apply this function $\|\phi\|$ to any information set I , we arrive at $\|\phi\|(I)$, which is the new information set $I' = I \cap |\phi|$ which results after updating our information with ϕ .

We would use this richer notion of meaning as our basic recursive notion, but this would not give us anything new. We would not really have enriched our interpretation, after all. The following considerations may serve to illustrate this. If $\|\bullet\|$ would be our recursive notion, we could now define the interpretation of conjunction as follows: $\|\phi \wedge \psi\| = \lambda X : \|\psi\|(\|\phi\|(X))$. Applied to an information set I we would get $\|\psi\|(\|\phi\|(I))$, which expresses that we arrive at a new context by applying the update interpretation of

the second conjunct ψ , to the information set we get by applying the update interpretation of the first conjunct ϕ to our original information set. So interpreting a conjunction would really amount to interpreting the first conjunct with respect to the information we already have, which results in a new information set, and interpret the second conjunct with respect to that, arriving at our final new information set. Any real dynamic notion of conjunction and update will *essentially* have this form, but although we *can* give them this form in a PL-framework, it is not essential at all to do it this way. For notice that our 'dynamic' definition of conjunction given above is equivalent with $\lambda X : \|\phi\|(X) \cap \|\psi\|(X)$ which in turn is nothing but $\lambda X : X \cap |\phi \wedge \psi|$. But these are global static notions of conjunction and update. For one thing, the order in which we take the two conjuncts will never make a difference! A semantics is dynamic if and only if its notion of conjunction is dynamic, and hence non-commutative.

In DPL the notion of the semantic content of a formula is different from that in PL. We associate with a formula a function from assignments to sets of assignments. (This is just another way of saying that it is a relation between assignments.) Starting with an assignment g as input, we get the output set denoted as $\|\phi\|_g$, the elements of which are the alternative possibilities for continuation of the interpretation process. In some cases this set will be a singleton set containing just the input assignment, which is the case if the formula in question does not have dynamic effects. In other cases it will be the empty set, which means that no continuation is possible after that input, i.e. it means failure. The interesting case is where we get a real set of alternatives. This will happen characteristically if we interpret an existentially quantified formula such as $\exists xFx$. With respect to an assignment g , the continuation set will contain the assignments g' which differ minimally from g in that they assign an object to x which has the property expressed by F .

Let us now look at conjunction in this dynamic scheme. Interpreting the conjunction $\phi \wedge \psi$ means that for every g' in $\|\phi\|_g$ we follow the interpretation procedure for ψ . In doing so we end up with a set of continuation sets $\{\|\psi\|_{g'} | g' \in \|\phi\|_g\}$. We get the continuation set of the conjunction $\phi \wedge \psi$ as a whole by taking the union $\cup\{\|\psi\|_{g'} | g' \in \|\phi\|_g\}$ of all the sets in this set of continuation sets. Then we are back again at a simple single continuation set, $\|\phi \wedge \psi\|_g$, that of the conjunction as a whole with respect to g . The important fact to note about this dynamic interpretation process is that the interpretation of the first conjunct of a conjunction plays an essential role in the interpretation of the second conjunct, hence DPL conjunction is not

commutative.

So, if we look at this procedure from an update perspective, we see that in fact this notion of interpretation does essentially incorporate a notion of update. For interpreting the second conjunct of a conjunction against the background of the interpretation of the first conjunct comes to the same thing as interpreting a formula against the background of available information. A comparison with updating makes this clear. Updating in this dynamic perspective would run as follows. We start out with an information set I , we interpret a formula ϕ with respect to each g in I , and arrive at a set $\|\phi\|_g$ for each such g . Our new information consists of the union of all these sets $\|\phi\|_g$. Hence, we may conclude that the updating aspect in the case of DPL is really incorporated into the notion of interpretation, and does not come on top of it, as in PL. That is precisely the reason why we call the semantics of DPL dynamic, and that of PL static.

This is an informal explanation of what conjunction and update in DPL amount to. We have indicated already that the really interesting things happen in case we meet existential quantification. Let us look at this in a little more detail. First, we look at standard, static predicate logic again. Suppose that in the course of the interpretation process we arrive at a formula $\exists x\phi$ which is to be interpreted with respect to an assignment g . Our next step will be to get to the interpretation of ϕ . We do so, however, not with respect to g , but with respect to a set of assignments $\{g'|g'[x]g\}$, the elements of which are like g , except that each g' in $\{g'|g'[x]g\}$ is one of all the possible reinitializations of g with respect to x only. I.e. any information about the value of x we might have had before meeting the quantifier, is cancelled. We update this set $\{g'|g'[x]g\}$ with the information contained in ϕ , $|\phi|$, in the way indicated above. The result will be that we end up with a set of assignments $\{g'|g'[x]g\} \cap |\phi|$. To get back at the interpretation of $\exists x\phi$ as a whole with respect to g , we look only whether this set is non-empty. If so, we can continue with g , if not, g is eliminated as a possibility. I.e. either we return to our original assignment g , or we meet failure. If we interpret $\exists x\phi$ with respect to an information set, we repeat this procedure for every g in this set, and will end up with a subset of the old information set, being the new information set. Suppose we consider the formula $\exists xFx \wedge Gx$, this story tells us that in interpreting the conjunction the existential quantifier will present us with a reinitialization of our assignment(s) with respect to x , and we will interpret Fx with respect to that. The second conjunct is evaluated with respect to our original assignment(s), i.e. the reinitialization caused by the existential quantifier in the first conjunct is no longer in force.

