

Propositional Attitudes and their representation

February 25, 1998

1 Introduction

In this chapter, I will present the main two theories of belief representation and show that they are jeopardised, for opposite but complementary reasons, by some famous puzzles developed in the philosophy of language. I will defend the claim that these problems are due to the ambiguous nature of intentionality, and that only a theory able to represent *both* the internal and the external nature of intentional content is adequate for representing the attitudes.

This claim has been defended in the framework of Discourse Representation Theory, by N. Asher([Ash86], [Ash87], [Ash89]) and H. Kamp([Kam85], [Kam90], [Kam97]). I will review and critically assess their work, and stipulate what is required of an adequate account to the representation of reasoning with attitudes.

2 Criteria for a logic of the attitudes

Propositional attitudes draw their very name from the fact that they are individuated by their contents: to believe or to wish something is to take some kind of attitude towards a proposition. But what are propositions? Instead of trying to answer this question directly, which would lead us into deep metaphysical problems, it seems reasonable to look for what *properties* propositions have, and to select the candidate entities according to these properties.

2.1 Logical properties

Firstly, propositions exhibit logical properties. When one combines them with the usual logical operators – conjunction, disjunction and negation –, they entertain relations of entailment. For instance, the proposition created by the conjunction of a proposition p and a proposition q entails p . It is crucial to note that the propositional attitudes inherit entailment relations from the logical structure of their contents. For instance, the following axioms should hold for any logic of beliefs:¹

1. Simplification.
 x believes $(p \wedge q)$ entails x believes p
2. Commutativity.
 x believes $(p \wedge q)$ entails x believes $(q \wedge p)$
3. Associativity.
 x believes $p \wedge (q \wedge r)$ entails x believes $(p \wedge q) \wedge r$
4. Existential Generalisation.
 x believes $\phi(a)$ entails x believes $\exists y\phi(y)$

¹From [AS93].

These entailments are due to the corresponding entailment relations holding between the propositions believed. For this reason, propositions themselves should possess some logical structure. We may interpret this in two ways: a syntactic way, which attributes a linguistic form to the contents of the attitudes; and a semantic way, which characterises the contents of the attitudes in terms of possible worlds.

2.1.1 Sentential theories

The simplest way to represent, say, the beliefs of an agent, is to identify them with a set of sentences. Intuitively, a sentence belongs to the belief set of an agent only if he would *accept* it. The sentences can be couched in any well defined language – the important thing being that the usual logical relations between sentences can be defined in this language.

To describe the cognitive states of an agent x , we shall therefore need an internal language L_0 – a first-order language, for instance –, and an extension of it, $L \supseteq L_0$, including a belief operator bel_x . Note that the internal language L_0 may include a belief operator as well. This is needed if we want to model cases where agents have beliefs about beliefs. Here are the central definitions:

- $bel_x\phi$ iff $\phi \in \mathbf{bel}(x)$
- $\neg bel_x\phi$ iff $\phi \notin \mathbf{bel}(x)$

where $bel_x\phi$ is a formula of L , and where $\mathbf{bel}(x)$ denotes the belief set of the agent x . Drawing on [Kon86], I will define an interpretation for the language L on the following lines:

Definition 1 A $B(L, \delta)$ model is a tuple $\langle D, F, \delta \rangle$, where the first two elements $\langle D, F \rangle$ are a first-order model structure, and where $\delta = \{d_0, d_1, \dots\}$ is a sequence of deduction structures, one for each agent named in L .

This definition presupposes the notion of deduction structure:

Definition 2 A deduction structure is a pair consisting of a set of base formulae and a set of deduction rules. The language of a deduction structure is L_0 .

For instance, consider the sentence:

$$(1) \quad bel_x(\phi(a)) \wedge \neg\phi(a)$$

It says that the agent x has the false belief expressed by the sentence $\phi(a)$. This sentence is satisfied in a model in which $\phi(a)$ is false, and in which the deduction structure with base set $\forall x\phi(x)$ and with deduction rule $\forall x\phi(x) \vdash [a/x]\phi$ is attributed to the agent x .

