

RENATE BARTSCH

**Propositional Attitudes in Dynamic
Conceptual Semantics**

LP-96-11, received: October 1996

ILLC Research Report and Technical Notes Series
Series editor: Dick de Jongh

Logic, Philosophy and Linguistics (LP) Series, ISSN: 0928-3307

Institute for Logic, Language and Computation (ILLC)
University of Amsterdam
Plantage Muidergracht 24
NL-1018 TV Amsterdam
The Netherlands
e-mail: illc@fwi.uva.nl

Propositional attitudes in dynamic conceptual semantics

Renate Bartsch
Institute of Language, Logic, and Computation (ILLC)
Department of Philosophy
University of Amsterdam

Abstract:

In this paper the shortcomings of the treatment of propositional attitudes in several branches of formal semantics will be discussed, shortly indicated by the names of Stalnaker, Hintikka, Carnap, Cresswell, Von Stechow, Asher, Barwise and Perry. One of the well-known shortcomings is that the structures of propositional contents in these approaches are either too coarse or too fine, and that the semantics contains syntactic-semantic hybrids like structured propositions or abstract situations. Furthermore, they cannot provide a good understanding of the quite flexible notion of identity of propositional attitudes. Another well-known shortcoming is the logical closedness of the set of beliefs in formal semantic theories. In traditional theories of cognition such as Fodor's, on the other hand, the notion of a propositional content as a representation formulated in some formal or natural language is not adequate, taking into account that beliefs are possible without being framed in a language. An alternative approach to propositional attitudes will be formulated in dynamic conceptual semantics, which can avoid the above mentioned shortcomings.

Propositional attitudes have been a tough topic in semantics and the shortcomings of formal semantics treatments of propositional attitudes are well-known. We will discuss the remedies proposed and then present a treatment of propositional attitudes in a new approach, dynamic conceptual semantics. It will be shown that and how, on such a rather cognitive level of semantics, we can make distinctions fine enough, but also coarse enough, to define different kinds and degrees of identity between attitude contents.

Dynamic Conceptual Semantics easily incorporates propositional attitudes because it includes, besides a model of concept formation, a constructivist approach to understanding. An important feature of this approach is in our research context, that partial individual concepts are parts of the life-histories of individuals, which are sets of experienced situations with an ordering structure on them according to space-time contiguity and other aspects of factual contiguity, such as causal contiguity and continuity of behavioural dispositions and roles and functions with respect to other individuals. The experienced situations, on the other hand, are also members of growing similarity sets of situations, whereby similarity can be due to identical internal properties or identical external relationships under certain perspectives. While the perspectives are presupposed, the different concepts, i.e. our knowledge of the different properties and relations that fall under them and define the similarity, are only the result of the concept formation. These similarity sets, formed and ordered in opposition under certain perspectives, are more or less stabilised in the sense that with adding new members under the same perspective, their internal similarity measure does not decrease noticeably. On a growing set of situations, as experienced situations, we have thus two evolving structures, the structures of contiguity resulting in histories, and especially life-histories or individual concepts, and the structures of general concepts. The situations are primitive, i.e. not interpreted or analysed, to begin with. They get analysed and hereby interpreted by being associated with other experienced situations in sets formed with respect to similarity under, to begin with, basic perspectives or similarity spaces, like colour, form, touch, smell, taste, motoric possibilities, biological based preferences and necessities, and in sets formed under contiguity and coherence in space and time and under other important relationships, like cause-effect, means-ends, instrument, action action-result. These perspectives are not concepts in the beginning, but can be conceptualised later on as second order concepts.

Later in concept formation, also other perspectives play a role which are formed with growing knowledge and abilities and are imposed culturally. By acquiring more advanced perspectives, situations, also the previously experienced ones, get further analysed by becoming members of similarity sets and contiguity sets, especially histories, under these new perspectives and new contiguity orderings. Thus also the interpretation of situations changes from ordering in sets of mere impressions, for example, 'horse-running' impressions and 'horse standing' impressions, to situations with a recurring individual horse and an event consisting of this horse running and an event of the same horse standing. This means that the situations are integrated into the same individual horse-history, i.e. the same individual horse concept, and are integrated into the general

concepts of (horse) standing and (horse) running. Furtheron, situations get analysed by general individual horse concepts as being a situation of some horse running or some horse standing. More details and some formal statements of this genetic model of concept formation can be found in Bartsch 1993 and 1996.

The propositional content of attitudes can be construed as a set of constraints on the possible fulfilment of the attitude. These are the attitude's satisfaction conditions. The constraints amount to the requirement that a possible satisfaction situation must be integratable *salva stability* into the growing sets, which make up the conceptual structure available at a certain stage of development. As far as the general concepts are concerned, this means that the sets which represent these, can be extended by the satisfaction situation in question without being destabilised in their internal similarity degree under the relevant perspective, and as far as individual concepts are concerned, the possible satisfaction situation can be added to the respective partial individual concepts without destabilising the coherence and the space time contiguity of the partial life-history. To entertain a certain attitude then does not imply that the subject of the attitude must have a relationship to a certain sentence or proposition expressed symbolically in some language, though this need not be excluded in special cases. Rather the subject can be assumed to have activated within an attitude state a constellation of sets of experienced situations which represent the individual concepts and the general concepts involved. This constellation configured in a process of setting such sets in relationship to each other can have syntactic properties in parallel with constituent structures. Therefore, when linguistic expression is involved, such structurings of constellations can be mapped onto syntactic structures of utterances, while the sets themselves can be mapped on content expressing lexical material. The other way around, a linguistic expression can be understood by mapping its parts onto concept representing sets of situations and having the syntactic structure guide the hierarchical structure of relating the addressing of the sets, thus forming the constellation in the process of 'activating' the sets into which the possible satisfaction situation must fit *salva stability* (Bartsch 1996). This view on propositional contents is germane to Ramsey, Stich, and Garon (1991) in that it does not assume a fixed mental representation of a propositional content in a language of thought, like the Fodorian position (1975, 1987) does. But it does not discard syntactic structure of attitude contents altogether, though this structure is not a structure of symbols formed in some mental language or some formal language generally. Rather it is a manner of gathering satisfaction constraints in addressing growing, internally structured sets of data, to which the possible satisfaction situation has to fit *salva stability*. This manner of gathering constraints is provided by forming syntheses of constraints, following in

principle the constitution of actions. This view is quite in the line with Piaget's (1947) ideas that senso-motoric schemata, routines, are formed and generalised in the process of developing from the stadium of egocentric action and perception towards growing flexibility in taking other perspectives, understanding transitivity, and achieving a structured view on situations.

The situations, i.e. the data in the growing sets that represent the individual concepts and the general concepts, are experienced situations and especially satisfaction situations for utterances, experienced by the language users. Agreement between communicating individuals about the fact that they perceive the same situations make these situations count as intersubjective, as far as the linguistic description goes about which one agrees. Intersubjective agreement on linking certain situations or constellations of situations to certain utterances as satisfying these, makes concept formation for a large part guided by linguistic usage and makes concept formation to that extent intersubjective. On a second level of concept formation, theoretical concepts, in the broad sense, are intersubjectively construed by agreement about holding true the same general sentences in the process of acquiring a common ground of knowledge. That the experienced situations are embedded in extendible sets of perceived situations and that there is bodily contact with parts of the situations in motoric activity, makes them to be understood as parts of real situations, which are the realistic counterparts into which the experienced situations of different individuals are understood as being embedded. On this basis, experiential concepts are realistically based and are understood as being represented in reality by sets of situations into which the experienced situations are embedded and which are seen as extensions of these *salva* stability. That the formation of general sentences held true is guided by insight in their reliability in inferring verifiable predictions about how things are going in the world, makes concepts explicated in these general sentences realistic, i.e. objectively satisfiable according to accepted procedures.

The approach to concept formation sketched above is incorporated into a treatment of propositional attitudes according to the philosophy of intentionality, which, especially in the version of Searle (1983), has suffered from employing a not well understood notion of intentional content as being a representation, which gives us the satisfaction conditions when being conjoined with the non-representational background of beliefs, abilities and skills. The status of this representation, or intentional content, remains totally unclarified in Searle's philosophy; it is not linguistic representation itself but some representation parallel with it. We are left with the question of what this representation might be and we might be tempted to take refuge to a Fodorian language of thought (Fodor 1975), in which the representations are formulated. Luckily, there is, in the present

model, an alternative way of understanding propositional contents in terms of structures of growing sets of data with some notion of stabilisation of the internal similarity measures and the contiguity structures on these growing subsets, which are taken as general concepts and as individual concepts, respectively. The similarity measure under certain perspectives, a primitive notion on this level of semantic theory, can be founded on connectionist or probability based models of concept formation. Also contiguity needs in principle be based on a measure of a kind of associative closeness concerning space-time perception and other factual relationships, by which coherence within individual concepts and other kind of histories is provided. Although this all is far from being achieved in connectionism, let us assume anyhow, that the addressing of growing sets of experiential data in our model of dynamic conceptual semantics can in principle be paralleled by patterns of activation in connectionist networks, and that these networks can in principle serve as models for patterns of activation of neural nets in the brain when it entertains some attitude content (Bartsch 1996).

In the theory of propositional attitudes advocated here, we shall take the intentional content to be the satisfaction conditions themselves, and not, as Searle 1983 and 1992 does, take it as a representation that in conjunction with a background of a non-representational kind provides the satisfaction conditions. In our theory, the background is our whole knowledge set, containing an experiential knowledge base in terms of experiences of situations, partially restricted by linguistically explicated knowledge. The linguistically explicated knowledge, which is also part of our knowledge set, consists of particular sentences held true, which describe situations, and histories, and of general sentences held true, which form coherent sets which are theories. Our concepts are not more than structurings within, or on, our knowledge set. The satisfaction conditions of propositional attitudes are established on our conceptual structure by addressing parts of it, and by this the satisfaction conditions are structures on sets of data in our knowledge set, including our experiential base. They are, on the one hand side, syntactically structured as hierarchically ordered patterns of the process of gathering of satisfaction constraints in our conceptual structure, and because of this they can be represented by linguistic expressions in some formal or natural language. However, the propositional contents are not representations themselves in the sense of traditional, Fodorian, cognitive theory. They can be entertained without any language, and depending on how involved the structuring of addressing concepts can be, also speechless animals can entertain attitude contents of more or less complexity. A further important property of attitude contents is that their components are selected within our knowledge set and hereby

they are fully endowed with context-dependence and information concerning experiential impact, which need not all be linguistically explicit. In this way our intentional contents include in themselves both, the syntactic properties of representations, and all the aspects Searle adds to his representations or intentional contents by applying them to the background and hereby achieving the satisfaction conditions.

1. Structured propositions as contents of attitudes

The treatment of propositional attitudes has always been a problem for referential semantics. Certainly it was impossible in extensional semantics, where sentences are semantically equivalent if they have the same truth value. The semantic object that is the interpretation of a sentence is a truth value, and this semantic object cannot be the object of a propositional attitude. If it were, then all true beliefs were identical and all false beliefs were identical. But intensional semantics does not fare much better. The semantic objects expressed by sentences are propositions, which are functions from possible worlds, or more generally from indices, to truth values. Sentences with the same truth value in our world can have different truth values in other possible worlds, and hereby they are distinguished semantically: they have in this world the same truth value, but nevertheless they express different propositions. Still, sentences that are tautologies and all necessarily true sentences express the same propositions, i.e. the same functions from worlds to truth values. Likewise, all contradictions and all necessarily false sentences do. Therefore, if the propositions of intensional logic were the semantic objects of propositional attitudes, then to believe that a dog is an animal would not be different from believing that a rose is a flower. Since they get the same truth value in every possible world, both express the same function from possible worlds to truth values, i.e. express the same proposition. Likewise, believing in one tautology would be the same as believing in all other tautologies. As has been recognised for a long time, propositions, as they are conceived of in intensional logic, are still semantically too coarse to be the objects of propositional attitudes.

Stalnaker (1987), of course, realises these problems of the intensional logic notion of proposition, which arise in attitude contexts. But he points out the advantage of such a view: If a belief state is represented by a set of possible worlds, and believing that p means that p is the case in these worlds, then being in such a state can be understood as the capacity to discriminate worlds that fulfil

the belief from those that do not, while taking the real world to be an element of the set of satisfying worlds. Implicit beliefs are treated semantically on a par with linguistically explicit ones. The discrimination between belief-fulfilling situations and others shows, pragmatically, in those actions and activities of a rational believer and actor which presuppose such belief.

Stalnaker's position would also be very attractive from an epistemological point of view, if it would be possible to say what this capacity to discriminate possible worlds into those fulfilling, and those not fulfilling a belief amounts to, and it would be also semantically more satisfying, if it could distinguish between believing different necessarily true sentences. Unfortunately, this cannot be done in a formal semantics based on total extensions of expressions in worlds. It can be done in a dynamic conceptual semantics, i.e. in a data-based semantics on the level of understanding. And the treatment on this level has to be such that on the level of interpretation in worlds, Stalnaker's model-theoretic semantics results. This chapter, accordingly, addresses the question what the so-called propositional contents of attitudes are on the level of understanding, such that they make up the capacity to discriminate between fulfilling and not fulfilling situations and worlds in all individual cases of belief.