Let us contrast this with DPL existential quantification. Suppose g is our input assignment. If we meet $\exists x\phi$ we also reinitialize g with respect to x , thus arriving at some set $\{g'|g'[x]g\}$ of assignments. Now we interpret ϕ with respect to each g' in $\{g'|g'[x]g\}$. For each such g' we arrive at another set $\|\phi\|_{g'}$ of assignments. We take the grand union $\{g''|\exists g'g'[x]g\&g'' \in \|\phi\|_{g'}\}$ of these sets. The important thing is that this set, the continuation set of $\exists x\phi$, still carries the information caused by the reinitialization the quantifier has brought about. So, if we interpret $\exists xFx \wedge Gx$ in DPL, the binding properties of the quantifier are still in force when we arrive at the second conjunct.

What has happened to the basic semantic notion of truth, and with the notion of semantic content in terms of truth conditions in DPL? It is still there, but no longer as a basic notion, but as a derived one. DPL's truth definition says that a formula is true with respect to an assignment g iff the continuation set of ϕ with respect to g is non-empty, i.e. iff there is an assignment $g' \in \|\phi\|_g$. The PL-notion of semantic content as the set of assignments under which a formula is true is then easily obtained: $\lambda g \exists g' g' \in \|\phi\|_g$. The notion of truth conditions is a global notion in DPL, defined in terms of the richer notion of meaning as a relation between assignments that oils the wheels of the recursive interpretation machinery.

Let us now turn to DPL-negation. With respect to an assignment g , the continuation set of $\neg\phi$ is either $\{g\}$, in case ϕ is false in g , or it is the empty set, in case ϕ is true in g . This means that negations do not have dynamic properties in DPL. On the contrary, $\neg\phi$ will block whatever dynamic features ϕ may have. An existential quantifier in the scope of a negation cannot bind variables outside the scope of that negation. And this seems to be correct. For the pronoun in the second sentence of 'No man walks in the park. He whistles' cannot be anaphorically related to the negated existential quantifier in the first. DPL-negation is just as static as PL-negation.

Disjunction, implication and universal quantification are defined in DPL in terms of conjunction, existential quantification and negation in the usual way: $\phi \vee \psi = \neg(\neg\phi \wedge \neg\psi)$; $\phi \longrightarrow \psi = \neg(\phi \wedge \neg\psi)$; $\forall x\phi = \neg\exists x\neg\phi$. Since in all these cases the definiens has negation as its outermost logical operator, and since as we have just seen negations are static, disjunction, implication and universal quantification are also treated as static operators. Quantifiers inside their scope cannot bind variables outside their scope. And further, we can notice that the antecedent of an implication can have binding effects in its consequent, but such binding relations are impossible among the disjuncts of a disjunction. In many cases this static treatment of these operators leads

to correct predictions about anaphoric relations, but, as we shall see in the next section, in some cases it does not.

At this point we discuss an issue put forward by Bunt in section 4 of his paper. Bunt notices, quite correctly, that the two formulas $\exists xPx$ and $\exists yPy$ have a different meaning in DPL. The choice of variable causes a difference in meaning. In this respect variables in DPL behave more like constants than 'ordinary' variables. Bunt's objection is that the two formulas do not change an interpreter's information in different ways: "*Both formulas express that there is some object in the universe of discourse that satisfies the predicate P*" (*this volume*). The latter is certainly true, and is accounted for in DPL. Although the two formulas differ in meaning, their truth conditions are the same. And if we identify what a formula expresses with its truth conditions, a proposition, they do indeed express the same proposition.

The essential point about dynamic interpretation, however, is that the corresponding notion of meaning encompasses more than the proposition expressed by a formula. And so it should be. If we add a conjunct Qx to each of the two formulas $\exists xPx$ and $\exists yPy$, arriving at $\exists xPx \wedge Qx$ and $\exists yPy \wedge Qx$, the meanings of the two conjunctions (and also the propositions they express) are radically different. The first, $\exists xPx \wedge Qx$, means that there is some object in the universe of discourse that satisfies both the predicate P and the predicate Q . The second, $\exists yPy \wedge Qx$, means that there is some object in the universe of discourse that satisfies the predicate P , and that the object assigned to x satisfies the predicate Q .

This is precisely what is nice about DPL. And, of course, this difference in meaning of the two conjunctions can only be obtained if there is a difference in meaning between the the formulas $\exists xPx$ and $\exists yPy$. Informally, the difference amounts to this: the first formula gets you to a state such that some object d that satisfies the predicate P is the value of x ; the second to a state such that some object d that satisfies the predicate P is the value of y .

In a sense, the difference is only there for the time being, only until we consider (that part of) the discourse be closed off. It is only there as a temporary device for the interpretation of new occurrences of the same variable. It helps to bind them to quantifiers even if such occurrences are outside their normal, syntactic scope. Once we can be sure that no such new occurrences are to be expected this information can be dispensed with. Then we can narrow down the meaning of the two formulas to the proposition they express, and this is indeed the same proposition in both cases.