A very interesting feature of this kind of theory is that some properties of belief systems are reflected by constraints imposed on $B(L, \delta)$ models².

For instance, the minimal properties quoted above are obtained in the models in which the corresponding formulae are valid. Less trivial properties of beliefs find a formal counterpart as well. Consider introspection, for instance. According to this property of belief systems, an agent cannot both believe something, and fail to believe that he believes it. This property translates easily in terms of belief sets and their interpretations:

- Positive introspection: if ϕ is an element in x 's belief set, so is $bel_x(\phi)$. Characterising schema: $bel_x(\phi) \rightarrow bel_x(bel_x(\phi))$
- Negative introspection: if ϕ is not in x 's belief system, then $\neg bel_x(\phi)$ is. Characterising schema: $\neg bel_x(\phi) \rightarrow bel_x(\neg bel_x(\phi))$

Other interesting properties can be so characterised:

- consistency: $(bel_x(\phi)) \rightarrow \neg bel_x\neg\phi$

²See [Kon86], pp. 35–38.

- knowledge: $bel_x(\phi) \rightarrow \phi$

Note that some of these conditions are really strong, and certainly not realistic from a psychological point of view. Positive introspection, for instance, entails that each time an agent a has a sentence ϕ in his belief box, he also has $bel_a(\phi)$, $bel_a(bel_a(\phi))$, ... and therefore must end up with an infinite number of sentences in his belief box³. Negative introspection seems too strong as well. However, the sentential theory is not committed, in itself, to any of these axioms.

2.1.2 Possible worlds theories

According to the second approach, a proposition is a partition of a set of possible worlds, or to put it in a more familiar way, a *function from possible worlds into truth-values*. The epistemic interpretation of this construction is straightforward: a proposition is seen as the narrowest set of possible worlds in which an agent is certain to find the actual world - or, equivalently, as the largest set of possible worlds compatible with his convictions.

In this model, a proposition p is accepted by an agent's total epistemic context c iff c is a subset of p ; it is rejected otherwise. An epistemic context, relative to a particular possible world and individual, is given by an accessibility relation: the worlds entertained as possible alternatives by the agent are said to be accessible from the actual world.³ A formalisation of this idea is given in the following definitions:

Definition 3 A frame is a tuple $\langle W, A \rangle$, where W is a set of possible worlds, and A is a set of binary accessibility relations on W , one for each agent.

Kripke models, or K-models, are defined in terms of frames.

Definition 4 A K-model based on a frame $\langle W, A \rangle$ is an ordered triple $\langle W, A, V \rangle$, where V is a value-assignment function.

Definition 5 If R_x is an accessibility relation for the agent x , $bel_x(\phi)$ is true in a world w relative to a model M iff for all w' such that wR_xw' , ϕ is true in w' relative to M .

According to these definitions, the propositions stated above as minimal conditions a theory of belief should satisfy are all true. As an example, consider Simplification:

- (2) x believes $(p \wedge q)$ entails x believes p

If $bel_x(p \wedge q)$ is true in a world w , then $p \wedge q$ is true in every world accessible from w . Hence, p is true in every world accessible from w , and therefore $bel_x(p \wedge q)$ is true.

Moreover, different conditions on the accessibility relation of an agent x correspond to different axiomatisations of our logic of belief, and to different properties of belief systems. Table 2.1 gives a summary of the belief systems obtained, with the corresponding accessibility relations and axiom schema.