In order to solve the latter question, linguistic philosophers and other semanticists have been looking for something more fine grained than the propositions of intensional logic. In *Meaning and Necessity*, Carnap (1947/1956) proposed a stricter criterion for identity of propositions as objects of propositional attitudes, namely that besides having the same intensions, these would also have to be built up in the same way. The intensions are built up in the same way if the expressions by which they are expressed are built up in the same way: Two sentences express the same meaning if they are intensionally isomorphic. Sentences that only differ by replacement of synonyms are intensionally isomorphic. The isomorphism between them preserves the same syntactic built-up of the intensions and thus, structured intensions can be distinguished in a more fine grained manner. The difficulty is that a model-theoretic notion of synonymy itself identifies all necessarily true statements. Therefore here synonymy must be restricted to lexical synonymy, or at least to synonymy of predicates that do not contain a conjunction of contradicting or of tautological predicates. In order to distinguish this kind of synonymy from synonymy generally we have to recourse to criteria of form, i.e. syntax. This not only leads to a mix up between semantics and syntax, but it is not principled enough. Either we restrict synonymy ad hoc to certain forms of predicates, or we restrict it to simple predicates, which are expressed by a word. But if we restrict synonymy to lexical synonymy, it is just an accidental matter how far this

notion applies. This is so because it is accidental, whether some property is expressed by one word or by a complex description, up to a full sentence. Other semanticists, for example Lewis (1972), and later Cresswell (1985) and von Stechow (1984), followed Carnap (1947/1956). They formulated a notion of 'structured propositions'. These are supposed to be semantic objects, and for Cresswell they consist out of referents and extensions of property-expressions in the real and in possible worlds, but they are still syntactically structured. They are a kind of syntactic-semantic hybrids, being intended as semantic objects of intensional logic, while still carrying on their sleeves the way in which they were composed. Their status in intensional semantics is far from clear. Cresswell compares them to Perry's situations, which are constructs made up from individuals, properties or relations, and truth values, according to whether the property or relation applies to these individuals or not. But this more realistic picture of propositions does not solve the problem of what their status is in intensional logic, which has to have a compositional semantics; and this includes that within a composition, also within an attitude context, the syntactic forms of two structures do not matter as long as they are semantically equivalent. What matters compositionally is merely the semantic value of the structure and its parts, and for propositions this is a set of possible worlds. How can a semantic object be both, a set of possible worlds and a syntactically structured construct out of other semantic objects, entities and properties, without both being interchangeable in all contexts? Analogously, in arithmetics, 5 can be represented also as $(2+3)$ or $(1+4)$, etc. But they are interchangeable in all contexts. However, this is not so for sets of possible worlds and structured propositions, and this means that semantic compositionality of intensional logic breaks down if attitude contexts are included in its language and one wants to squeeze their identity conditions with respect to meaning into a model-theoretic framework. It is methodologically misguided to mix up semantic and syntactic aspects and to mess up model-theoretic semantics, which deals with interpretation as evaluation of linguistic expressions in worlds, with conceptual phenomena that play a role in understanding and have, of course, repercussions in model-theoretic (extensional and intensional) semantics, because understanding is involved as a prerequisite of interpretation. But can understanding and other conceptual phenomena not be treated within modeltheoretic semantics without damaging the guiding principles of dealing with extensions and obeying semantic compositionality in dealing with them. Several aspects that can be distinguished on the level of concept formation and understanding collapse in the extensionally based intensional semantics and cannot be recovered. The relationship between the structures on the level of understanding and those on the realistic and

possible worlds level of interpretation is weaker than an isomorphism; it is a homomorphism in which a finer structure, the conceptual structure, is mapped on a coarser one, fine structure is lost at some points and coarse structure is preserved in the mapping.

Propositions, structured according to natural language syntax, are too fine grained as objects of propositional attitudes: To believe that the dog has bitten the cat and to believe that the cat was bitten by the dog would be different beliefs, since such structured propositions are different. In order to make such proposals work, one has to be able to take into account all kinds of paraphrases of quite different linguistic structures. These have to be identified on the level of structured propositions. But then the notion of structure gets obscure. Have all linguistically possible paraphrases the same propositional structure? If not, which do and which do not? Are there semantic criteria to answer this question in a principled way? A sufficient answer has not been given.

Even if the previous problem could be solved, there is another group of problems, which are concerned with deciding about identity of beliefs. These problems are due to the influence a context or background knowledge has on which belief is expressed when a sentence is used for representing a belief content. They cannot be tackled by a combination of syntax and intensional semantics, in terms of structured propositions, though Stalnaker's approach at last in principle takes care of context-dependence of expressing a belief by a sentence. The solution in the present essay to the problem of identity of beliefs will be to strictly separate something like structured propositions from the realistic semantics of possible worlds or situations. However, it will not be a Fodorian approach (1976, 1987) with propositions as mental entities formulated in a language of thought. It will also not be a Discourse Representation Theory (DRT) approach like Asher 1986, where representations are used in the format of discourse representation structures (DRSes). Though there the information is built up in a dynamic way by the successive introduction of referent markers and conditions on these, the referents are notional objects, which are pegs on which the informations expressed in the DRSes are hanged. In *de re* beliefs the notional objects are anchored to actual objects. It will be shown that pegs, which will have to be identified across the beliefs of different persons if they have beliefs about the same objects, and which will have to be identified with referents that are anchored in the real world, are too coarse to represent objects of belief. Also it is unclear what ontological status representations in a DRS-language, like the Fodorian language of thought expressions, can have in the minds of people. Propositional structures, and also structures like the DRSes, should not be treated as symbolic representations which minds manipulate according to the Fodorian

account. Rather such structures can be found in the ways of relating concepts, i.e. in our procedures of addressing stabilising sets of experienced situations when we gather the constraints for possible satisfaction situations. These procedures are syntactically structured, and they act on parts of the structured sets of data, which make up our conceptual system. The structured sets of data are kind of distributed representations that are not constant over time and over various individual believers, but they are located in sets that play a role in partly stable structural relationships. In these ways conceptual structures erected on growing sets of data, such as experiences and changing theoretical knowledge, have some structural parallels to distributed representations in connectionist networks. Likewise, the syntactically organised way of addressing concepts on the data level can be paralleled on the connectionist level by using recurrent fibres and parallel processes. Via such models we can imagine global sequences of patternings that support physically the propositional attitudes we are aware of. The upshot of these considerations is that we keep the level of semantics proper, i.e. model-theoretic interpretation, clearly separated from dynamic conceptual semantics, which, as it is conceived of here, has to respect the facts stated by Stalnaker for the model-theoretic, interpretational, treatment of propositional attitudes and understanding linguistic expressions. Understanding has to result in making the right discriminations in interpretation. On the other hand understanding and entertaining propositional contents in attitudes has additional structures to those structural aspects it inherits from the requirements of modeltheoretic interpretation. It is fed from the rich structure that an understanding individual has available from ordering previous experience and knowledge, according to contiguity and similarity under perspectives. Here, more structuring is available than is linguistically expressed in single utterances, and even in language generally. These additional structures play a role in understanding and in entertaining propositional attitudes. This whole capacity of understanding is rooted in processes and states of the neural networks of a person. Nevertheless, on the level of dynamic conceptual semantics we can parallel in a cognitive model some of these processes and states. This model is provided by a cognitive theory which is a theory of concept formation as ordering growing sets of data under different perspectives, a theory of understanding new data as integrating them *salva stabilitate* into the sets that form the established and also changing structures on the data, and a theory of propositional contents as ways of addressing or 'activating' such sets of data, which amounts to forming combinations of the concepts they represent.

2. Identity of beliefs in formal and in conceptual semantics

Beliefs people entertain are partially made public linguistically by uttering belief-sentences. Such sentences represent only some aspects of the beliefs, and only these aspects can be under public attention. Depending on how much the hearer knows about the history of the believer he can guess that certain further aspects are involved as part of the belief, but are not expressed. If we say that two people have the same belief on the basis of their assenting to the same utterance we can only take into account what is expressed and what is its public reference. We then restrict ourselves to communicated belief and there is more to a belief than that. In a study of the semantics of belief sentences this is all that can be dealt with, but in a cognitive study, a study about belief and conceptualisation, we also can deal with further aspects of beliefs and can get into the picture finer grained identity relationships.

In order to decide about identity of belief, we have to take into account available background knowledge and different perspectives under which judgements about identity of beliefs are made. A belief is not expressed simply by a sentence, rather by a sentence together with a presuppositional context or background. These together make up a proposition, in the model-theoretic semantics of Stalnaker (1987). We shall now analyse the notion 'having the same belief' by considering the following group of examples and comparing how the different approaches can incorporate these.

People judge that John's belief that his brother will come is the same belief as John's belief that the president of the big oil company will come, if he knows that this president is his brother. A notion of structured propositions cannot capture this kind of identity judgement about beliefs. The referent itself, or its deictically fixed individual concept, would have to be part of the proposition. But for that the notion of structured propositions is too fine grained. It makes propositional attitudes sensitive to the logical form or macrostructure of the sentences (Cresswell 1985) and on this level identity of beliefs is determined, though in evaluation the identity of the referent across both sentences can be taken into account by the accompanying identity statement, which identifies the referents involved in the two different belief contents. But this fact is not available on the level of structured meanings, though it can be inferred. In Stalnaker's semantics the senses of the respective referring expressions, serve to identify their referents in each world, which are identified in all of John's belief-worlds. This would explain, why we may hold that the two beliefs are identical. However, this eliminates the difference which is nevertheless involved in John's

attitudes with respect to both belief-contents. In Asher's (1986) DRT-treatment of propositional attitudes we would have in the DRSeS two different referents with different conditions on them, according to the two different descriptions, and the two referents would be identified. But this identification is too coarse, as will be evident from the example in the following paragraph, where the different referents have to be anchored to, i.e. identified with the same external actual referent. On the other hand, in dynamic conceptual semantics the senses of the two referring expressions are understood as two different partial individual concepts, which get identified merely in the sense that they are seen as parts of the same complete individual concept in John's belief-worlds. In John's conceptual system, the two partial individual concepts do not collapse; they are part of one individual history and in his interpretation in his belief-worlds they assign the same individual. This differentiation is important for understanding the identity and also the difference between the two beliefs, as it is more clearly illustrated in the next example.

The structured proposition approach cannot capture that John and Bill are considered to believe the same if John believes that his brother will come to the conference and Bill believes that his former schoolmate will come to the conference, and if the two expected men happen to be the same person. Here a single situation in the real world will satisfy both beliefs. Taking their respective references in rigid designator interpretation, we may say that in this sense both, John and Bill, believe the same. But if Bill does not know that Fred, who will come, is the brother of John, and if John does not know that Fred is Bill's former schoolmate, we are more likely to doubt whether we want to say that both believe the same thing; and even if they do know about their different acquaintance relationship to the same person, their beliefs have different contents in terms of different kinds of predictable behaviour. Though a single event satisfies both belief contents, this event is placed in a different series of other experienced or expected events for both, John and Bill. In this way it is fulfilling different satisfaction conditions, or constraints posed by the different conceptual systems of John and Bill. The dynamic conceptual approach is compatible with taking the reference by John to his brother, and by Bill to his former schoolmate, as a rigid designator in interpretation, which explains our intuition of identity in the first place and also explains our intuition that both have different beliefs, even if they know that they refer to the same person. In John's and in Bill's respective conceptual system a different partial individual concept plays a role for identifying the person in question in a possible satisfaction situation, but both partial individual concepts are part of one completion which is found in interpretation in the world.

A general problem for a theory of propositional attitudes is, whether beliefs are identified according to what they refer to in the world, or according to the structuring of the proposition expressing the belief, or whether the proposition together with background knowledge constitutes a belief and both together have to be taken into account for judging about sameness of beliefs. In a certain sense John and Bill, in the above example, believe the same thing, namely that a certain man will come. He is the common referent of their beliefs. Depending on whether they know that they refer to the same person, and depending on the difference in their personal relationships to that person, we are more or less willing to say that they believe the same thing. Even if they don't know that they refer to the same person we would say that, in a certain sense, they believe the same thing, but they don't know that they believe the same thing. A theory of propositional attitudes should explain the different intuitions we have about identity of beliefs under different points of view.

Even if there is structural identity and referential identity, there remains some doubt whether we would say that both persons have the same belief. Imagine John and Bill have a common best friend, and both believe that their best friend will come. Do they really have the same belief? They view the situation each from their own perspective, from the point of view of the history of their own friendship, and they have different expectations because of this. Thus it seems that we have in fact two coinciding situations they refer to, but not an identical one. There is the situation of John's friend coming and the situation of Bill's friend coming, which coincide in place and time. What satisfies Bill's belief is something other than what satisfies John's belief because it is placed into different sets of situations, different partial histories, *salva* stability and continuity. If we would analyse situations in the world as intensional entities we could make this distinction. There are possible worlds in which Bill's best friend is not John's best friend. Imagine, the common friend plays two totally separated roles. Then we would say that John and Bill have different beliefs. This common friend may even have two names, one known to Bill but not to John, and the other name known to John but not to Bill. However, assuming the best theory of proper names and taking them to be rigid designators, i.e. having the same referent in all possible worlds, we could not distinguish both beliefs in intensional semantics. This is so because also the belief worlds are members of the set of possible worlds.