So it is precisely the fact that in DPL variables get some 'discourse'

meaning of their own, that lends it the power to deal with pronouns and variables in a more satisfactory way than is possible in PL.

Still, this does give rise to a problem that should be paid attention to, a problem that is noted by Bunt in section 4. In the former he discusses compositionality. The problem is that if we are to translate a piece of natural language discourse into DPL, we do not have a uniform translation for different occurrences of one and the same quantified phrase or pronoun. E.g. *a man* could be translated as $\exists x_1 \text{ man}(x_1) \dots$ and as $\exists x_2 \text{ man}(x_2) \dots$. And as we just noted this marks a difference in meaning. To get correct results, we have, in the case of an (existential) quantifier, to choose a 'new' variable, one that is not yet in active use by a quantifier at that particular stage in the discourse. This fact is a potential threat to compositionality in a strict interpretation of that notion. What a suitable translation is depends on the context. To our mind, this is only a minor issue, and, unlike Bunt suggests, it certainly does not call for a 'drastic' revision of the system.

First, we want to point to the fact that in Montague Grammar, the compositional framework *par excellence*, similar kinds of decisions have to be made as to what variable we are to choose to arrive at intended interpretations. This has never been taken to be a serious threat to compositionality. Once given some (arbitrary) indexing of pronouns and quantifying phrases, the translation/meaning of the whole is a function of the meanings/translations of its parts.

What is important is that, contrary to what Bunt suggests, in this respect DPL is more compositional than DRT. Even if we allow for this kind of preliminary indexing, DRT is still not compositional, whereas DPL is.

Still, there is the interesting question whether or not we could live up to a stricter form of compositionality which would forbid pre-indexing. Bunt seems to argue that this is impossible. The reason he presents for this is that what variable to choose is dependent on context. But context dependence and compositionality do not bite each other. As we have seen above, the general strategy is to build-in features of context dependence as parameters in the notion of meaning. If interpretation of an expression E depends on value V of some contextual feature F , we construct the meaning of E as a functionally dependent on the possible values F can take.

We can try to follow this strategy in the following way. A variable can be viewed as the name of an address. Each variable refers to a particular address, in different states an address may contain different values. Instead of choosing a new variable in order to refer to a new address, we could introduce a pointer to addresses, and use it in the translation of a quantified

phrase to refer to ‘the next free address’. This would eliminate the context-dependency of the translation of such phrases. It is the value of the pointer that is context-dependent, and this context-dependency would be built into the context-independent meaning of a quantified term. Different quantified phrases within the same (sub-)discourse would automatically get associated with different addresses; the choice of address would no longer be one that affects translation.

Things are different, or at least more complicated for the translation of pronominal elements. In interpreting pronouns we have to relate them to an address that is active at that particular point in the discourse. We have to choose a suitable variable that is bound at that point. But, of course, there might be several alternatives, the pronoun may be ambiguous with regard to its anaphoric relations. The pointer approach would not work in this case, it would have to point to a unique active address, but there may be several of them.

Notwithstanding this fact we could come a long way in sticking to compositionality. The general strategy is to design an algorithm for pronoun resolution. A lot of factors may play a role in this. In arriving at a compositional interpretation, one would build this resolution process into the meaning of the pronoun. So if we succeed in this, we could again use a pointer in the translation. Its value in a situation, a particular address, would be the outcome of the resolution algorithm.

5 Limitations of DPL and how to overcome them

As a logical theory, DPL is concerned with the same ‘subject matter’ as first order predicate logic: DPL-models are ordinary extensional first-order models. The difference resides in the dynamic notion of semantic content that plays a role in the recursive definition of the semantics of the language. Since its logical subject matter is the same as that of first-order predicate logic, it is obvious that, where the latter is a poor tool to apply in the semantic analysis of natural language because it is existensional and first-order, DPL is too. On the other hand, the dynamic aspect of interpretation, the way it deals with variables and binding can, in principle, be expected to be portable to richer logical systems.

So, in discussing DPL’s limitations, one should make a distinction between, on the one hand, its shortcomings in describing or explaining things that it is designed to deal with and, on the other hand, its being confined to

first-order models. (Is it a limitation of a pram that it is not a sportscar?)

Intensionality

In fact, the most important limitation of DPL that Bunt brings to the fore in section 4 of his paper, really amounts to DPL being an extensional first order logic. This being so, predicates and individual constants receive a complete and rigid interpretation in a DPL-model. As a consequence, formulas build up from predicates and individual constants are just true or false in a model, that is all that can be said about their meaning. But then, nothing else is to be expected in a first-order extensional logic.

This means that, at least at the object language level, DPL does not offer the possibility to define a *general* notion of information and information change (update). It is not designed to do that. Of course, DPL *can* account for certain changes of information, but only information with respect to the values of variables. And, so we claim, it does that in a proper way.

Information of whatever kind is essentially partial. Bunt seems to take it for granted that for this reason alone we should allow DPL-models to be partial models; and in particular, to allow the assignment of values to variables and the interpretation of constants to be partial functions. As we have shown in section 2, first of all, this is not necessary; and secondly, it would not be sufficient either. In DPL, partial information about the values of variables is represented as a set of assignments. And that is a proper way to do it. There is no need here to turn to partial assignments, and, as we have argued in section 2, doing so even does not solve the problem.