Name	Condition on R_x	axiom schema
T	reflexive	$Bel_x(\phi) \rightarrow \phi$
D	serial	$Bel_x(\phi) \rightarrow \neg Bel_x(\neg\phi)$
S_4	transitive	$Bel_x(\phi) \rightarrow Bel_x(Bel_x(\phi))$
S_5	euchdian	$\neg Bel_x(\phi) \rightarrow Bel_x(\neg Bel_x(\phi))$

Table 1: Models for belief systems

The possible worlds semantics for belief is very elegant, but also extremely strong. Firstly, agents

³See [Hin62] and [Hin69] for early formulations of these ideas.

can have inconsistent beliefs in this framework, only if they believe every propositions. The possible worlds approach cannot distinguish between believing an explicit contradiction (*i.e.*, believing both ϕ and $\neg\phi$) and being logically inconsistent (*i.e.*, believing formulae that entail a contradiction, without having derived it explicitly). For suppose an agent believes two propositions that entail a contradiction. In this situation, the propositions can never be true in the same possible world, and the agent ends up believing everything, since the worlds he has access to are reduced to the empty set. Secondly, possible worlds semantics cannot distinguish between beliefs in logically equivalent propositions, since these propositions are true in the same worlds, and since beliefs are individuated by the set of possible worlds their objects are true at. Thirdly, agents believe every logically valid proposition. This very unfortunate consequence of the theory is due to the fact that a logically valid proposition is true in every possible world, so in particular in all the possible worlds accessible to an agent. Finally, the possible worlds approach is bound to see an agent as an ideal logical reasoner. Suppose that an agent believes p , and that q is a logical consequence of p . In each possible world where p is true, q will be true as well. But then, the agent believes q . This means that the agent believes all the logical consequences of her beliefs: her beliefs are closed under the relation of logical entailment.

2.2 Attitudes and behaviour

A second characteristic of the attitudes is that they cause behaviour. In fact, some authors insist that this is their most central property.⁴ At first blush, possible worlds theories seem to be better suited to account for this pragmatic aspect of the attitudes than sentential theories. Typically, actions are individuated by the facts they help to bring about, which can easily be represented in terms of possible worlds. The possible worlds machinery, however, does not turn out to be wholly adequate to capture the way attitudes interact with an agent's behaviour. The reason for this is that propositions construed as set of possible worlds are not fine-grained enough to account for the causal role of epistemic states. Consider the following story, due to J. Perry:

I once followed a trail of sugar on a supermarket floor, pushing my cart down the aisle on one side of a tall counter and back the aisle on the other, seeking the shopper with the torn sack to tell him he was making a mess. With each trip around the counter, the trail became thicker. But I seemed unable to catch up. Finally it dawned on me. I was the shopper I was trying to catch. I believed at the outset that the shopper with a torn sack was making a mess. And I was right. But I did not believe that I was making a mess. That seems to be something I came to believe. And when I came to believe that, I stopped following the trail around the counter and rearranged the torn sack in my cart. My change in beliefs seems to explain my change in behaviour. [Per93].

Now, how are we going to construct the propositions believed by Perry, in the possible worlds framework? There seems to be two possibilities: either we construct them descriptively, as incorporating descriptions of the objects of his thoughts, or directly, as incorporating the objects themselves. The first option does not lead to a paradox: Perry is first related to the proposition that *the man who is so and so* is making a mess, and then to the proposition that δ_{Perry} is making a mess, where δ_{Perry} express the way Perry would describe himself. However, this analysis does not capture our intuitions about the situation. Perry is causally related to the man that happens to be himself: he can see him, he may try to talk to him, to move towards him. And he does not need to know any description of the man in order for the passage above to be true. His belief is not mediated by a description of the man, but by his sensuo-motor abilities, by his capacity to maintain a cognitive grip on the man.

Consider now the second option. It respects our intuitions, at least to the extent that it considers Perry's first belief as expressing a proposition which includes the particular man he is thinking about, *i.e.*, himself. But it cannot capture the *change* in Perry's beliefs anymore: the epistemic

⁴See for instance [Str84], Chap. 1.

context, individuated in terms of possible worlds, does not change, since the messy shopper happens to be Perry himself. Hence, the belief that the messy shopper was making a mess, and the belief that Perry himself is making a mess, turn out to be the same according to the possible worlds approach.