In Hintikka's (1969) logic of belief, a difference can be made because Bill and John handle two different individuating functions as individual concepts for referring to their friend, and these individual concepts are the intensional referents. They have the same value in all epistemic alternatives of Bill (i.e. all

possible worlds compatible with what Bill believes) and they have, accidentally, also this value in all epistemic alternatives of John. If we take as the semantic value of a proper name the individuating functions which *x* uses in his belief, then as the content of a belief, i.e. as the proposition believed by *x*, we have to take the function from possible worlds to truth values, which for *x*'s epistemic alternatives yields the value true. However, with respect to distinguishing beliefs in different tautologies such a semantics of epistemic alternatives is too coarse, like possible worlds semantics generally is. And, of course, in this kind of semantics it holds that if someone believes a set of sentences, then he also believes all implications that follow from conjunctions of these sentences. This might be all right under a normative point of view, but it is not right descriptively. It is counterintuitive descriptively, because logical closure can only be the case if the believer realises all the implications from the sentences which he believes. This inadequacy of possible worlds semantics for the analysis of propositional attitudes is well-known: A set of beliefs is not logically closed, i.e. it need not contain all implications of these beliefs. Trying to find a solution to these problems within model-theoretic semantics itself can only lead to messing up an otherwise beautiful theory. Rather interpretational semantics should be complemented with a theory of dynamic conceptual semantics as a theory of understanding, where these phenomena belong and easily can be dealt with.

In order to see how a semantics with conceptually structured intensional situations, as a special sort of basic entities (cf. Bartsch 1995), works for propositional attitudes, now imagine that there are two (not half) brothers, Bill and John. Would then John's belief that his father will come be identical with Bill's belief that his father will come? In all possible worlds accessible from their origin, they would have the same father. Thus, according to a theory of situations as intensional entities, they necessarily believe the same if each believes that his father will come. To some degree this is so. But still we might want to say, that this is not quite so; more is required, namely that each not only believes that his father will come, but that each believes that their (common) father will come. This is a subtle difference of perspective, namely whether in believing a personal perspective is chosen, or a common perspective. The identity of situations then would not merely require identity in all possible worlds, but also identity under all perspectives. A personal perspective is a way of classifying: The man referred to can be classified under a concept bound strictly personally, which here is 'my father', or under a socially bound concept 'our father'. Under the first perspective, the situation will be one analysed in its relevance strictly for me, in the second case, in its common relevance for us. The object of belief can thus be as fine grained as situations can be, understood as

intended objects. The intentions involved here can be pre-structured by concepts (possibly expressed linguistically), and by perspectives of classification. Under the perspective chosen, the possible satisfaction situation is incorporated into different partial histories, i.e. series of experienced situations in a continuous and stable manner. In our example it seems that we can do with the intensionality in terms of possible worlds, because there is for John a possible world in which "my father" would not be "our father", namely a world in which his brother Bill would not have existed. Thus intensionality seems to scrutinise fine enough for this type of example, but still it does not capture the difference at hand. It can only be accounted for by incorporation of the possible satisfaction situation into the different partial histories of Bill's and John's respective relationships with their common father, recorded in the different partial individual concepts they have of their father. This is the ready at hand solution in dynamic conceptual semantics.

Sometimes also two beliefs are judged to be the same beliefs if they are conceptually quite the same and are even expressed in the same way, but refer to different things in the world. Here identity is meant as identity of the type of situations. For example, John and Bill both believe that their best friend will come to town. Both subscribe to the same sentence for representing their belief: "My best friend will come to town". But both refer to a different person as their best friend. Here the satisfaction situations do not coincide. Even if they are viewed under the same perspective and the same linguistic or conceptual categorisation, they are two different situations. Thus the objects of the two beliefs are different situations. In this sense, both men do not have the same belief. Still we often say that two people have the same wish, if they both desire that their best friend should come, although it is a different person and herewith, a different situation is the object of desire. We apparently have to distinguish between being (partly) in the same conceptual state and referring to the same state of affairs in the world. Both aspects play a role in judging about identity of propositional attitudes. The kinds of identity in these examples cannot be handled in a realistic or possible world semantics. Rather they require a cognitive approach in which one can speak of the same conceptual state even vis à vis different realistic situations of satisfaction, and also the other way around of different conceptual states vis à vis the same realistic situation.

In an intensional semantics in which we can refer to types of (intensional) situations we can provide for an analysis that makes distinctions parallel to identity and difference of conceptual states, as far as it is needed for these kinds of examples. In such a model with intensional situations as entities, structured propositions are superfluous for dealing with this kind of examples (Bartsch

1995). However, in such a semantics, tautologies still cannot be distinguished. In order to distinguish between tautologies and also between necessarily true sentences generally, structured propositions would still be called for. These hybrids, in between language and its semantic interpretation in terms of (truth) evaluations, do not really fit into an intensional semantics which works compositionally within models where the meanings are functions from possible worlds to entities, sets of entities, sets of entities to sets of entities, etc. The status of such hybrids, structured propositions, can only be mental. But this does not fit with a non-mentalistic semantics, such as possible world semantics and situations semantics of Barwise and Perry (1983). I have pointed out above that structured propositions cannot by themselves be belief contents. Information states of individuals have to be taken into account for explaining many judgements about identity of beliefs. These are conceptual states, some logical properties of which have been expressed in update (data) semantics (Veltman 1983, Landman 1986) and in other partial logics, where semantic evaluation takes place on information states rather than on complete possible worlds. The point is to treat conceptual entities, like information states, in an also logically satisfactory way. Partial logics and update semantics have been developed and further get elaborated in order to do this. But they still have the property that tautologies which are expressible with respect to one and the same information state such that their presuppositions are satisfied there, are semantically indistinguishable; in discourse semantics or dynamic intensional logic, additional differences between tautologies can be made as far as different discourse referents are involved. - See for this the combination of update and discourse semantics by Dekker (1993). - But after static closure of a text interpretation, these differences become semantically irrelevant: All tautologies over discourse referents that had been introduced in the course of interpretation are finally semantically indifferent, and also discourse referents get identified like in Asher's DRT-approach, and by this differences between partial individuals that are part of the same real individual get lost. Partial semantics and constructivist semantics both are able to distinguish semantically between tautologies, though in different ways. For example, in dynamic update semantics and in situation semantics you always can find a situation, or information state, in which two tautologies about different individuals can be distinguished in the way that one has parts that have to be evaluated with the third truth-value ('undecided'), or are not defined there. Nevertheless, this still leaves many tautologies undistinguished, namely all those that have the same (contextually introduced) presuppositions.

Constructivist semantics, following for example Dummett's (1976) approach to the theory of meaning and understanding, can distinguish tautologies generally. The construction of the proof or justification of an assertive sentence, and not simply the resulting evaluation, is captured within a constructivist semantics approach. There a distinction is made between meaning, i.e. the construction of the semantic evaluation against evidence, and the final evaluation, i.e. the result in terms of truth values. Of course, so-called structured propositions are intended to be something like this, but they do not fit into a semantics where evaluation only takes place with respect to complete worlds; there the construction cannot make a semantic difference, if semantic compositionality is upheld. Rather they must have a partial semantics in which also partial individuals can be referred to such that they are not identical with the complete individuals they are part of. Such a semantics is the dynamic conceptual semantics, or data-based semantics, advocated here.

The conclusion of the previous discussions is that some structuring of propositions is necessary to get a distinction between propositions fine grained enough in order for propositions to play a semantically specifying role in the analysis of reports about propositional attitudes. But this structuring should not be squeezed into semantics proper, namely model-theoretic interpretation with complete, realistic, entities and complete extensions for predicates.

3. Intentional contents in dynamic conceptual semantics

We will now look into a different approach, namely the analysis of propositional attitudes within a philosophy of intentionality. Linguistically, *that John will come* is the object term of, e.g., *believe*, in *Fred believes that John will come*. The *that*-clause is object term of the attitudinal verb; this is so by syntactic analogy with other verb-object constructions, but it does not say much about how far the semantic parallel goes. It certainly does not imply that propositions have to be the objects of propositional attitude verbs, like an apple is the object of eating as represented in sentences like *John ate an apple*.

In the philosophy of intentionality, for example in Searle's (1983) *Intentionality*, and also in Husserl's (1913/1920) *Logische Untersuchungen*, the intentional content which we express by a *that*-clause in an attitude report is understood as representation of the satisfaction conditions of the propositional attitude it specifies. For example, John's belief that Fred will come will be fulfilled if Fred would come. The so-called intentional content of the attitude is not thought of as

an independent object standing in the (external) relationship of 'believe' to the subject. Rather it is part of the intention (attitude) of the subject. Also the intention or attitude is not something the subject has, like he can have a car; rather the subject is in an intentional state. This 'being in' means that the intentional state, i.e. the attitude, is part of the whole state the individual exists in at this moment. The treatment of propositional attitudes in conceptual semantics, i.e. within the current theory of dynamic concept formation and understanding will be compatible with this view held in the Philosophy of Intentionality.

A propositional attitude is a realisation of an attitudinal type, for example belief, wish, hope, which is specified by a propositional content. The content of a belief consists, we can provisionally say, in an activation of parts of the set of information, the previous data, in such a way that it is an activation of parts of the conceptual structure erected on the data as a certain ordering. This implies that with changing information, different activated sets of information can internally represent a single propositional content, though the activation region or pattern still has an at least partially identical external and internal structure that contains the identities and differences necessary in order to identify and discriminate the Propositional contents, the individual concepts and the general concepts in intersubjective agreement, as well as across different times in the history of the believer. Thus he can, across time and in conformity with other language users, identify the same (types of) satisfaction situations. Merely the activated conceptual structures have to be identical, though they are built up of different growing sets of experiences and information generally. Although there are different activation patterns on our overall set of data, our information set, which make up the constraints on possible satisfaction situations, we nevertheless can speak of the same belief, if the same objective public situations are taken to satisfy them. The intersubjectivity, the objectivity, and the temporal continuity of what it is to hold the same belief is provided by the same public pattern of satisfaction/dissatisfaction situations and not by a mental state. Nevertheless, on the conceptual level there have to be structurally partly identical activation patterns or conceptual structures causally related to recognisable patterns on sets of public situations of evidence and satisfaction, and justification of statements generally, which explain why two different cognitive states still are discriminators of the same satisfaction/dissatisfaction pattern between situations in the world. This partial structural identity on the conceptual level and its causal boundedness to experienced real, often public, situations amounts to entertaining same propositional contents, or understanding sentences that expresses it. We can also identify two mental states as being conceptually identical if they show the same conceptual structuring, though they may be built up from different

knowledge bases. People have different life-histories, i.e. partly different experiences and partly different knowledge. Keeping this in mind, we may state that two beliefs are the same under a certain perspective, or linguistic expression, if they contain the same conceptual structure selected by the perspective, or by the linguistic expression: They may be conceptually identical as far as general concepts go, or additionally even as far as the completions of their partial individual concepts go. The strongest identity we find if, furthermore, also the partial individual concepts, which are involved in the belief, are the same. But that is never completely, though partially possible, if two persons have the same kinds of experiences about their common object of belief.

For the analysis of desires we also have to take into account sets of experiences of one's own activities and actions and, as orderings built on these, routines of activities and actions. - Routines are, on the set of motoric experiences and other experiences of own activities, what concepts are on the set of perceptual experiences. They are of the same kind, and they are built up together in concept formation such that established routine structures and other conceptual structures are associatively connected. Together with certain perceptual experiences, routines are the basis for action and activity concepts. - Expectations and beliefs play a role in forming successful desires: We have to adapt our activities to facts, in order to avoid clashes with reality and in order to exploit facts for our purposes.

Searle (1983) analyses a propositional attitude minimally as a direction of fit between a content and the world, whereby certain side conditions and emotional standpoints provide the more specific characterisation of different attitudes. Belief and desire are the basic propositional attitudes, distinguished by an opposition in the direction of fit. A belief is fulfilled if the propositional content fits the world; a desire is fulfilled if the world is made such that it will fit the propositional content. In order to reach satisfaction of a desire we tend to produce associated activities and actions that make happen the desired course of events.

The propositional content has two different functions, the second of which depends on the first. The first is that the propositional content specifies the intentional state or, in Husserl's terms, the intentional act such that a certain kind of situation would satisfy it. The second function is dependent on this: by means of the propositional content we can refer to something outside the intentional state, or act, as its intentional object. I call the satisfaction situation itself the intentional object of the attitude. If such a situation does not exist, then there is no intentional object, though there is still the type of intentional object, namely the propositional content specification of the attitude itself. But this specification

itself is not an intended object. We are directed by it to the outside world in a certain way by having certain expectations of fulfilment of the intention, but we do not refer to it; it just is the conceptual structuring of the attitude and thus it is internal to the attitude or act; it constitutes the attitude or act as being of a special type. The same set-up we find in Husserl's philosophy, where the noema, or the pure intentional meaning together with itsthetic function, is the type of every noetic act, or act of knowing, in which it gets fulfilled.