Probably misled by the machine metaphor, which compares an assignment to a machine state, Bunt must have overlooked this. A machine is just in one particular (information) state, and executing a program just gets it from one particular state into another. The comparison of the role of assignments in DPL with information states of a machine, breaks down at just this point. 'DPL-programs' are in some cases essentially non-deterministic, relational, not functional. Executing them in one particular state can result in a number of alternative possible ones, where in each of them one or more variables are assigned a certain value. That is where partial information comes in. If the program is continued with a new statement, this continuation is carried out in each of the alternatives. If we do want to stick to a comparison with programming, we might have to take the high flight into parallel programming.

Once this much has been acknowledged, it is not difficult to see that

Bunt's extension of DPL in section 5 of his paper will not do. Partializing the interpretation function of the model is of no help for exactly the same reasons why partializing assignments is not either. Intensionalizing the model is the easiest way to arrive at the required result that not only updating information about the values of variables becomes possible, but also updating information about the possible extensions of descriptive constants of the language. Extensional first-order models, partial as Bunt proposes, or complete, do not give enough logical space.

As we have indicated, sets of complete worlds can be used to represent partial information about the world. Bunt conceives of an extensional partial model itself as a representation of such information. This is reminiscent of the early days of modal logic, when logical modalities were treated at the meta-level, as quantifications over models. Intensional possible world semantics has given us the tools to deal with modality at the object language level, and in a much more refined and flexible way. The way we see it, if meaning is viewed in terms of context update, we should account for that inside a model, and should not try to deal with it in terms of a meta-relation between models. But should it be considered necessary, e.g. for computational reasons, to use the latter, then one needs to keep in mind that information cannot be represented in a single partial extensional model, for, as we have illustrated above, in order to be able to represent doubt we still need sets of models.

There is one more point connected to this that deserves attention. Contrary to Bunt's suggestion that they can be assimilated, we think that there are essential differences between the kind of information and information change involved in dealing with variables, i.e. in dealing with anaphoric relations, and that involved in interpreting descriptive constants. Information about the extension of the latter is what really counts in the end, it is what we are really interested in, it is the information we want to keep in memory. Information about the values of variables is needed only while we are still involved in the interpretation of a text. In fact, the sooner we can free our memory of this temporal information the better. And, happily, even within the process of interpreting a text, we can often forget about it pretty soon. E.g. information about possible values of variables which are introduced within the scope of a negation, of a disjunction, or of a universal quantification, can in most cases be forgotten about once we are outside their scope.

Assignments and possible world play a different role as contextual parameters. Possible values of variables are characteristically logically inde-

pendent. We can change the value of one particular value without having to worry whether or not this will affect the values of other variables. I.e. the notion of an assignment g' differing from an assignment g at most with respect to the value it assigns to a variable x , makes sense. Things are radically different for the interpretation function that serves to give the meanings of our descriptive constants. In a sense, the interpretation function is much less under our own control. It is pretty hard, rather implausible, and perhaps downright impossible to make sense of a notion like that of an interpretation function F' differing at most from an interpretation function F with respect to the predicate P in that $a \in F'(P)$. How can we be sure that such a unique F' exists? The interpretations of non-logical constants are not necessarily independent of each other, and it is rather implausible that we will be able to control the dependencies completely.

Another metaphor: the assignment of values to variables is an internal affair. We can do as we like. The interpretation of descriptive constants is an external affair, which we cannot control all by ourselves. For it relates language to the external world, and that, to a large extent, just is what it is, and not what we tell it to be.

Compositionality and NP's

DPL has a second limitation which is important in view of our wish to give a compositional semantics. DPL does make it possible to decompose a sequence of sentences such as 'A man walks in the park. He whistles' into the two sentences 'A man walks in the park' and 'He whistles', assign each a meaning of its own, and still give an account of the anaphoric link in question. In this respect it is superior to ordinary predicate logic, and to DRT. But like PL (and DRT) it is not capable of decomposing the first sentence into the NP 'A man' and the VP 'walk in the park'. In predicate logic quantified terms cannot receive an independent translation, they can only be translated in the context of sentences in which they occur.

There is a standard solution to this problem of compositionality. That solution is to use higher-order logic, or, more generally, a typed logic with λ -abstraction. An NP like 'a man' is then translated into the expression $\lambda X \exists y [\text{man}(y) \wedge X(y)]$, denoting the characteristic function of the set of properties such that at least one man has them. Such an object is called a 'generalized quantifier'. This function can take the property of walking in the park as an argument, resulting in the proposition that the property of walking is among the properties such that at least one man has them, which

is indeed equivalent to the proposition that there is a man who walks in the park, i.e. which is equivalent to the ordinary first-order translation.

So, as was the case for the first limitation, this one has a standard solution as well, and both can be solved at the same time by taking recourse to a typed intensional logic such as Montague's IL.

More on anaphora and compositionality

There is a third limitation of DPL that is of an altogether different nature. It does not have to do with DPL being an extensional first-order logic, but it concerns phenomena that are in principle within DPL's empirical domain, viz. a proper account of anaphoric relations involving first-order objects. In other words, in this case we meet a real shortcoming of DPL (which it shares, by the way, with its source of inspiration, DRT). In this case, too, we restrict ourselves to an informal discussion of a small number of examples to illustrate our point.