Perry's story shows that the possible worlds theory cannot both account for the fact that our thoughts may be directly directed towards their objects without a descriptive mediation, and for belief change. The force of the argument obviously depends on whether or not one thinks that direct beliefs are indispensable. In the limits of this dissertation, I will assume that they are.⁵ Sentential theories also face problems as far as the pragmatic aspects of the attitudes are concerned. Consider the following twins story:

In one room of a house, Pierre is looking at an apple, and forms the belief:

- (3) This apple is red.

At the same time in another room, his twin, who happens to be exactly identical to him, is looking at an apple exactly similar to the first one; he forms the belief:

- (4) This apple is red.

(this should not be a surprise, since Pierre and his twin have the same internal mental states)

It seems clear that Pierre and his twin must have the same sentences in their belief box - and so, according to the sentential theory, they should believe the same proposition. The problem is that they simply don't, since Pierre's belief is directed toward a particular apple, and his twin's thought toward another, different, one.

2.3 The puzzles of intentionality

Attitudes are intentional: they are directed towards their objects. Here, however, looms an ambiguity: what objects are we talking about? Is a belief, for instance, directed towards a real individual in the world? Or is it directed towards a 'mental object' that the agent has in mind? In my opinion, the puzzles I introduced in the previous paragraph show that neither of these answers is enough in itself. Following Fodor ([Fod94]), I think that it is possible to classify them, as well as most of the philosophical puzzles developed in the literature, in two categories: Frege's cases and Twins cases.

In Frege's cases, two different beliefs of an agent are directed, unbeknown to him, towards the same external object. Perry's story of course describes such a situation. The conclusion of these cases, as we saw about Perry's argument, is conditional: if one accepts some widely held assumptions about the semantics of singular beliefs, propositions are not fine-grained enough to individuate beliefs. To illustrate this point, let us consider Kripke's puzzle about Pierre, which is another example of a Frege's case.⁶

Pierre is a French speaker who lives in France and does not speak English. He has heard about London, and has admired pictures of it in magazines. Relying on these experiences, he would utter, in French, 'Londres est jolie' ('London is pretty'). On this basis, we could safely attribute to Pierre the belief that London is pretty. Later, Pierre moves to London, in a very poor part of the city, in which he stays without ever going to the centre. Because the population always complain about the city, and because of his own unfortunate experience, he does not connect London with Londres. Therefore, he would still utter 'Londres est jolie', but would claim as well that 'London n'est pas

jolie'. Kripke's conclusion is that one can ascribe to Pierre both the belief that London is pretty *and* the belief that London is not pretty, without seeing him as inconsistent or irrational.

Here, the puzzle arises when one considers that the occurrences of 'London' and 'Londres' in Pierre's utterances should contribute by an individual rather than by a description to the propositions expressed. If one also assumes that Pierre believes what he says, the puzzle immediately follows. A very important conclusion one can draw from Kripke's puzzle is that linguistic entities like natural language sentences are not fine-grained enough either to individuate beliefs. The same sentence, in different context, may be used by speakers to attribute different beliefs.

Twins cases illustrate a symmetric situation, where a unique internal state may be anchored to different objects in the world, and therefore may have different contents in different circumstances.

One could argue that perceptual states have this property: the content of a perceptual belief - for instance, the belief stated in (5) - does not only depend on descriptive information, but on some causal links which anchor it onto a real object.

- (5) There is an apple in front of me.

Therefore, two individuals thinking the thought expressed by (5) in front of *different* apples form different beliefs. To illustrate how a twins case could arise in real life, let me tell you the following story, whose broad lines are due to D. Lodge:

Percy, a young academics attending a congress on English literature, meets Angelica, a bright, pretty, and very serious-looking PhD student. He falls in love with her but, unfortunately, she disappears. Sad and depressed, he enters in a cinema which is showing an erotic film. To his surprise, he recognises ? in the main actress. Stunned, he cannot help saying aloud:

- (6) This girl is a sinner!