When we say that John has the same belief as Bill, we can mean that the attitude, internally, in its conceptual structuring, is constituted in the same way. If it is also directed referentially towards the same external objects under the same perspective, this means that it has the same type of intentional object. To have the same type of intentional object does not include that the intentional object in fact does exist. It just means that the possible situations of satisfaction and disappointment of the propositional contents of the two beliefs are the same as far as the type is concerned. If the type is fulfilled, then the two attitudinal states have the same intentional object.

For two attitudinal states to have completely the same type of intentional object means that the propositional content is conceptually, perceptively, and referentially the same. Letting identity also depend on subjective perspective implies that no two beliefs by different persons are exactly the same. Even for one subject, the same belief expressed in different periods of his life will not be completely the same, since it is embedded into different partial individual histories with the same completion, i.e. which are part of the same individual concept, but properties and relationships in the partial histories have changed in the meantime. Above we have considered an example, where neither the internal structure of the attitude, nor the intentional object was the same, merely the real object or real situation in which satisfaction occurred was the same. What here is the same is, in Husserl's terminology, not 'the intended object', but 'the object which is intended'. This 'real' object or 'real' situation comprises all intentional objects and intentional situations that have been, are now, and possibly will be aspects of it. They are selected under different perspectives and at different times.

The *intended object or situation* is a partial object or partial situation, the *object or the situation which is intended* is a complete object or complete situation. The complete object or situation is the real object or real situation which is intended. For us it never is complete, though it is, as being real, thought of as the object which is intended, namely a partial intended object together with its horizon of still open epistemic possibilities of determination and characterisation in space and time. This imbedding of a partial object into the still available epistemic

possibilities, and into further courses of events, makes, according to Husserl, from an intentional object the object which is intended. By each new specification, the partial situation or the partial object gets embedded into new horizons of possibilities of specifications. It is embedded into possible sequences of specifications, and accordingly with each new information, into further horizons of further possible specifications. This process is guided by the questions which can be asked, given that new state of information. This successively embedding is what makes partial objects to be part of a real object. It constitutes the reality, or objectivity, of the partial object. Husserl (1939/1985) in *Erfahrung und Urteil*, calls this the "Fundamentalstruktur" of experience and knowledge.

In our discussion about what it can mean to say that A and B have the same belief, sameness was determined on three accounts, first with respect to real and possible objects and situations in the world, second with respect to conceptual structure as far as general concepts go, and third with respect to conceptual structure as far as general and partial individual concepts go. In the last respect it was important what the partial individual histories or concepts are, into which the possible satisfaction situations have to be integrated in a continuous and stable manner. Indefinite and definite descriptions provide not only means for identifying and re-identifying individuals, but also provide a perspective under which the partial individual concepts, imbeddable into the complete individuals, are selected. Thus the function of descriptions can be completely accounted for in dynamic conceptual semantics, while other semantic accounts such as Discourse Representation Theory, or the Dynamic Semantics of Groenendijk and Stokhof (1991), treat indefinite descriptions merely as introducers of a complete referent, and definite descriptions as re-identifiers of the complete referent, which is intended in an extensional interpretation. On all three accounts of same belief mentioned above we have a notion of same belief, which is practical in the theory of decision and action. Furthermore, there are combinations of these notions of sameness in cases where the same realm of possible situations is intended under identical conceptual structure, possibly expressed by an identical description, and under the same perspective of a personal history.

4. Intentionality and a constructivist approach to meaning

A propositional content provides structure for an attitude or intentional state or act. This content is the set of constraints a possible satisfaction situation must

fulfil. In a belief, as well as in a desire, there is involved a certain structure of expectations about possible satisfaction situations. The constraints are the structures and content specifications of these expectations. The expectations can be fulfilled or disappointed, and here-in consists the satisfaction or dissatisfaction of the attitude, of the intentional state or intentional act. The set of constraints can be expressed by a sentence, but they need not be expressed; without being expressed they can still be a bundle of constraints the mental reality of which is the subject's experience of certain dispositions or inclinations due to a certain pattern of activation in his overall data base, according to the conceptual structure erected on the base. This pattern specifies the attitude. The experience ('das Erleben') due to the pattern realisation is the content of the attitude, if it is conscious. But not the pattern itself, nor the subject's awareness of it in self-experience or 'Erleben' is what the attitude is directed at. The attitude is directed towards an expected satisfaction constellation or situation in the world, and if that is not there or does not arise there, then still there is or has been the direction of fit and the expectation determined by the activation pattern. In this case the expectation, and thus also the intentional state, ends in disappointment instead of being fulfilled by an intentional object or situation in the world. This, of course can be a matter of degree. Some aspects might be satisfied, while others are not. Some aspects might be linguistically expressed, and others not. If publicly expressed aspects get satisfied, it might still be that unexpressed aspects connected with these are not fulfilled. Then the attitude is only partially fulfilled, though its public linguistic content is completely fulfilled. The minimal requirement for any rational attitude is that the set of constraints is such that there can be a possible satisfaction situation. This means that it must fit into the conceptual system *salva* stability, which implies that the bundle of constraints is internally consistent. It must also be externally consistent in the sense that it fits with the information one has about the world. Only then there is a real possibility of realisation of a satisfaction situation of the required kind. The set of constraints, furthermore, should be internally coherent and also externally coherent with information about the world as well as with goals of actions and activities, such that relevant predictions about courses of actions and events can be made. This coherence is a pragmatic coherence and depends on what is felt to be relevant given certain goals. The restrictions on a propositional content of an attitude, internal consistency (= conceptual adequacy), external consistency (= factual adequacy), and internal and external coherence (pragmatic adequacy), define different aspects of rational adequacy of a propositional attitude. An example of the important consequences of requiring internal coherence in a belief is that we can believe the following tautology, which for us is informative:

If Mary promises something, she does it or she doesn't do it. If we replace the tautology in the consequence by another tautology, for example, *it rains or it doesn't rain* the resulting sentence *If Mary promises something, it rains or it doesn't rain* still is a tautology, and therefore still true in all possible worlds, but it is non-informative. The interesting question is, why the first sentence is informative, though it is true in all possible worlds. In data semantics or update semantics (Veltman 1990), this sentence would be equivalent with *If Mary promises something, she might do it.* This stands in opposition to the possible information *If Mary promises something, she does it* and to the possible information *If Mary promises something, she doesn't do it.* The above tautology rules out these two possible informations and is therefore informative. In update semantics the above sentence, therefore, would not be a tautology. However, the same consideration holds for the examples we get by replacing the consequence with another tautology. It is the additional requirement of internal coherence which explains why *If Mary promises something, it rains or it doesn't rain* is not informative. We know that there is no general relationship between someone promising something and raining. But normally there is a relationship between promising something and doing it, and this background assumption establishes the coherence required. Our sentence is informative because it negates this relationship in the case of Mary, and we then may believe in Mary's irresponsible way of dealing with promises. Also the sentence *If Mary promises something it rains or doesn't rain* may be part of our information state in the sense that this sentence is true there; but it would not be something we believe in. Our conceptual structure does not contain that promising normally has anything to do with rain. Therefore, although such information is consistent with our conceptual system, and also with our knowledge (as long as it does not contain anything to the contrary), it is not coherent, and therefore under these conditions it is merely a (trivial) information that cannot be a candidate for becoming the content of a rational belief or desire.

Not all information we have, and in this sense know, is or will explicitly be believed. Believing implies awareness about what one believes and thus is an intentional state, but there can be lots of information in our data base and overall knowledge set we are not aware of, though it plays a role in safeguarding our behaviour. There is a connectedness by coherence by which parts of the knowledge set get enhanced such that they become relevant in determining what are possible beliefs. Only some parts of our data base are activated at a certain moment, and not all of the activation is strong enough to result into conscious belief. Activated parts may be more or less close to awareness and function in directing our behaviour. They afterwards, in a process of explaining our

behaviour as rational action, can be reconstructed and in this way be detected as unconscious motives and background belief. Activation happens because content is addressed by experiences, motoric activity, and internal attitude states. These get linked, under selecting perspectives, to sets of data which are themselves structured by similarity and contiguity.

We have seen that in partial semantics some sentences that are tautologies with respect to total models can carry information and are semantically distinguishable from each other because they may be evaluated differently with respect to partial information. Generally, also tautologies at least always carry the information that their presuppositions are fulfilled. We have noticed also that in order to understand the notion of internal coherence, we have to recur to conceptual structures. In order to show that conceptual semantics is fine grained enough for providing a reasonable notion of propositional content for propositional attitudes, I now shall discuss how tautologies can be distinguished in dynamic conceptual semantics. In doing this I shall compare the treatment here with a possible treatment in Situation Semantics (Barwise and Perry 1983), the so-called realistic brand of partial semantics. I choose, as simple examples, two pairs of tautologies and shall show how they are distinguished in dynamic conceptual semantics.

- (1) This is a tulip or this is not a tulip.
- (2) This is a cat or this is not a cat.

A difference between the two can be made in the procedures of understanding, and of verification accordingly. Sentence (1) means that a possible satisfaction situation either fits into the concept 'tulip' or it does not fit into it. Sentence (2) means that such a situation either fits into the concept 'cat' or it does not fit into it. We have here a test for (1), and a different test for (2): The sets of constraints which the satisfaction situations for these two sentences must fulfil are different. A set of constraints or conditions on possible satisfaction situations is a test for deciding, whether or not some situation is a satisfaction situation of a given type. Still, each world or situation that satisfies (1) also satisfies (2) and vice versa. On the realistic level of truth-evaluation the distinction cannot be made. But we can make, on the realistic level of Situation Semantics, a selection of possible satisfaction situations, parallel with what we can do on the conceptual level. If we would integrate perspectives of classification and judgement also on the sets of situations, like done by Seligman (1990), who treats perspectives as channels for selecting information, we would get a difference on the realistic level of evaluation with respect to situations: (1) is seen

under the perspective provided by the question *Which flower is this?* and (2) presupposes the perspective provided by *What animal is this?* This means that for (1) we select only situations in which flowers are at issue; for (2) we select situations where identification of animals is at issue. The selection is guided by the respective presuppositions of (1) and (2), which have to be assumed as being satisfied. Nevertheless, this does not provide for a difference in all cases. Where two classifications are placed under a single perspective, such a distinction cannot be made, because the same set of situations is selected as domain of evaluation. This is so in the following pair of examples:

- (1) This flower is a tulip or this is not a tulip.
- (3) This flower is a narcissus or it is not a narcissus.

Therefore also in Situation Semantics, even aided by the perspectival or presuppositional selection of the domain of evaluation, the distinction in meaning between two sentences cannot be made in terms of the sets of real situations that support the sentences; rather the difference must be found on an abstract level: the abstract situations, which are constructed partial situations, account for the difference. They are not more than the information that certain conditions, properties and relations, are realised or not realised at certain locations. Even if in our example the locations are the same, fixed by the same reference for the demonstrative, the information sets with respect to this location are different. So, abstract situations are like structured propositions, with the difference that they are not semantically represented by evaluation with respect to complete worlds, but with respect to parts of the world. Different abstract situations can be embedded into one real situation. A real situation is a locally delineated part of the real world, which is complete with respect to information which holds on that location. Actual situations are merely possible situations that are informationally complete on their location parameter. They are a limit of information growth about what is the case on a location.

For our pair (1)-(3) the distinction cannot be made on the level of semantic evaluation under perspectives and at locations. In all possible worlds, in all actual situations, but also in all real situations, under all relevant perspectives both sentences are true. The difference can only be made by taking into account the construction of the sentences. The constructions can be transformed into some regimented logical form, or hybrid form which includes evaluation at a location, as it is done in the format of abstract (i.e. partial) situations. The constructs amount to a disjunction of two abstract basic situations. These are different in (1) and (3), respectively. Therefore, (1) and (3) are different

tautologies. These hybrids, disjunctions of abstract situations, are included into semantics. But what their status is in a semantics that calls itself 'realistic', as Situations Semantics does, is unclear. Abstract situations are not simply parts of the real world. Only some that can be embedded into real situations are in this sense parts of the world. They are also not sets of (complete) possible worlds as propositions in intensional semantics are. They are merely abstract constructs, types constructed out of types, and even if the basic types in our respective constructions are fully spelled out in the language of set theory, the non-realised types are semantically nothing but the empty set. They are therefore just syntactic entities transposed into the *language* of set theory. As such they are publicly accessible merely like all expressions are. If we ask further what, besides their linguistic form, their semantics is, we realise that sentences (1) and (3) have the same semantic value and are thus semantically indistinguishable. This forces even those semanticists who adhere to a realistic semantics, to use abstract situations, like structured propositions, and to hold that there are in fact two kinds of semantics, extensional semantics proper, and a level of form where syntactic forms play a role, made up, allegedly, from semantic objects such as entities, properties, relations and truth values, though they are only made up from the names of these. Such are Barwise and Perry's abstract situations or situation types, and likewise Cresswell's structured propositions. They cannot be made up from real semantic objects, namely extensions in our or other worlds, since these cannot be contents of propositional attitudes.