DPL tells us that there is an important difference in anaphoric behavior between existentially quantified terms such as *a man*, and universally quantified ones such as *every man*. Unlike our example (1), a sequence of sentences such as (3) is hard to interpret in such a way that the pronoun in the second sentence is an anaphor of the universally quantified sentence in the first:

- (3) Every man walks in the park. He whistles.

In special contexts, however, such sequences do make sense, such as in the rules of a game:

- (4) Every player chooses a pawn. He puts it on square one.

Another prediction that DPL makes, is that implications are static, i.e. that terms occurring in an implication cannot bind pronouns further on in the discourse. Again, besides examples that confirm this, there are also examples that go against it. Consider the following discourse, which concerns an implication with an existentially quantified term in the antecedent, and a pronoun in the next sentence:

- (5) If a client turns up, you treat him politely. You offer him a cup of coffee and you ask him to wait a moment.

As a last example, consider the following disjunction:

- (6) Either there is no bathroom here, or it is in a funny place.

In DPL we cannot simply take this as a disjunction of the two sentences ‘there is no bathroom here’ and ‘it is in a funny place’. DPL cannot account for anaphoric relations across disjuncts, even apart from the fact that the existential quantifier in the first sentence of (6) occurs in the scope of negation, which blocks anaphoric binding by that quantifier of a variable outside the scope of the negation. The remarkable thing about (6) is that it has a ‘donkey’ paraphrase (7), which, as we have seen, we can treat properly in DPL:

- (7) If there is a bathroom here, it is in a funny place.

Like the donkey-sentence (2) and the simple discourse (1), we also know quite well how to represent the intended readings of (4) and (5) in predicate logic (and the same holds for (6), of course):

- $$(4a) \forall x[\text{player}(x) \longrightarrow \exists y[\text{pawn}(y) \wedge \text{chooses}(x, y) \wedge \text{put_on_start}(x, y)]]$$
- $$(5a) \forall x[[\text{client}(x) \wedge \text{turn_up}(x)] \longrightarrow$$
- $$\quad [\text{treat_politely}(h, x) \wedge \text{offer_coffee}(h, x) \wedge \text{ask_to_wait}(h, x)]]$$

And we can come up with DPL-representations, too, that express their meanings:

- $$(4b) \forall x[\text{player}(x) \longrightarrow [\exists y[\text{pawn}(y) \wedge \text{chooses}(x, y)] \wedge \text{put_on_start}(x, y)]]$$
- $$(5b) \exists x[\text{client}(x) \wedge \text{turn_up}(x)] \longrightarrow$$
- $$\quad [\text{treat_politely}(h, x) \wedge \text{offer_coffee}(h, x) \wedge \text{ask_to_wait}(h, x)]$$

But it is quite obvious that if we would translate these sequences into these DPL-formulas, we break the rules of decomposing them in a natural way. In order to arrive at DPL-formulas which express their meanings, it is necessary, in both cases, to take the second sentence, not in conjunction with the first sentence as a whole, but in conjunction with the consequent of the first sentence. But clearly, this goes against the natural decomposition of these discourses. If we want to keep faithful to the principle of compositionality, we have to find a way out that does treat these cases as simple coordination of the second sentence with the first sentence as a whole. I.e. we want to end up with translations like:

$$(4c) \forall x[\text{player}(x) \longrightarrow \exists y[\text{pawn}(y) \wedge \text{chooses}(x, y)]] \wedge \text{put_on_start}(x, y)$$

$$(5c) [\exists x[\text{client}(x) \wedge \text{turn_up}(x)] \longrightarrow$$

$$\text{treat_politely}(h, x)] \wedge \text{offer_coffee}(h, x) \wedge \text{ask_to_wait}(h, x)$$

Similarly, for (6) we would to be able to translate it as (6a):

$$(6a) \neg \exists x[\text{bathroom}(x) \wedge \text{here}(x)] \vee \text{in_funny_place}(x)$$

But although these are proper DPL-formulas, they will not do as translations, since they do not express the intended meanings. The reason why they do not, is that in (4c) and (5c), the pronouns translate as variables that are not bound by the quantifiers they should be bound by. Unlike conjunction and existential quantification, implication, universal quantification, negation, and disjunction, do not allow quantifiers inside their scope to bind variables outside their scope. Furthermore, a quantifier in the first disjunct cannot bind variables in the second. To put it differently, the dynamic semantics of DPL treats only the existential quantifier and conjunction as really dynamic logical operators, all the others are interpreted statically, as in PL. The examples just discussed seem to point out the need, at least for some cases, of a yet richer, a more dynamic notion of meaning than DPL offers, a notion of meaning which also allows for real dynamic versions of the other logical operators. This time logical tradition does not offer a direct solution, we have to invent one ourselves. The solution is to further enrich the notion of dynamic meaning.

6 Hyper-Dynamic Predicate Logic

In DPL the interpretation of a sentence is a function from assignments to sets of assignments. In our ‘hyperdynamic’ semantics for predicate logic, which we will refer to as HDPL, we associate with a sentence a function from assignments to sets of sets of assignments. A set of sets of objects of a certain type is a generalized quantifier over objects of that type, so in this case we have generalized quantifiers over assignments. They correspond to natural language phrases like *an assignment such that ...*, or *every assignment such that ...*, or *the assignment such that ... or else no assignment at all*. (Likewise, the sets of assignments we have in DPL correspond to noun-like expressions such as *assignment such that...*). A generalized quantifier over a certain type of object denotes a set of properties of that type of object.