In Lodge's convoluted story, the actress happens to be Angelica's twin sister. Therefore, the adequacy of Percy's beliefs depends on the situation. When he thinks the belief expressed by (6), looking towards the image of the actress, he is right (if we endorse his rather strict moral values for the sake of the argument). But suppose that he met Angelica again, and that the thought that she is a sinner occurred in his mind. This time, he would be wrong.

The origin of Percy's confusion seems to be that a mental representation, containing some descriptive information, gets anchored in context to different individuals.

2.4 The criteria

After these discussions, the reason why I think that neither the sentential theory nor the possible worlds theory are satisfying should be clear. For opposite reasons, each is jeopardised by one of our puzzles of intentionality. In a nutshell, the sentential theory is unable to account for the fact that our attitudes have an external aspect, whereas the possible worlds theory cannot represent differences due to their internal properties. Therefore, we need a theory which would be able to:

- model the logical properties propositions have
- model the internal aspect of the attitudes
- represent the anchoring of the attitudes in the real world

I think that these criteria are met by dynamic theories of the attitudes.

⁵For a review of most of the arguments motivating this assumption, see [Bac94] and [Rec93].

⁶See [Kri79] for the original formulation of the puzzle, and see [Ash86] for an extensive DR-theoretic analysis of its consequences.

3 Attitudes modelled in a dynamic framework

The dynamic conception of content was originally motivated by puzzles concerning the meaning of anaphora. Consider this famous example, due to Barbara Partee:⁷

- (7) a. Four marbles out of five are on the table.
 b. * It is under the sofa.
- (8) a. Every marble is on the table, except the red one.
 b. It is under the sofa.

The problem with this minimal pair is that (7a) and (8a) have exactly the same truth-conditions. Why do our judgements concerning their possible continuations differ in the way they do?

The answer of dynamic semantics is that sentences convey information bearing on the world, but also discourse information, which constrains the way further sentences are to be interpreted.

Of course, we should not expect to find anything exactly similar to anaphora in the domain of the attitudes. It has been argued, however, that the dynamic notion of content used in the context of discourse semantics could be fruitfully applied to this domain.

3.1 Discourse referents and mental individuals

The *prima facie* reason for this claim is that discourse referents, such as they are used in dynamic semantics, seem to be well suited to represent the internal side of content. One can see a belief system, internally, as a system of files⁸. When one forms a belief about an individual, one opens a file where information can be stored, and retrieved on suitable occasions. There does not need to be a bijection between such files and the real objects they are about. Everyone familiar with administrative problems knows that many unrelated files may be created about the same person. Something similar seems to happen in our cognitive system: we can mentally represent the same object in very different ways, especially when we have formed beliefs about it on very different occasions.

On the other hand, the same file may be re-used many times, for different objects, just as different mental representations may be linked with different things in the world. Discourse individuals are very similar to files as well⁹. One may consider the universe of a discourse representation as specifying the entries of files in which the information is incrementally stored. Moreover, a discourse referent may be anchored onto a real individual, just as a file may be attached to a real person.

3.2 The holistic nature of epistemic states

Let us draw on the file metaphor. To be efficient, files must be inter-connected: from a file, one should be able, when needed, to access the information stored in other files. One of the main claim of DR-theoretic approaches of attitudes representation is that mental states are internally connected. Just as the dynamic approaches to interpretation claim that utterances of particular sentences have a truth-value only relative to the context of previously interpreted sentences, Kamp claims that attitudes have truth-conditional content only relative to a whole epistemic background:

In general it is not possible to specify the truth-conditional content of each of the individual beliefs that are held by a given person; only the belief state as a whole (or, at any rate, substantial parts of it, which comprehend a large number of his beliefs) determines well-defined truth-conditions ([Kam90], p. 40)

⁷This example is quoted in [Hei82], Chapter one.