The abstract situations of Situation Semantics would simply be logical forms of sentences in a formal language, if location would be included in the logical forms as an extra argument in every condition, and if the information that a relationship does not hold on the location would simply be expressed by negation on the condition. Then all properties and relations would be relations with location as one of its argument, and the partial semantics would be reduced to a complete extensional semantics, just like Davidson's (1967) logical form of action- or event-sentences. Of course, this loss of partiality is not wanted in truth-evaluation. But still, in retaining abstract situations within the semantics one houses there a kind of logical forms. The semantics is in fact two-levelled: It is an information 'semantics' on the level of abstract situations or types, and a realistic proper semantics on the level of real situations. The real situations are on their location complete, but partial with respect to the world as a whole. Sentences can be evaluated with respect to information, and the syntactically structured information can be 'evaluated' with respect to real situations by possible embedding into these. Our sentences (1) and (3) cannot be distinguished semantically on either level of evaluation.

If a semantics nevertheless wants to make a distinction between sentences (1) and (3), it has to take the test for truth into account, the procedure typical for constructivist semantics. Our sentences express a different test for truth, and in this sense they have a different semantics. Situation Semantics, if it includes a semantics of attitudes, has to be understood as a constructivist semantics, and this not only because the meaning of a sentence has to be constructed in a relationship between utterance situation and satisfaction situation. This relationship also is included in every other semantic theory that takes into account Tense and deictic expressions. Rather the constructivist character must be located in the syntactic-semantic construction of abstract situations. But what are these constructs ontologically, if they are not mental constructs? Realistic semantics is restricted to public constructs. But as public constructs they would have to be expressions of some public artificial language for building up abstract situations or structured propositions. Such expressions cannot be the contents of propositional attitudes, simply because propositional attitudes are not expressed publicly in this way. Of course, like any language, such syntactic expressions can be used to represent these contents. But to take the expressions themselves to be the contents of propositional attitudes would amount to the quotational approach, where to believe that *p* simply is analysed as believing a sentence that expresses *p*. The unacceptable consequence would be that we always have to believe something in some language or other. This would not pose an obstacle for an adherent of an innate language of thought; the language of abstract situations or structured propositions would have to be the universal language that philosophers have always dreamt of. The quotational approach, if it were not committed to a universal context-independent language, also suffers from the fact that sentences can have different meanings, depending on the contexts, due to pronouns and indexical expressions, polysemy and homonymy. These phenomena force upon us the conclusion that the belief content must be the meaning, rather than the sentence itself. The universal language would have to be fully context-independent and theory-independent, which would require to even assume a universal theory about the world into which all other theories can be embedded or at least translated. This old philosopher's dream misses any support from the history of science and culture, and is at odds with the human epistemic condition, which is characterised by a variety of interests, perspectives and ways of life, on which theories depend.

For a semantics of propositional attitudes we can do with a weaker constructivist semantics than the one of partial semantics sketched above, or the one of Dummett (1976). There is no need to claim, like Dummett does, that a sentence at a certain time can only have a truth value if the test procedure can be carried

out. Like all sentences, sentences about the future and universally quantified sentences can be true, without us having made the necessary tests for verification, though we cannot really be sure whether the sentence is true. Dynamic conceptual semantics accepts a realistic notion of truth for interpretation by distinguishing the epistemic concept of holding true on the basis of proof and justification from the realistic concept of truth, which we presuppose in our notion of holding true. The constructivist approach in a theory of understanding does not preclude a realistic approach in the theory of interpretation. A complete semantic theory requires to take into account the interdependence of both aspects, understanding and interpretation, the latter as extensional evaluation, and especially truth evaluation of expressions in a world. All we really need to accept from constructivists' insights, is that in a semantics of attitude reports the propositional content is a specification of the test for judging satisfaction of the attitude. A propositional attitude constrains partly the specifications of the methods or tests of its satisfaction. And so the semantics of propositional attitude reports must contain a notion of proposition as such a partial test specification. The test is not specified in its full procedure. Several procedures are still possible, but they are constrained by the specification. This specification is provided in dynamic conceptual semantics on the level of conceptual structures on growing sets of data.

To view the proposition in attitudinal contexts as a test-type is compatible with taking it as an intension in contexts of modal operators. A test-type is also a function from possible situations or worlds to truth values, but it is a function that partially specifies a procedure. The function is not just specified by its extension. This procedure-related view is taken in dynamic conceptual semantics, which is a semantics of understanding and concept formation. Our method of interpretation, i.e. our evaluation of utterances with respect to the world, is dependent on our understanding and its results are partial. But to understand these as partial presupposes to understand them as being part of, i.e. being embeddible into, a complete interpretation, where also universally quantified sentences and sentences about the future have a truth value. A theory of understanding has to be based on the role understanding plays in language to world interpretation. It need not to be anti-realistic in the sense of Dummett, rather it presupposes realism by accepting reality as providing the objective constraints on concept formation, which figure besides the intersubjective constraints. Concept formation is constrained by truth and norm. The dynamic theory of concept formation and understanding is, in this sense, based on semantics proper. To the extent of excluding attitude contexts, natural language interpretation can be analysed by a purely extensional method (whereby

intensions are thought of as relations between possible worlds or indices with extensions there) and thus is, and had to be developed independently of understanding. Understanding cannot be understood without being modelled in its role in interpretation, i.e. in evaluation with respect to worlds or situations. Meta-theoretically therefore, a model-theoretic or realistic theory of public interpretation is prior to the theory of understanding. Only through the restrictions implied by its role in interpretation can we analyse understanding in a semantically interesting way.

On the level of dynamic conceptual semantics or understanding we do not take as semantic values of expressions the sets of situations or worlds that satisfy them, but the constraints on possible satisfaction situations themselves. For different tautologies, the sets of constraints are different, but the worlds that satisfy them are the same, and so are the satisfying situations, if the expressions are about these. This is the case if their presuppositions are fulfilled. Thus, though tautologies are equivalent in terms of truth, their truth conditions are different. Difference in meaning is thus consistent with equivalence in truth, even in all possible worlds or in all situations about which two different sentences, such as (1) *This flower is a tulip or it is not a tulip* and (3) *This flower is a narcissus or it is not a narcissus*, are expressible. Herewith, meaning is different from intension in the technical sense of possible world semantics, and from evaluation with respect to information generally. Meaning of sentence (1) includes that the situation referred to by the demonstrative is matched *salva stability* with a conceptually complete, i.e. stabilised, set of satisfaction situations for *tulip*, while the meaning of (3) includes that it is matched *salva stability* with such a set for *narcissus*. In this match, every situation, if it is referred to by the demonstrative, comes out as a satisfaction situation because it either matches *salva stability* with the set of 'tulip' situations' or it does not. Likewise, according to (3), we match situations referred to with the set of 'narcissus' situations, and they all make (3) true, because they match *salva stability* with this set or not. Thus (1) and (3) have the same realistic semantics in terms of sets of satisfaction situations, but their tests for that differ.

The crucial point about the ontological status of this meaning, or test-specification, is that there cannot be a matching to the complete set of 'tulip' satisfaction situations or 'narcissus' satisfaction situations, because they are never available in the matching process. On the conceptual level we have as conceptually complete sets those sets of experiences of previous satisfaction situations, for which the internal similarity degree does not decrease with further extensions. This means that with respect to defining the use of the expression, here *tulip* and *narcissus* respectively, all further extensions are treated as

redundant. This is what it means that matching, or adding new situations, must be performed *salva* stability of the concept. To judge whether something fits into a set *salva* conceptual stability is a cognitive operation, which can be thought of as a result emerging from automatic matching processes on the sub-cognitive level of content-addressable connectionist patterns. As far as this judging is done publicly, social control presses towards group conformity of the measure of stability, but each participant of the group has to make his own judgement of stability in the view of the measure he handles at that moment in his process of concept formation. Matching can only be done against his own concepts, i.e. his systematisations of his experiences. The results can be made public and thus be tested against the results of other members of the group. This can give rise to corrections in his conceptualisation, even to such an extent that re-systematisation of his experiences may be required and induced accordingly.

The ontological level of meaning cannot be a level that merely contains (experiences of) satisfaction situations and sets of these, but it must contain sets that each are held together under a notion of stability. These sets are concepts on the experiential level, corresponding to sets on the realistic level held together in public agreement by the individual language users' sense of stability, which by social pressure conforms to more or less the same measures. The notion of matching or testing includes as its social element that the possible outcomes of these procedures have to be acceptable for nearly all members of the speech community. In other words, meaning is not simply based on the notion of truth and satisfaction, but it contains all aspects which are required for judgements about truth and satisfaction. These are the cognitive requirement of matching under conceptual stability, the social requirement of having the same measures of stability, i.e. the pressure towards forming the same concepts, and the requirement of reaching a realistic result in interpretation, an object or truth value. Summarising, meaning includes the subjectivity of judging conceptual stability against one's own conceptual system, the intersubjectivity of conformity in the public results of these judgements, and the objectivity of reference towards objects, real situations, and truth values.

We still can adhere to the old slogan "The meaning of a sentence utterance are its truth conditions", though having the same truth conditions does not, of course, boil down to having the same extension, and even not to having the same intension, i.e. the same extension in every possible world. Different tautologies have the same extensions and intensions, but have different truth conditions. These truth conditions are equivalent under the point of view of truth, but they are not equivalent under the point of view of matching with concepts, the operations necessary for establishing truth values for sentences.

On what kind of ontological level do truth and satisfaction conditions exist? It cannot simply be the realistic ontology with all the set theoretic constructs of intensional semantics, nor that of real objects and real situations and set theoretic constructs on these. But could it be the latter together with publicly agreed upon operations about testing? Social control only pertains to the results of these operations in terms of accepting or rejecting situations as satisfaction situations for public utterances. There are not really operations involved in testing which are regularly publicly controlled step by step. Use of language is publicly judged as correct or incorrect not because syntactic-semantic operations are followed correctly or not, but because of the public results of the privately performed operations. The whole testing as matching new data under the stability of a measure of internal similarity of sets over old data, and their extension by more data *salva* stability, can be described as cognitive operations, for a large part understandable as emerging from subcognitive processes. These processes and the resulting operations are not public and are not themselves socially controlled, only their outcomes are. The public requirements on the outcomes are restrictions on the cognitive operations. The operations have to meet the restrictions that are defined in terms of public relations between utterances and constellations of situations, especially in terms of correlations between utterances and satisfaction situations. But these by themselves, without including the cognitive operation of matching *salva* stability, are not enough. This matching everybody has to do for himself, though the result is publicly approved of, or it is corrected. This approval or correction is based on nothing more than the subjective results of the conceptual matching operations of the norm giving and norm promoting members of the speech community, which are, in non-revolutionary situations, the established members against the newcomers, and especially those that have to teach against those that have to learn.

Concluding this discussion, the ontological realm of meanings, as they function in attitudes and understanding, is the cognitive level, which is a construct of our, the theoreticians, theorising about propositional attitudes, concept formation, and understanding. Certain basic theoretical terms on this level of theory construction, such as similarity-measure and stabilisation, are structurally defined on this level by pointing out their role in the theory of concept formation, but they have to be understood as having subcognitive foundations which can now to some extent be modelled in connectionist theories. However, everything we find on the cognitive level, such as the meanings of linguistic expressions, .e. concepts understood by the theoretician in terms of orderings on growing sets of data, is restricted by public requirements of intersubjectivity, i.e. social conformity, and by the objectivity of the results of applying expressions

according to these meanings. Objectivity means that truth and satisfaction have to be assigned to utterances depending on the state of the world. Intersubjectivity means that this truth assignment to utterances on the basis of the state of the world has to be such that the community can agree upon it, given their conceptual systems by means of which they can understand the state of the world as being such that it supports the judgement. Agreement in judgements and objectivity of judgements are the public restrictions on concepts, which by themselves are subjective mental entities. These subjective entities are not more than sets of constraints on possible satisfaction situations for utterances and also for propositional attitudes. These constraints require the integratability of possible satisfaction situations into the subject's conceptual systematisation salva stability. The public results of integratability tests are compared to the results of the tests performed by other members of the group, and in case of mismatch between the members' results, the conceptual systematisations of persons who are in the social position of learners have to be adapted to the new data provided by public disapproval and correction.

5. Propositional attitudes and partial conceptual systems

In what follows I shall return to the analysis of propositional attitudes, firstly to give a partial definition of belief and desire, discuss some of their logical properties and the problem of deductive closure of sets of beliefs. Beliefs are not conjoined with the whole knowledge set someone has available as background information; rather they are positioned in certain neighbourhoods of experiences under certain perspectives, which select from these neighbourhoods aspects which figure in conjunction with the belief, or rather as non-expressed parts of the belief. These parts are, as associated contents, directly addressed by the belief contents. This partial knowledge set consisting of experiential and linguistically explicit knowledge determines what is implied by a belief in terms of further belief, preferences, and action routines. I shall then return to the problem of speaking about identity of beliefs, now under the point of view of inferences and predictions deducible from beliefs. Thirdly I shall discuss the problem of understanding utterances and reports about beliefs of persons who seem to have a conceptual system that partly deviates from our own system, or from what is considered to be commonly established linguistic practice.