The set of properties of assignments denoted by a formula ϕ with respect to assignment g , we write again as $\|\phi\|_g$. It can be looked upon as a characterization of the properties of assignments resulting after processing ϕ in g . If X is a property of assignments, $\|\phi\|_g(X)$ expresses that X is such an output property if we interpret ϕ with input g . And by abstracting over g , we get $\lambda g : \|\phi\|_g(X)$, which denotes a specific property of assignments, the property an input assignment g should have if X is to be an output property of the assignments that may result if we interpret ϕ in g .

One such property is the property of being a possible assignment, which is represented by G , the set of all assignments. If processing ϕ in g is to be successful at all, being a possible assignment should certainly be among the output properties of ϕ with respect to g . Truth of ϕ in g , now amounts to the requirement that ϕ has this output property with respect to g , i.e. ϕ is true in g iff $\|\phi\|_g(G)$. So, $\lambda g \|\phi\|_g(G)$ presents us with the assignments under which ϕ is true, i.e. with its truth conditions.

A characterization of conjunction $\phi \wedge \psi$ now amounts to the following. With respect to input g , the conjunction has output property X iff being an input property of the second conjunct ψ which gives X as output, is a proper output property of the first conjunct ϕ . The set of these properties is represented by $\lambda X : \|\phi\|_g(\lambda g' : \|\psi\|_{g'}(X))$. Again, this is a truly dynamic notion, which could not be defined just on top of a recursive definition of truth simpliciter. According to the definition we just gave, conjunction is not commutative. It does not hold in general in our dynamic semantics that $\lambda X : \|\phi\|_g(\lambda g' : \|\psi\|_{g'}(X)) = \lambda X : \|\phi\|_g(X) \& \|\psi\|_g(X)$. As was the case in DPL, the notions of truth and truth conditions are global derived notions,

defined in terms of the richer update notion of meaning.

Conjunction is static and commutative in some cases, viz. in case the conjuncts ϕ and ψ do not have dynamic properties. This is so, for example, if ϕ is an atomic formula. An atomic formula such as Fa , is interpreted as follows: $\lambda g \lambda X : I(a) \in I(F) \& X(g)$. With respect to any particular assignment g , we either get all the properties g has, in case $I(a) \in I(F)$, i.e. in case the object denoted by the individual constant a has the property denoted by F ; or we get the empty set of properties of assignments. In other words, the generalized quantifier over assignments denoted by Fa in g can be paraphrased as ‘the assignment g in case a has the property F , or else nothing at all’. It can easily be checked that if we conjoin two such static formulas ϕ and ψ , there is no difference between $\lambda X : \|\phi\|_g(\lambda g' : \|\psi\|_{g'}(X))$ and $\lambda X : \|\phi\|_g(X) \& \|\psi\|_g(X)$.

Let us now turn to the definition of existential quantification in HDPL. The interpretation of $\exists x \phi$ is given by $\lambda g \lambda X : \exists g' g'[x]g \& \|\phi\|_{g'}(X)$. In words: X is an output property of processing $\exists x \phi$ in g iff for some g' which differs from g at most in the value it assigns to x , X is an output property of interpreting ϕ with respect to g' . The generalized quantifier over states denoted in g by $\exists x Fx$ can be paraphrased as: ‘An assignment g' which differs from g at most in the value it assigns to x , and such that the object that g' assigns to x has the property denoted by F ’. In effect, the interpretation of existential quantification is the same as that in DPL.

The difference between HDPL and DPL comes in if we look at negation. DPL-negation, which is static, can be represented in HDPL as $\lambda g \lambda X : \neg \|\phi\|_g(G) \& X(g)$. Remember, $\|\phi\|_g(G)$ stands for ϕ being true in g . So, if we interpret $\neg \phi$ with respect to g , we either get the set of properties of g , in case ϕ is false in g , or the empty set.

If we would use this notion of negation, and were to define disjunction, implication and universal quantification in the usual way in terms of negation, conjunction and existential quantification, we would get static interpretations again, and hence, exactly the same results for anaphoric treatment as in DPL. But, in the richer scheme of HDPL that we are exploring now, there is a clear-cut alternative for interpreting negation, viz. $\lambda g \lambda X : \neg \|\phi\|_g(X)$. This interpretation, which we call ‘denial’, gives you all the output properties ϕ does not have. It amounts to something like an error-message. If the denial of ϕ is true in g , i.e. if $\neg \|\phi\|_g(G)$, it says that output assignments characterized by ϕ in g are not possible output assignments. In other words, it claims that ϕ cannot be successfully accommodated in g . As we just saw, the static DPL-negation characterizes the assignment g if it is true in g ,

i.e. processing $\neg\phi$ means updating with the fact that ϕ is not true. The dynamic negation of HDPL denies that update with ϕ is possible, which is quite a different thing.