⁸This metaphor has been extensively used in the philosophy of direct thought. See for instance [Bac94, [Rec93].

⁹They resemble files so much that Irene Heim dubbed her theory ‘File Change Semantics’. See [Hei82].

Of course, *some* of the individual attitudes of an agent have truth-conditions by themselves, as it were; that is, independently of the whole epistemic context. We could imagine, for instance, a punctuate mind believing that all men are mortal, and nothing else. Such beliefs, it seems, do not presuppose other attitudes.

Nevertheless, Kamp is right to maintain that in general, attitudes cannot live an isolated life: we shall see that singular beliefs, desires, and intentions all need a presupposed background. Moreover, taking this background information into account will enable us to solve the puzzles of intentionality.

4 Models of beliefs

In the first days of DRT, belief was modelled in a structural way, as a relation between an individual and a syntactic object: a DRS. Even though DRT seems a natural framework to implement such an idea, a more traditional conception of belief as a relation between an individual and a propositional (semantic) object is also possible, and is in fact proposed by Geurts [Geu95]. I will compare these two approaches, starting with the semantic approach.

4.1 A DRT implementation of the possible worlds theory

According to possible worlds theories, a belief is accurately attributed to an individual if it is satisfied by his total epistemic state, represented by the set of worlds compatible with the agent’s opinions. The first thing to do, thus, is to upgrade the DRT syntax so that total cognitive states are represented. Geurts’ new definition of DRS-Terms is the following:


Definition 6 *DRS-terms*:

- *The set of DRS-terms* $T = T_1 \cup T_2 \cup T_3$
- $T_1 = \{u, v, w, \dots\}$, *the set of discourse referents*
- $T_2 = \{p, q, r, \dots\}$, *the set of propositional reference markers*
- $T_3 = \{p + K \mid p \in T_2 \text{ and } K \text{ is a DRS}\}$, *the set of complex propositional terms*

A belief, in this framework, is represented by a complex propositional term: its content adds some new information to the previous cognitive context of the agent, represented by a propositional reference marker. Thus, the sentence

- (9) Pierre believes that a man whistles

will be represented by:



What kind of semantics should propositional terms have? To answer, we should first define a semantics for DRSs. This is rather straightforward: the semantic value of a DRS should be an indexed proposition (Geurts), or, maybe more accurately, an open proposition (Kamp), that is, a set of worlds/embedding functions pairs¹⁰:

¹⁰Following Geurts, I will sometimes call *states* these worlds/embedding functions pairs.

Definition 7 Let $\langle W, D, I \rangle$ be an intensional model, where W is a set of possible worlds, D is a domain of individuals, and I is an interpretation function. Let f be an embedding function and w a world. Then:

$$[[K]]_{w,f} = \{\langle w, g \rangle \mid f \subseteq g \wedge \text{dom}(g) = U(K) \wedge \forall \phi \in \text{Con}(K), [[\phi]]_{w,g} = 1\}$$

Open proposition encode both information about the world, and discourse information. To illustrate this point, let us consider a simple example.

$M = \langle W, D, I \rangle$, with $W = \{w_1, w_2\}$, $D = \{\text{Pierre}, \text{Kim}\}$, and I the function such that

- $I(w_1, \text{'walk'}) = \{\text{Pierre}\}$, $I(w_2, \text{'walk'}) = \{\text{Pierre}\}$
- $I(w_1, \text{'whistle'}) = \{\text{Pierre}\}$, $I(w_2, \text{'whistle'}) = \{\text{Kim}\}$
- $I(w_1, \text{'man'}) = \{\text{Pierre}, \text{Kim}\}$, $I(w_2, \text{'man'}) = \{\text{Pierre}, \text{Kim}\}$

Consider now the sentence:

- (10) A man walks. He whistles.

and the discourse representations corresponding to each sentence:

$$(11) \quad \begin{array}{ll} \text{a.} & \boxed{\begin{array}{c} x \\ \text{Man}(x) \\ \text{walk}(x) \end{array}} \\ \text{b.} & \boxed{\begin{array}{c} y \\ x = y \\ \text{whistle}(x) \end{array}} \end{array}$$