We give a partial definition of the meaning of belief and desire reports by what counts as satisfaction of the attitudes. The partiality of the definition is due to the

fact that we can say what a satisfaction situation of an intention is and we can point towards such a situation, but we cannot say what an intention is, nor can we point towards it. We have to take the notion of intention as a primitive, and we can characterise it implicitly by all kinds of relationships in which intentions generally and specifically stand. This is standard in the philosophy of Intentionality (Searle 1983). Especially we can characterise intentions by their preconditions and consequences.

x believes that s is true in w_0 if and only if *x* is in the intentional state B_s in w_0 , whereby in w_0 the intentional state B_s of *x* is satisfied if and only if the set of constraints S , expressible by *s*, is adequate with respect to w_0 : ($S \rightarrow w_0$). Satisfaction means that there is a situation in w_0 which satisfies the constraints S expressed by *s*.

x desires that s is true in w_0 = *x* is in the intentional state D_s in w_0 , whereby the intentional state D_s of *x* is satisfied if and only if w_0 is made adequate with respect to S : ($w_0 \rightarrow S$). Satisfaction means that there is brought about a situation in w_0 which satisfies the constraints S expressed by *s*.

Besides a knowledge base also a routine base plays a role in cognition, especially as a background for the preconditions of having desires. This base contains routines of activities and actions. They are formed analogously to concepts. Routines of actions and activities are like concepts as far as they can be represented as growing sets of ones own conscious and unconscious action and activity experiences, under stabilisation according to success and similarity measures. Learning a routine consists in a stabilisation process in successful exercising behaviour under correction by failure. Together with the knowledge base and the conceptual structures erected on it, the routine base provides the side-conditions for actualising routines. We can construct new routines from old ones like we can compose new concepts from old ones, and we can devise modifications of routines like we can form new concepts by extending a subset of a concept representing set of situations in a new manner by similarity under a different perspective. If some routines which we tend to actualise cannot be completed or even started because certain side-conditions, expressible by *s*, are not fulfilled in our world, then we feel or say that we desire that *s*. Desires are constituted by tendencies to actualise certain routines for which preconditions must be fulfilled which are not true in the knowledge base. Then there is a need which causes the desire that *s* be the case. Following suggestions by Searle (1983), other propositional attitudes can be construed out of the two basic

attitudes of belief and desire in conjunction with additional conditions concerning emotions and social relationships, relevant with respect to the basic attitudes involved.

It has been pointed out above that the analysis of attitudes is such that the propositional content is not the object of the belief or desire. In case there exists an object which is intended, it is the satisfaction situation of the attitude. However, there is no abstract intentional object that also would be the object in case the attitude is not satisfied. There merely is the set of constraints of possible satisfactions situations and the intended direction of fit, both as specific characteristics of the intentional state itself. They are not the object of the intention. Rather they are the way in which the intention is formed, and in this sense they are the form or internal structure of the intention. This mainly formal specification determines what kind of concrete contents will fulfil the intention. They, so to speak, have to fit into this form via the causal connections provided by the experiential basis. However, this way of describing the relationship between an intention and its possible object is terminologically unfortunate: What traditionally has been called propositional "content" is in fact the form or structure of the intention, besides the direction of fit and other properties of the intentional state or act. It has been called "content" in opposition to the character of the intentional state, being belief, desire, hope, etc. The state or act character has traditionally been viewed as the form to which the proposition is the specific content. I will not discard this terminology. Nevertheless, if we view the intention as a formally specified state or act, the real content can only come from the external object that satisfies the intention. In a certain sense, all mental aspects are formal, though with a causal bond connecting them with outside objects. This bond is provided by the experiential level of concept formation. By this causal bond, concepts have a discriminatory relationship towards externally caused stimulations and thus are causally related to objective contents. The conscious effects of the stimulations, i.e. our experiences or 'Erlebnisse', we can call "subjective contents". The propositionally determined aspects of form, or structure, of the intentional state consist in all the connections made with the sets of experiential data or "subjective contents". Which causal connections are involved in the relationship between the intentional state and its satisfaction situation is defined in form of a collection of constraints for the possible satisfaction situations. Because of the role of previous experiential data in the constraints, we can speak of a subjective "content" that gets activated with the structural properties of the so-called propositional content.

If S is the propositional content specification of a belief or desire and it is not the attitude's object, then it must be the way in which the intentional state or attitude

is structured. The structure of an intention namely can never be its object. The object has to be outside, if it exists or comes to exist. But on a higher level, a belief can be an object of a secondary belief: $x B_2 x B_1 s$ is the form of a second order attitude report, for example of *John believes that he believes that s*. Here the collection of constraints expressed by $x B_1 s$ is the internal proposition-induced structure of B_2 . Such a reflexive secondary belief cannot be false, it always has an object, namely the intentional state $B_1 S$ of x . A secondary belief about one's own belief is not informative and therefore is not a knowledge state that likely is a candidate for an attitude of belief to be directed to.

If we would artificially put ourselves into an attitude which is such a secondary belief, the following would hold: $B_2 B_1 s \rightarrow B_1 s$, but not the opposite. If I have the attitude of believing that s , I need not to also have the secondary attitude of believing that I believe that s . Of course, I always can have it, and if I have it, it is satisfied. Belief is an attitude in which a propositional content is actualised in activation, while having the information that s does not imply actualisation. An attitude includes an activation of some pattern, and for belief it must belong to the information set. Thus, not everything that is part of the data base or knowledge set is believed. It is also not generally the case that if 'if p then q ' is in the knowledge set and ' p ' is believed then also ' q ' is believed. Only if the conditional 'if p then q ' is itself believed and ' p ' is believed then ' q ' is believed in the sense that if ' p ' is put under attention of the believer, he will agree that he believes it also. Also a logic of beliefs of the kind of Hintikka's where x believes that p is true if and only if in all worlds that are compatible with x 's beliefs (i.e. in all epistemic alternatives) p is true, has the consequence that all tautologies are believed and that all consequences of something believed are also believed. This logic seems better suited for the notion 'having information that', because this notion does not include that the one who has the information has been aware, or is aware, of having it. For belief it holds that if x believes that p he also has the information that p . Believing that p is an attitude, but having information that p is not. It merely is a more or less permanent state which is the result of having accepted p or having accepted some information that implies p .

Like other verbs, we use the verb *believe* also in an habitual interpretation. We say that John walks if we mean that there are occasions on which he has been walking and there will be situations furtheron in which he will be walking. Likewise, if someone loves somebody he will not always be in the state of actually entertaining this feeling, but there have been and there will be such actualisations. If there are suitable occasions, the actualisation will be called upon. It is not different with belief. If someone believes that p , there have been and there will be situations in which this belief will be actualised.

It has always been a problem for logical analyses of belief sentences that they predict that with believing something also all its consequences are believed. This is, of course, not in accordance with the facts. In dealing with beliefs within the framework of dynamic conceptual semantics we can make clear how far logical closure goes. It is a question of how far an actualisation of a belief goes, whereby the main point is that information cannot be presupposed to be discrete in a way fixed once and for all. Rather it always can be made discrete in different manners and to different degrees. Logical closure is always relative to a certain manner and degree of discreteness. It is not a notion that can be applied to information per se. To substantiate this claim we have to be more explicit about the nature of our knowledge base.

The basis of our knowledge consists of experiences of situations. These complex perceptual and motoric experiences are primarily scanned by our bodily organs and parts of the brain connected with these. Experiences are made in forms which are provided by pre-conceptual similarity spaces and their ways of measuring input. Whether experiences are discrete on this primary level we cannot say, I guess that they are most probably not. Anyhow, on a conceptual level these experiences of situations, to begin with, are not discrete. This is so, because there are no concepts to begin with in terms of which the experiences can be analysed and then be represented in a discrete fashion and thus become part of our memory. Only by operations of concept formation, i.e. by integrating, under certain perspectives, the experiences into similarity sets and contrast sets of previous situational experiences, the new data become analysed into parts. They thus become discrete as far as the analysis goes. Each analysis of experiential data consists in relating these to some parts of the previous data in the way that the new data are integrated within, and put into contrast with similarity sets on the previous data. A new experience under some perspective can also be the beginning of singling out new similarity and contrast sets formed by old and new data, and thus can give rise to forming new concepts. Standardised ways of analysis, triggered by certain perspectives under which situations are perceived, can directly be integrated into the perceptual activity as additional similarity spaces by which the data are scanned. Depending on which subsets of previous experiences are involved in analysis, we get a different manner of discreteness. The manner of discreteness can be extended to different degrees and in different directions, depending on how fine similarity sets and contrasts are made and under what perspective they are made.

Within the information we have gathered on the experiential level, there always is implicit information, i.e. information that has not yet been made explicit by analysis; it has not yet become part of a conceptual structure and therefore it

cannot be part of our potentially conscious memory. Nevertheless it might pre-conceptually be singled out on the neuro-physiological level by being connected with certain motoric routines and thus play a role in activities and actions we perform. This implicit information is not something we can believe or remember. It belongs to our 'know how' and not to our 'know that'. For being able to make it the content of a conscious attitude it first has to be captured in an analytic way, i.e. it has to be recognisable as a part in a certain manner and degree of discreteness on the conceptual level. Also particular information gathered from reports and stories which, in understanding, has been integrated with our experiential level, contains implicit information. It is included within the information incited in our knowledge base in understanding the text. Furthermore, per occasion in which we get a certain piece of information, much information is kept implicit. It can, in principle, be made explicit by relating it to further subsets of previous or even later experiences. This means it becomes analysed conceptually, and thus it gets understood. A logic of information can only be applied to discrete data, i.e. it can only be applied to explicit information. This has consequences for logical closure.

A substructure of our conceptual system used in the analysis of data makes that these are discrete relative to it. If we have a language that can express the concepts involved, we can formulate logical consequence in this language. Now, ideally all implications from a set of beliefs P, expressed in this language, should also be believed. The problem, of course, is predicting from the use of a portion of a certain language for the belief report, how far the conceptual system is indeed used in the analysis. Especially it is undetermined whether, in the analysis of his information at a certain moment, somebody uses all or some implicational relationships to which he once has subscribed, or sometimes will subscribe. If someone uses standard English at a certain occasion to express a certain belief, we cannot predict that he uses, at that occasion, or generally, all categorisations that can commonly be implied by the use of standard English. He always will make some or other selection, and this determines the discreteness of his information at that occasion. So by hearing him express certain beliefs in English we may conclude that he must also have certain implicit information, but it may not be explicit to him and thus cannot be content of his belief-attitude, though it may be part of his information on the experiential level. Only if we would exactly know the manner and degree of discreteness of a person's information at a certain occasion, could we make predictions about which implications his belief has for him. Of course, in artificial knowledge bases with only discrete information we have logical closure of the information. But if pictures are part of the base then, of course, the manner and degree of analysis,

i.e. the discreteness, has to be determined in which the information stored by the pictures gets broken down. Logical implication, or inclusion between aspects of the pictures, is likewise possible relative to this discreteness.

The basis of our inferences is not that we use logical inference rules on sets of sentences. Rather the structural relationships, especially structural inclusions on the whole knowledge set are at issue. On its experiential base, inclusions are structurally discernible, which may give rise to linguistic statements that afterwards can be seen in terms of inferences. Only on the linguistic level, logical inference rules play a role directly, and this not so much in that we would use them to draw inferences, rather they play a role in judging correctness of the inferences arrived at on the more basic level of structural inclusion.

An example of different manners and degrees of discreteness of information is the following: I believe that a certain path is more than three kilometres long. Does this imply that I then also believe that it is more than two kilometres fifty metres and five centimetres long? No. The degree of discreteness of my information about the path is in kilometres, not in metres and centimetres. Of course, I can also use the finer measures, but I normally do not, if I estimate the length of a path. In fact, I normally do not measure the paths I walk regularly in kilometres at all. I measure them in time. And then I do not measure in seconds or minutes, but in hours, half hours and quarters of an hour. This is an other manner of having information about the length of paths. It presupposes that I have a standard speed, which I automatically realise when walking in the dunes. If I believe that this path is forty-five minutes long, do I then also believe that it is about four kilometres long? I normally do not, even if I know that my walking speed is about five kilometres per hour. Only if I would really calculate, I can transfer the information 'I shall walk path A' between the two ways of measuring. Normally, I would conclude that I shall walk about forty-five minutes, without concluding that I shall walk about four kilometres. My beliefs about the paths' length are in time and not in distances. My perspective in these judgements is time, not space, and the perspective selects the activation domain directly in addressing the content.