We note two things. First, if we use the denial meaning of negation in defining universal quantification and the other connectives in terms of existential quantification and conjunction in the usual way, we end up with dynamic versions of the former. This means that we are able after all to respect the principle of compositionality, translate examples like (4), (5) and (6) as (4c), (5c) and (6a), and have the latter express the meanings of the former. Second, we need not be afraid of having to give separate dynamic and static definitions; the latter can be obtained from the former by applying in each case one and the same ‘stabilizer’, a semantic operation that amounts to closing off (part of) a text from the remainder, blocking its possibilities to pass through anaphoric relations. This operation is defined as follows: $\lambda g\lambda X : \|\phi\|_g(G) \& X(g)$. Applying this operation to a text ϕ with respect to an assignment g , it will denote the assignment g in case ϕ is true in g , or else no assignment at all. E.g., if we apply it to dynamic negation, it results in the definition of static negation, as is easy to check.

Our richer version of dynamic predicate logic makes it very transparent that there are three notions of context at work in the interpretation process. For in HDPL the three turn up as distinct kinds of, hierarchically ordered formal objects. First there is the notion of an assignment as a contextual parameter. Then there is the notion of context as a global information set, represented as a set of assignments, a set of contexts in the first sense. And third, there is the notion of local, dynamic information set, embodied in the notion of a generalized quantifier over assignments, i.e. represented as a set of contexts in the second sense.

7 Dynamic Montague grammar

Like DPL, HDPL remains an extensional first-order system, and so it inherits from DPL the other two major limitations we discussed above. But in this case, too, these two limitations can, in principle, be overcome in the way we indicated, i.e. by using an intensional typed language.

We have developed an approach along these lines in Groenendijk and Stokhof (1988), which we call ‘dynamic Montague grammar’ (DMG). Like (H)DPL, it takes as its starting point a semantics of programming languages. But this time it is not dynamic logic, but Janssen’s intensional semantics

for programming languages (see Janssen, 1986) that serves as a source of inspiration. We will sum up the major ingredients of DMG, and indicate the ways in which it differs from (H)DPL.

DMG follows the main strategy of Montague Grammar. A syntax is provided for a fragment of natural language. The fragment is interpreted via a compositional translation in a logical language, in this case dynamic intensional logic. Each lexical item in the fragment is translated into an expression of the logical language, and each syntactic rule is accompanied by a semantic rule, which specifies the translations of the complex expressions produced by the syntactic rule as a function of the translations of their parts.

Compared to the semantic set-up of (H)DPL, the following differences should be noted. The intensionality of DMG brings along that a model contains, besides a domain of discourse, another set of primitive objects, called 'states'. Its ontology is essentially richer than that of (H)(D)PL. States are the basic contextual parameter of DMG. Assignments appear also, but their role is no different from the one they play in standard static semantics. So, the fundamental notion of information here is information about states. A state has an internal and an external aspect, it embodies both the notion of an external possible world, and that of an internal state. I.e. states play two roles at the same time, that of possible worlds in ordinary intensional semantics, and the role played by assignments in (H)DPL. Formulated in terms of information, we can say that information defined in terms of states is at the same time information about the external world and about the internal state.

The interpretation function in a DMG-model assigns an intension to the constants of the language, i.e. it makes them state-, i.e. context-dependent. Corresponding to the two aspects of states, the language has two different kinds of constants, ordinary descriptive constants and so-called 'discourse markers'. Besides these two kinds of constants, the logical language has variables as well, but their role is not a special one here. It is now the discourse marker, a special kind of constant, that carries the information we need to solve anaphoric relations. The discourse markers appear in the translations of (quantified) nominal phrases and pronominal expressions of natural language. These markers can be viewed as names of addresses in an internal state. They refer always to the same unique address in any state, in this they are rigid, but in different states they may have a different object as value; in this they are variable.

Discourse markers also share the logical independence of variables, and in this they differ from descriptive constants. Logical independence is guar-

anted by postulates that a model has to satisfy if it is to be a proper DMG-model. These postulates require that in any given state we can change the value of one particular discourse marker only, in a specific way, and arrive at a unique ‘updated’ state. These constraints are not imposed on descriptive constants. I.e. they concern the internal aspect of states only, and not their external, possible world aspect.

We need logical independence of discourse markers to make sense of a new kind of logical operators, called ‘state switchers’, that are introduced in the logical language. A state switcher, of the form $\{\alpha/\underline{x}\}$, can be put in front of any expression β . The interpretation of $\{\alpha/\underline{x}\}\beta$ with respect to a certain state s amounts to the instruction: interpret β with respect to the state s' which results from s by putting the semantic value α has in s in the address denoted by discourse marker \underline{x} . State switchers are an essential ingredient in the translation of quantified nominal expressions of natural language that DMG offers. But before we give an impression of how this works, we first say a few words about the interpretation of sentences.

Full sentences are not translated as expressions that denote a truth value and express a proposition. They are taken to denote generalized quantifiers over states, so their meaning (intension) is a function from states to sets of sets of states, i.e. sets of propositions, i.e. sets of properties of states. This is clearly directly related to the interpretation of sentences in HDPL, where we ended up with generalized quantifiers over assignments. The important difference is that being generalized quantifiers over states, sentence meanings embody not only information pertaining to the solution of anaphoric relations, but also information about the world that sentences carry.