According to the definition above, the denotation of the first DRS is the open proposition $\{\langle f, w_1 \rangle \mid \langle f, w_2 \rangle\}$, where f is the function sending x onto Pierre. Since a man walks in each world of the model, the sentence does not convey any information about the world. However, some discourse information is conveyed by the embedding function f . The presence of the function constrains possible continuations of the discourse. Consider now the denotation of the final DRS, after the interpretation of the second sentence. If we interpret the pronoun anaphorically, each pair in the new open proposition must include a function g extending f , such that $g(y) = f(x) = \text{Pierre}$. This constraint leads to the elimination of one possibility, since Pierre does not whistle in w_2 . The open proposition will therefore be $\{\langle g, w_1 \rangle\}$. Note that in this situation, the discourse information that *the individual talked about whistles* indirectly conveys world information, since a possible world is eliminated in the open proposition. This is a case of discourse information leading to world information.

From this, we can draw a semantic definition for complex terms:

Definition 8 Let f be an embedding function and w a world. Then:

$$[[p + K]]_{w,f} = \{\langle w', h \rangle \mid \exists \langle w', g \rangle \in [[p]]_{w,f} : \langle w', h \rangle \in [[K]]_{w',g}\}$$

Intuitively, $[[p + K]]_{w,f}$ is the open proposition obtained by taking the extensions of the embedding functions in p , and the set of worlds in which these extended embedding functions satisfy the DRS K .

The next step is to find a way of extracting the 'purely' propositional information from an open proposition:

Definition 9 Let s be an open proposition (that is, a set of embedding functions/worlds pairs).

The proposition defined by s , $\text{Prop}(s)$, is:

$$\text{Prop}(s) = \{w \mid \exists f : \langle w, f \rangle \in s\}$$

It is then possible to give a Hintikka style semantics for the *bel* relation:

Definition 10 Let I be the interpretation function giving the semantic value of each predicate in each world, s an open proposition, and dors_x a function giving the set of worlds considered by an agent x as drastic alternatives to the real world. We have:

$$I(w, \text{bel}) = \{\langle a, s \rangle \mid \text{dors}_a(w) \subseteq \text{Prop}(s)\}$$

This purely semantic approach of the attitudes, however, does not really improve on the possible worlds theory as far as the puzzles of intentionality are concerned. The reason for this is that only complete belief contents are represented. The theory does not include a way of encoding the familiarity agents have with the particular objects of their beliefs. The main advantage of the theories I am going to present now is that they deal with the aspect of the attitudes.

4.2 Asher's structural account

In a number of papers ([Ash86][Ash87], [Ash89], [AS93]), N. Asher defends a representational framework in which the objects of beliefs are modelled by DRS-like structures, rather than by propositions. The motivation for this choice seems to be what I called above the ambiguity of intentionality, *i.e.*, the fact that belief contents have both an internal and an external aspect. Consider again Pierre's belief, reported in 12, and suppose that we represent it content by a DRS K :


- (12) a. Pierre believes that London is pretty.

$$\text{b. } K = \boxed{\begin{array}{c} x \\ \text{London}(x) \\ \text{pretty}(x) \end{array}}$$

Firstly, and because of the causal links which hold between Pierre's mind and the world, this belief is directed towards a real geographical individual, London. This connection has nothing to do with the fact that London is *named* 'London'; we could very well imagine a situation in which London would be named 'Paris' but where, nevertheless, Pierre would still believe 12, because he would still believe of the individual London that it is pretty. To account for this 'external' aspect of beliefs, Asher introduced what he called 'external anchors' in DRT. An external anchor is a function from a discourse referent to an individual, which constrains any embedding function to take the individual as value for this discourse referent. Thus, the discourse referent x , in the following representation, is anchored to the individual London:

$$(13) \quad \langle \langle x, \text{London} \rangle, \boxed{\begin{array}{c} x \\ \text{London}(x) \\ \text{pretty}(x) \end{array}} \rangle$$