A further application of conceptual semantics to a logical problem is the defeatability of default conclusions. This problem is treated quite a lot in information- or belief logics. However, I want to show that it is due to the internal diversification of conceptual structures belonging to a single term. If (all) birds can fly, and Twiggy is a bird, then Twiggy can fly. But suppose, Twiggy is a penguin. Then it cannot fly. This sequence results in a contradiction. From the point of view of dynamic conceptual semantics, a change in concept has taken place. In the sentence *(All) birds can fly*, the term *birds* is used in its every

day usage for normal birds, based on our experience with flying birds and their bodily proportions. Penguins, small chickens and very fat geese, and also ostriches are not included. And also disabled birds are at the periphery of what we call a bird; in any case it is not a normal bird. In the sentence *Twiggy is a bird*, the term *bird* is used in the biological sense, i.e. as a formal concept of biology, based on biological taxonomic theory. That notion does not include that birds can fly, precisely because there are creatures that are, biologically speaking, birds, which cannot fly. What happens in the above sequence of statements is that the conceptual impact of the term changes; it expresses a different, though related concept. Our experiential concept of a bird is that of a normal bird, and it is not the same as the theoretical concept from biology. In our language we handle, depending on occasion, the one or the other concept. In everyday situations we would not say that penguins are birds. If suddenly a penguin, or an ostrich is in my garden, next to five birds, I would not say that there are six birds in my garden. No, there are five birds and a penguin, or five birds and an ostrich. If someone grows up in the Antarctica and only knows seagulls and penguins he probably would not use the term *bird*, or he would use it instead of *seagull*, for a concept possibly narrower than ours. If he would use the term *bird* for seagulls and penguins together then, of course, he would never believe that birds generally can fly. Thus the problem of defeatability of default sentences is due to interference between different conceptual structures. Here, of course, logic breaks down, since the term *bird* is used ambiguously in a single piece of text. Either the sentences each have to be interpreted relative to the partial conceptual structure by which the term *bird* is used. Then the above sequence of argumentation vanishes, and the logical problem vanishes too, because it is not about the same extension of birds. Or we have to use different terms for the everyday notion of birds and the biological notion. Then no logical problem can arise. The fault arises only if we apply logic to everyday language without taking the precautions which are necessary because of the context-dependence of meanings of lexical items. It therefore it is not advisable to mess up logic by devising complicated default logics. Rather we can take care of the fact by realising that a single term can have different meanings in different context-types, and in that case the premises are about different extensions when the term *bird* is used, representable by indexed terms *bird₁* and *bird₂*, and therefore the inference does not go through.

Under the point of view of implications we now return to the different notions of having the same belief. They can be localised in different ways in the combined system of understanding and interpretation. If two people believe that a certain Fred shall come and their beliefs both are expressed by *Fred will come*, they

believe the same referentially and conceptually as far as their common language goes. Both may have had different experiences with Fred and thus one is glad and the other is worried about Fred's arrival. Then the belief that Fred will come has different implications for both. On their different backgrounds they draw different conclusions and thus the belief in this sense means something different to both. One believes that everything will be fine, the other believes that everything will be a mess. If both believe about the same person *My best friend comes*, of course, their beliefs are the same as far as common language goes, which implies in this case that the conclusions are the same for a large part: Both are happy and expect everything good of the arrival of the person.

If both think *My best friend will come* about different persons, they merely believe the same kind of thing, but they do not believe the same. The conclusions they draw can be very similar as far as the types of situations go that are expected, though they are about different situations. In terms of general concepts involved, they believe the same, but in terms of individual concepts, and with that with respect to referents involved, they believe something different. Since the individual concepts of the two involved referents are different also on the conceptual level, there is a difference in belief. But the type of individual concept used in the acts of reference, expressed by *My best friend*, is the same. Thus on the level of expressed types of individual concepts, and herewith, on the level of general concepts their beliefs are identical. Additionally, there may be differences between their respective beliefs due to the very different histories of experiences both have had with their respective friends.

If two people believe about the same person that he will come, without knowing that they refer to the same person, they do not have the same belief. If, about the same man, John believes that his brother will come and Bill believes that the man will come who bribed him three years ago, then they will apply quite different conceptual structures in the analysis of the same real situation. The different partial situations both people expect coincide objectively in a real situation. If they get to know that they refer to the same man, though in a different way, each can extend his respective partial situation such that it comprises the other's partial situation. In this way, they objectivise their perception of the situation to more than their own personal perspective. After such a step we can say that they know that they believe the same situation to come about, though it has quite different implications for each of them. In this sense it still is a different situation to them. Different aspects of the same situation are relevant: It is analysed in relation to a different set of previous data with a different conceptual structure on them. The manner of discreteness, or analysis, is different for them and herewith they draw

different conclusions about the situation. For one it will be a friendly reception, for the other it will be the opposite.

Putnam's (1975) example of water on Earth and Twin-earth and what the people on the two planets believe about their water belongs into this series of examples. Both people believe that water is wet, and many other things about how water is used in daily life. Do they believe the same if the water on Earth is H₂O and the water on Twin-earth XYZ? The term *water* refers to two different substances which look and behave quite alike. If the people do not know yet the scientific concept of water they have conceptually the same belief. Their experiential concepts are identical. If they already have available the respective scientific concept, their beliefs are distinct if the respective concept is taken into account. This is not necessarily so, because a belief need not be connected to the whole conceptual system, but can be connected just to some part of it. In the everyday context of washing their hands they will have the same beliefs about water, even if their scientific concepts differ. These just don't matter there. Referentially, they accidentally believe something different, if they do not visit each others world. But if they visit each other, they normally will take the water of the other as water like their own water and handle a split extension for water, without making any problem about this mixed extension. Only in scientific contexts will they distinguish two extensions for the respective scientific terms, and for their respective scientific use of the term *water*. In scientific situations they believe something different when they believe that water is wet. Not only is there a referential difference, but also conceptually the belief is different, since they here use the term *water* to refer to something by means of the concepts H₂O and XYZ, respectively. In deciding securely whether a substance is of the same kind as the water they referred to up to now, they have each to use their respective scientific concept. Against Putnam, I stress that concepts determine reference: To which things we refer depends on the concept we entertain at an occasion, be it the everyday concept or the scientific concept. Still we can uphold, with Putnam, that under the perspective of natural kind identification, the development of the scientific concept depends on what we find out about the things we take to be its referents, under the scientifically relevant perspective.

We now turn to the last topic of this section. Many years ago, Barbara Partee put forward a problem about belief reports that arises when there exist differences between the lexica of members of a speech community such that certain lexical items seem not to have the same meaning for different language users. This way of presenting the problem is, of course, already an interpretation of it, but in some way or other it has to be presented in its generality. The problem becomes obvious when one person utters a sentence which to others seems to contain a

lexical contradiction or violates one of the so-called selection restrictions. Suppose the child Mary believes that clouds are alive. The report *Mary believes that clouds are alive* is problematic because we must ask ourselves in which language it is put. In Mary's language the sentence is okay, but in our common language we now doubt whether for Mary the word *alive* has a different meaning than it has for us, or whether the word *cloud* has a different meaning for her. Anyhow, she either uses the word *alive* or the word *cloud* differently than we do. We cannot decide which is the case, or whether both are the case. If we would use the words according to the public meaning, her belief would be a contradiction, i.e. impossible. A difference in belief between her and us is at the same time a difference in meaning. This implies that we cannot represent her belief in our public language. Although put into a single language, English, this example nicely illustrates how interpretation and translation can be impossible if different conceptual schemes are involved. Either we have to report in Mary's English, or we have to attribute a contradiction to her, which she did not commit and which Quine's (1960) principle of charity would prevent. Even agreeing in reference with her, we are stuck with an indeterminacy of possible boundaries between meaning and belief. It is, of course, Quine's point that we cannot draw any sharp boundaries here.

To some extent Mary seems to speak another language than we do. On the other hand, we see that she is able to identify clouds correctly; she is able to distinguish them from other things and applies the word *cloud* correctly. She also applies the word *alive* generally to living things and not to others, with, up to now, the only exception that she applies it to clouds. The only peculiar thing is that clouds sailing along the sky fit, *salva stabilitate*, into her set of experiences of satisfaction situations for *cloud*, and they also fit into her set of experiences of satisfaction situations for utterances containing the word *alive*. There is no reason to doubt that her experiential concept of 'cloud' is the same as ours. Then there are two possible explanations: There can be a theoretical difference; she might have a theory in which she attributes the clouds' movement and their slowly changing forms to some hidden property, namely that they are alive, whereby she has the same experiential concept of 'alive' as we do. This can be corrected by correcting her theory, and thereby we correct her theoretical concept of what a cloud is, though the connected experiential concept is like ours and has served to determine the common referents. The other explanation is that her experiential concept of 'alive' is different from ours. It must be broader than ours, an example of over-generalisation in concept formation. This can be corrected by telling her that clouds are not alive because they are not called *alive*. But there is also a third manner of reaction, namely to assume that she uses the

adjective *alive* metaphorically. This would mean that we accept her statement that she loves clouds because they are alive, and we understand it by creating within our conceptual system a new use of the adjective *alive*, namely for characterising the clouds' ways of moving and changing form. Thus we extend our conceptual structure with respect to the word *alive* from a concept 'alive' to a set of concepts, a polysemic complex, for this word, which also contains Mary's use of the word. This amounts to a linguistic change in our common language, probably restricted to situations of communication with Mary. In order to mark the case of metaphoric use for the sake of clarity about our common conceptual system, we might additionally point out to her that clouds are not really alive like animals are, but that she is right that they look a little bit so as if they were alive like animals, and that therefore she is right to say that they are alive. Additionally we stress that *alive* for clouds means something different than *alive* for animals.

- We would not correct our own conceptual system by changing our concepts, because Mary is a child and has to learn the language, while we are in the position of teachers.-

The different roads we can take amount to either correcting her theoretical concepts, or to correcting her experiential concepts, or to reorganising our own conceptual system such that it agrees with hers. The last we do without changing our own concepts, rather we create polysemy. The first two reactions, namely correction, are conservative; the last reaction is progressive, i.e. being open to change. I shall not evaluate these solutions, but merely point out that one of these roads has to be taken in order to make communication possible, i.e. in order to employ a common language. Only within a common language can we correctly make attitude reports about each other, and can we understand each other. Correction, i.e. making the other adapt to our language use, or adapting our language use to that of others, is a necessary meta-linguistic activity. Adaptation to others' language use does not necessarily mean that we have to change our beliefs. If we now agree with Mary about the truth of her sentence *Clouds are alive* in the context of being fascinated by their movement and changing forms, we have not changed our belief with respect to clouds; we have not changed our concept of 'alive'. We merely have changed our language by adding another concept, also expressed by the word *alive* and restricted in its use to the kinds of contexts of admiring the lively cloud formations. We continue this line when we say that the sea is alive, because of its rolling waves, or that a cornfield is alive in the wind. This addition by metaphorical use amounts to continuing a subset of our set of satisfaction situations for the use of *alive* by adding cloud-situations, rolling wave situations and possibly others. This new

growing set then represents the new concept expressed by *alive*, which together with the old one forms a polysemic complex.

6. Conclusion

In the preceding considerations it has been shown that what people can believe depends on the conceptual structures they have established and are still establishing on their experiential and theoretical data. A belief at a certain occasion is formed on a partial structure of the system. Rather than complete holism, partiality plays a role in the use of terms and also in the formation of beliefs. This also is the prerequisite for making attitude reports about other people. Because parts of people's conceptual systems are identical they can agree in beliefs which are based on these, and they can make and understand reports about other's beliefs so far. Attitude contents are always based on some part of the system: From here the satisfaction conditions are gathered in a certain way. The way of gathering the constraints on possible satisfaction conditions consists in selecting certain (quasi-)concepts, i.e. sets of satisfaction situations, into which the possible satisfaction situation has to fit *salva stabilisation*. The syntactic structure of the possible satisfaction situation consists in procedures of forming certain syntheses of the constraints on basic situations. They determine the fit of the possible satisfaction situations into the respective concepts, *salva stability*. On the level of dynamic conceptual semantics, a level of data organisation, meanings and the contents of propositional attitudes especially are such procedural constructions, which specify what the corresponding satisfaction situations must be like when they are experienced.

It has been pointed out that logical closure is not valid for sets of beliefs because of the partiality in our use of our conceptual system. Our individual and general concepts are rooted and embedded in neighbourhoods of classificatory and historical organisation of our experiential and linguistically expressed knowledge. Metaphorically speaking, our knowledge seems to be a kind of topological structure on our, always growing, set of experiences. The topology is formed by historically organised neighbourhoods of partial individual concepts and situations, whereby the whole topological space acquires an overall uniformity by a kind of metric, established as a structure of classificatory and ordering relationships, in terms of general concepts. This structure is extensionally built up on the set of experiences, whereby topological sub-spaces contain more specific sub-concepts of the overall concepts, as parts of the overall

extensions. Thus in neighbourhoods of certain experiences, in certain partial histories, under the perspective of certain desires and preferences, and also in theories, and under special perspectives generally, not only our partial individual concepts but also our general concepts have a more specific content because they amount to, in these contexts, to constituting stable extensions or continuations of the sets of those data which are selected under these circumstances. The smaller the selection the more specific the conceptual content is, compared with the conceptual content of a more globalised set of data. Traditionally, the addition of associated content to the so-called lexical content in special circumstances has been called 'connotation'. Since in a theory of dynamic concepts there is no essential difference between lexical meaning and connotations, we simply speak of a normal, more specific, sub-concept of the whole concept. A sub-concept comprises the so-called lexical content and connotational content. New possible satisfaction situations are additionally constrained in that they have to get incorporated into these more specific sub-concepts *salva stabilitate*, i.e. they are expected to fit into the set of experiences which had been made in the respective circumstances. Therefore the specific neighbourhoods involved in a belief determine the implications it has.