The interpretation of the natural language quantifiers and connectives is structurally the same as the one we introduced in discussing HDPL. In DMG, however, the meta-language characterizations of HDPL can be expressed in the object language. Consider the following example. The translation of a simple sentence containing an indefinite term now has the following form: $\lambda p \exists z \{z/\underline{x}\} [F(\sim \underline{x}) \wedge \sim p]$. In a state s this expression denotes the generalized quantifier ‘a state s' such that for some object d in the extension of F in s , s' differs from s at most in that the value of the discourse marker \underline{x} in s' is d ’. Compare this with the interpretation of $\exists x Fx$ in HDPL.

The three notions of contexts make their appearance in DMG as well. They occur as the same kinds of formal objects as in HDPL, only now they have ‘more content’. The basic notion of contextual parameter is that of a state, which plays the role of a possible world and that of an ‘assignment’ with respect to discourse markers. The notion of a global information set

appears as proposition, a set of states. And the dynamic notion of information takes the form of a generalized quantifier over states. Since states carry more information, so do propositions and state quantifiers.

We conclude that incorporating intensionality into a dynamic semantics is really an orthogonal task. The essence of the dynamic semantics of HDPL carries over into the richer intensional schema of DMG, without problems.

8 Concluding remarks

We end this paper by returning briefly to the point raised in the introductory section about the relationship between dynamic semantics and what we have called the ‘pragmatic interpretation approach’, which was characterized as an attempted unification of the three views on meaning formulated at the beginning of this paper.

It should be obvious from the above that the various systems of dynamic semantics, DPL, HDPL and DMG, amount to a different treatment of the notion of semantic content. How then does dynamic semantics relate to Bunt’s notion pragmatic interpretation, which distinguishes between semantic content and communicative function? Part of the role of the pragmatic aspect of meaning that is embodied in the notion of communicative function as a function that changes information states, is covered in dynamic semantics in the semantic content of an expression. For as we have seen above, a notion of update of information is an ingredient of the semantic interpretation of sentences in the dynamic semantics framework. And it is an essential one, since without it we cannot give an adequate account of the semantics of anaphoric relations, and of a great many other phenomena, which we have not said anything about in this paper.

However, it should be added right away that this notion of update is more restricted than Bunt would have it. It only presents the update of the information of the addressee with the content of a discourse, but it does not concern what we might dub ‘second-order’ information, such as information about the information of other speech participants. And a second difference is that where communicative function comes into the picture of dynamic semantics at all, it is only in the form of one particular function, that of assertion. But as far as we can see, this is an accidental, rather than an essential shortcoming. We have no doubt that some orthodox, ‘static’ semantics of interrogatives can easily be transposed into the dynamic semantics framework.

So, it seems that a lot that has been allocated to the pragmatic component of a theory of meaning can and should be incorporated into semantics proper, i.e. into the analysis of semantic content. Does this mean that the view that meaning consists of semantic content and communicative function, has been falsified? That would be too rash a conclusion. A full-blown account of the way in which interpretation in natural language proceeds, does indeed need a pragmatic component on top of a semantic one. The way we see it, this pragmatic component will deal with the appropriateness conditions for various communicative acts, and with the information that is derived by reasoning on the basis of the assumption that these conditions are fulfilled. Whether we want to call such matters aspects of meaning, or of interpretation in a wider sense of the word, is largely a matter of terminology. It should be noted, though, that our definition of the pragmatic component covers only part of what some would claim that pragmatics should be concerned with. For we expect that a lot more of what is traditionally counted as being pragmatic territory, will be conquered by model theoretic semantics.

References

- Bunt, H.C. (1988) Towards a dynamic interpretation theory of utterances in dialogue. In: B.A.G. Elsendoorn and H. Bouma (eds), *Working models of human perception*, London: Academic Press.
- Groenendijk, J.A.G. and Stokhof, M.J.B. (1984) *Studies on the semantics of questions and the pragmatics of answers*, Amsterdam: Department of Philosophy.
- Groenendijk, J.A.G. and Stokhof, M.J.B. (1987) Dynamic predicate logic. Toward a compositional and non-representational discourse theory, Amsterdam: ITLI/Department of Philosophy.
- Groenendijk, J.A.G. and Stokhof, M.J.B. (1988) Dynamic Montague grammar. Going back to MG and beyond DRT, Amsterdam: ITLI/Department of Philosophy.
- Groenendijk, J.A.G. and Stokhof, M.J.B. (*forthcoming*) Compositionality, discourse and dynamic semantics. *Unpublished paper*.
- Harel, D. (1984) Dynamic logic. In: D. Gabbay and F. Guentlner (eds), *Handbook of philosophical logic, Vol. II*, Dordrecht: Reidel.
- Janssen, T.M.V. (1986) *Foundations and applications of Montague grammar. Part I: philosophy, framework, computer science*. Amsterdam: Centre for Mathematics and Computer Science, CWI Tract 19.
- Kamp, H. (1981) A theory of truth and semantic representation. In: J. Groenendijk, T.M.V. Janssen and M. Stokhof (eds), *Formal methods in the study*

of language, Amsterdam: Mathematical Centre. Reprinted in: J. Groenendijk, T.M.V. Janssen and M. Stokhof (eds), *Truth, interpretation and information*, Dordrecht, The Netherlands: Foris Publications.