An interesting facet of this whole picture is that the conceptual structure, on the extensional level of growing sets of experiences, contains an unlimited amount of sub-concepts, as partial individual concepts and as partial general concepts. They are themselves not lexically expressed and are denoted only by the names for the maximal individual concepts and maximal general concepts they are part of. The enormously fine-grained conceptual structure is linguistically representable by the labels for the conceptually complete sets, i.e. by expressions of the individual and general concepts that are common to the speech community. The maximal concepts are mapped onto the layer of corresponding linguistic labels for individual and general concepts, and the relationships of inclusion, overlap, and opposition between concepts can be expressed on that layer linguistically. The sub-concepts often have no expression of their own, but have to be expressed by the labels for the maximal concepts into which they can be continued. On this linguistic layer, seen by itself, linguistically unexpressed conceptually fine-grained differences are lost, though they are available in the extensional rooting of the concepts on the set of experiences. This, of course, implies that not two individual conceptual systems can have exactly the same beliefs. The identity of beliefs reaches always as far as the experiential basis is the same, and as far as the relevant theories are the same. For communicational purposes the identity goes as far as the common maximal individual and general concepts go, and as far as certain perspectives and situational conditions are


commonly available, which make context-dependence of lexical understanding also a public, though local, phenomenon.

7. References

- Asher, N. 1986
Belief in Discourse Representation Theory. *Journal of Philosophical Logic*. Vol. 15::127-189
- Bartsch, R. 1993
Concept Formation and Understanding: An Overview. D. Nauta, A. Nijholt and J. Schaake (eds.), *Pragmatics in Language Technology. Proceedings of the fourth Twente Workshop on Language Technology*: 71-86. Department of Computer Science. University of Twente. NL 7500 AE Enschede.
- Bartsch, R. 1995
Situations, Tense, and Aspects. Mouton de Gruyter. Berlin-New York.
- Bartsch, R. 1996
The Relationship between Connectionist Models and a Dynamic Data-oriented Theory of Concept Formation. ILLC Research Report LP-94-19. Institute of Logic, Language and Computation (ILLC), University of Amsterdam. To appear in: *Synthese* 1996.
- Barwise, J. and J. Perry. 1983
Situations and Attitudes. Cambridge, Mass.: MIT Press.
- Carnap, R. 1947/1956.
Meaning and Necessity. Chicago: The University of Chicago Press.
- Cresswell, M.J. 1985
Structured Meanings. The Semantics of Propositional Attitudes. Cambridge Mass.: M.I.T. Press.
- Davidson, D. 1967
The Logical Form of Action Sentences. N. Resher (ed.). *The Logic of Decision and Action*. Pittsburgh: The University of Pittsburgh Press. Reprinted in: D. Davidson, *Essays on Actions and Events*. Oxford: Clarendon Press 1982: 105-122.
- Dekker, P. 1993
Transsentential Meditations. Ups and Downs in Dynamic Semantics. Ph.D.-dissertation. University of Amsterdam.
- Dummett, M. 1976
What is a Theory of Meaning?(2). G. Evans and J. McDowell (eds.), *Truth and Meaning*. Oxford University Press: 67-137.

- Fodor, J. 1976
The Language of Thought. Hassocks: The Harvester Press.
- Fodor, J. 1987
 Psychosemantics. The Problem of Meaning in the Philosophy of Mind.
 Cambridge, Mass.: MIT Press.
- Groenendijk. J. and Stokhof. M. 1991
 Dynamic Predicate Logic. *Linguistics and Philosophy* 4: 39-100.
- Hintikka, J. 1969
 Semantics for Propositional Attitudes. Linsky (ed.), *Reference and Modality*. London: Oxford University Press: 145-167.
- Husserl, E. 1913/1920
Logische Untersuchungen. II. Band. Teil 1 und 2. Halle: Max Niemeyer.
- Husserl, E. 1939/1985
Erfahrung und Urteil.. Hamburg: Felix Meiner Verlag.
- Landman, F. 1986
 Towards a Theory of Information. The Status of Partial Objects in Semantics. Dordrecht: Foris Publications.
- Lewis, D. 1972
 General Semantics. Davidson and G. Harman (eds.), *Semantics of Natural Language*. Dordrecht: Reidel: 169-218.
- Piaget, J. 1947/1967
Psychologie der Intelligenz. Zürich und Stuttgart: Rascher Verlag.
 (French original: *La Psychologie de l'Intelligence*. Paris: Librairie Armand Colin 1947.
- Putnam. H. 1975
 The Meaning of 'Meaning'. In: *Philosophical Papers, Vol. 2*. Cambridge University Press.
- Quine, W.V.O. 1960
Word and Object. Cambridge, Mass.: MIT Press.
- Ramsey, W., St. Stich, and J. Garon 1991
 Connectionism, eliminativism, and the future of folk psychology.
 Greenwood, J.D. (ed.), *The Future of Folk Psychology, Intentionality and Cognitive Science*: 93-119. Cambridge: Cambridge University.
- Searle, J. 1983
Intentionality. Cambridge: Cambridge University Press.
- Searle. J. 1992
The Rediscovery of Mind. Cambridge. Mass.: The MIT-Press.

- Seligman, J. 1990
Perspectives in Situation Theory. R. Cooper, K.Mukai, and J.Perry,
eds., *Situation Theory and its Applications. Vol. 1*: 147-192. CSLI
Lecture Notes 22. Stanford University.
- Stalnaker, R. 1987
Inquiry. Cambridge Mass.: The MIT Press.
- Stechow, von A. 1984
Structured Propositions and Essential Indexicals. Landman, F and F.
Veltman (eds.), *Varieties of Formal Semantics*. GRASS-Series no.3:
385-404. Dordrecht: Foris Publications.
- Veltman, F. 1984.
Data Semantics. In: J. Groenendijk, et al.(eds.), *Formal Methods in the
Study of Language*. Dordrecht: Foris Publications: 43-63.



ILLC Research Reports and Technical Notes

Coding for Reports and Dissertations: *Series-Year-Number*, with LP = Logic, Philosophy and Linguistics; ML = Mathematical Logic and Foundations; CT = Computation and Complexity Theory; X = Technical Notes; DS = Dissertations. All previous ILLC-publications are available from the ILLC bureau. For prepublications before 1994, contact the bureau.

- CT-95-01 Marianne Kalsbeek, Yuejun Jiang *A Vademecum of Ambivalent Logic*
- CT-95-02 Leen Torenvliet, Marten Trautwein *A Note on the Complexity of Restricted Attribute-Value Grammars*
- CT-95-03 Krzysztof Apt, Ingrid Luitjes *Verification of Logic Programs with Delay Declarations*
- CT-95-04 Paul Vitányi *Randomness*
- CT-95-05 Joeri Engelfriet *Minimal Temporal Epistemic Logic*
- CT-95-06 Krzysztof Apt, Rachel Ben-Eliyahu *Meta-variables in Logic Programming, or the Praise of Ambivalent Syntax*
- CT-95-07 Frans Voorbraak *Combining Unreliable Pieces of Evidence*
- CT-96-01 Peter van Emde Boas *The Convenience of Tilings*
- LP-95-01 Marten Trautwein *Assessing Complexity Results in Feature Theories*
- LP-95-02 S.T. Baban, S. Husein *Programmable Grammar of the Kurdish Language*
- LP-95-03 Kazimierz Świrydowicz *There exist exactly two Maximal Strictly Relevant Extensions of the Relevant Logic R**
- LP-95-04 Jaap van der Does, Henk Verkuyl *Quantification and Predication*
- LP-95-05 Nataša Rakić *Past, Present, Future and Special Relativity*
- LP-95-06 David Beaver *An Infinite Number of Monkeys*
- LP-95-07 Paul Dekker *The Values of Variables in Dynamic Semantics*
- LP-95-08 Jaap van der Does, Jan van Eijck *Basic Quantifier Theory*
- LP-95-09 Jeroen Groenendijk, Marin Stokhof, Frank Veltman *Coreference and Modality*
- LP-95-10 Jeroen Groenendijk, Martin Stokhof, Frank Veltman *Coreference and Contextually Restricted Quantification*
- LP-96-01 Renate Bartsch *Understanding Understanding*
- LP-96-02 David Beaver *Presupposition*
- LP-96-03 Theo M.V. Janssen *Compositionality*
- LP-96-04 Reinhard Muskens, Johan van Benthem, Albert Visser *Dynamics*
- LP-96-05 Dick de Jongh, Makoto Kanazawa *Angluin's Theorem for Indexed Families of R.E. Sets and Applications*
- LP-96-06 François Lepage, Serge Lapierre *The Functional Completeness of 4-value Monotonic Protothetics*
- LP-96-07 Frans Voorbraak *Probabilistic Belief Expansion and Conditioning*
- LP-96-08 John Case *The Power of Vacillation in Language Learning*
- LP-96-09 Jaap van der Does, Willem Groeneveld and Frank Veltman *An Update on Might*
- LP-96-10 Jelle Gerbrandy and Willem Groeneveld *Reasoning about Information Change*
- LP-96-11 Renate Bartsch *Propositional Attitudes in Dynamic Conceptual Semantics*
- ML-95-01 Michiel van Lambalgen *Randomness and Infinity*
- ML-95-02 Johan van Benthem, Giovanna D'Agostino, Angelo Montanari, Alberto Policriti *Modal Deduction in Second-Order Logic and Set Theory*
- ML-95-03 Vladimir Kanovei, Michiel van Lambalgen *On a Spector Ultrapower of the Solovay Model*
- ML-95-04 Hajnal Andréka, Johan van Benthem, István Németi *Back and Forth between Modal Logic and Classical Logic*
- ML-95-05 Natasha Alechina, Michiel van Lambalgen *Generalized Quantification as Substructural Logic*
- ML-95-06 Dick de Jongh, Albert Visser *Embeddings of Heyting Algebras (revised version of ML-93-14)*
- ML-95-07 Johan van Benthem *Modal Foundations of Predicate Logic*
- ML-95-08 Eric Rosen *Modal Logic over Finite Structures*
- ML-95-09 Hiroakira Ono *Decidability and Finite Model Property of Substructural Logics*

ML-95-10 Alexei P. Kopylov *The Undecidability of Second Order Linear Affine Logic*

ML-96-01 Domenico Zambella *Algebraic Methods and Bounded Formulas*

ML-96-02 Domenico Zambella *On Forcing in Bounded Arithmetic*

ML-96-03 Hajnal Andréka, Johan van Benthem & István Németi *Modal Languages and Bounded Fragments of Predicate Logic*

ML-96-04 Kees Doets *Proper Classes*

ML-96-05 Søren Riis *Count(q) versus the Pigeon-Hole Principle*

ML-96-06 Angelo Montanari, Alberto Policriti *A Decidable Theory of Finitely-Layered Metric Temporal Structures*

ML-96-07 Angelo Montanari, Adriano Peron, Alberto Policriti *Decidable Theories of ω -Layered Metric Temporal Structures*

ML-96-08 Johan van Benthem, Angelo Montanari, Giovanna D'Agostino, Alberto Policriti *Modal Deduction in Second-Order Logic and Set Theory - II*

ML-96-09 Angelo Montanari, Maarten de Rijke *Decidability in Metric Temporal Logic*

ML-96-10 Vladimir Kanovei *On a Dichotomy related to Colourings of Definable Graphs in Generic Models*

ML-96-11 Domenico Zambella *Forcing in Finite Structures (revised version of ML-96-02)*

ML-96-12 Jon Barwise and Johan van Benthem *Interpolation, Preservation, and Pebble Games*

X-95-01 Sophie Fischer, Leen Torenvliet *The Malleability of TSP_{2Opt}*

X-96-01 Ingmar Visser *Mind Rules: a Philosophical Essay on Psychological Rules and the Rules of Psychology*

DS-95-01 Jacob Brunekreef *On Modular Algebraic Protocol Specification*

DS-95-02 Andreja Prijatelj *Investigating Bounded Contraction*

DS-95-03 Maarten Marx *Algebraic Relativization and Arrow Logic*

DS-95-04 Dejuan Wang *Study on the Formal Semantics of Pictures*

DS-95-05 Frank Tip *Generation of Program Analysis Tools*

DS-95-06 Jos van Wamel *Verification Techniques for Elementary Data Types and Retransmission Protocols*

DS-95-07 Sandro Etalle *Transformation and Analysis of (Constraint) Logic Programs*

DS-95-08 Natasha Kurtonina *Frames and Labels. A Modal Analysis of Categorical Inference*

DS-95-09 G.J. Veltink *Tools for PSF*

DS-95-10 Giovanna Cepparello *Studies in Dynamic Logic*

DS-95-11 W.P.M. Meyer Viol *Instantial Logic. An Investigation into Reasoning with Instances*

DS-95-12 Szabolcs Mikulás *Taming Logics*

DS-95-13 Marianne Kalsbeek *Meta-Logics for Logic Programming*

DS-95-14 Rens Bod *Enriching Linguistics with Statistics: Performance Models of Natural Language*

DS-95-15 Marten Trautwein *Computational Pitfalls in Tractable Grammatical Formalisms*

DS-95-16 Sophie Fischer *The Solution Sets of Local Search Problems*

DS-95-17 Michiel Leezenberg *Contexts of Metaphor*

DS-95-18 Willem Groeneveld *Logical Investigations into Dynamic Semantics*

DS-95-19 Erik Aarts *Investigations in Logic, Language and Computation*

DS-95-20 Natasha Alechina *Modal Quantifiers*

DS-96-01 Lex Hendriks *Computations in Propositional Logic*

DS-96-02 Angelo Montanari *Metric and Layered Temporal Logic for Time Granularity*

DS-96-03 Martin H. van den Berg *Some Aspects of the Internal Structure of Discourse: the Dynamics of Nominal Anaphora*