Secondly, however, Asher insists that our cognitive access to real individuals is always conceptually mediated. When we entertain a singular belief about a real object, we do so by representing the individual as possessing a set of properties. Even more, we can entertain different beliefs about the same individual, when we do not notice that very different sets of conceptual familiarity conditions happen to point towards the same external object. In Kripke's story, for instance, Pierre has two beliefs directed, without his notice, towards the same external individual because he has, as it were, two mental objects representing London. Asher introduced internal anchors to represent the internal aspect of intentionality. Pierre's predicament is rendered as follows:



In this representation, h is an external and/or sending both x and y onto London, and the V_i s and V'_i s are two sets of familiarity conditions. The fact that x and y are not *internally* connected explains why Pierre, although he is not irrational, may entertain two beliefs which cannot be true at the same time. According to Asher, the internal conceptual individual can be shared by several attitudes. For this reason, he extends the DRT language so that it includes what he calls delineated DRSs, which comprise:

- a set of DRSs $\{K_1, \dots, K_n\}$
- a set of discourse referents $\{x_1, \dots, x_m\}$ representing sharable conceptual individuals
- a (possibly empty) set of conditions $\{C_1, \dots, C_k\}$

A delineated DRS represents a relevant portion of an agent's total cognitive state. Thus, the following delineated DRS represents a portion of Pierre's cognitive state:



How should the attribution of a belief to an agent be represented according to Asher? Like Geurts, Asher expands the DRS language to include propositional markers, $p, q, r, \dots$ and DRS reference markers $K_1, K_2, K_3, \dots$. A belief is interpreted as a relation between an individual and a propositional marker, $believe(x, p)$, and the content of a belief is represented by a condition $p : K$. Contrary to Geurts, however, Asher takes propositional markers to denote delineated DRSs rather than open propositions. Thus, a condition $p : K$ is satisfied if the DRS K matches the delineated

DRS denoted by p .

I won't examine here the detailed truth-conditions for belief attributions in Asher's theory. I only want to emphasise the mixed nature of this theory. It is partly a syntactic theory of beliefs, since delineated DRSs, which are syntactic objects, serve as the denotations of belief contents. These syntactic objects, however, also have a denotation in the model. In particular, it makes sense, in Asher's framework, to compare the content of the denotation of a propositional marker with the content of a DRS.

The obvious advantage of Asher's approach is that it combines the good points of the sentential and of the possible worlds theories of beliefs. Because objects of beliefs are modelled by DRS-like representations, they are more fine-grained than propositions defined solely in terms of possible worlds, which leads to obvious solutions for Frege style puzzles. And because these DRSs may be externally anchored, and hence have external contents, Twins style puzzles can be solved as well.

4.3 Kamp's intermediary position

However much seduction it displays, Asher's theory seems to have two main drawbacks. The first one comes from the fact that the theory was first conceived to deal with belief reports rather than the representation of attitudes in general. As a consequence, a delineated DRS does not include the types of the attitudes it represents. This is a major problem, because the potential to treat the relationships between different kinds of thoughts is one of the main advantage of DRS-like representations. We would like to be able to represent, for instance,

(16) Pierre's belief that a man whistles in the park.

(17) Pierre's desire that this man stops.

in a way that would exhibit the fact that the desire presupposes the belief as a kind of basis. The second problem is more philosophical: Asher seems to be committed to the existence of discourse structures, since he introduces them among the possible values of linguistic signs. This may be unfortunate, because we don't seem to be referring to any object of this kind when we talk about other people's attitudes. Of course, one could reply that DRSs, according to Asher's views, should be seen as idealisations of mental structures, and that consequently, they should be allowed in a model. I think that we should at least attempt to avoid this step. DRSs are a theoretician's invention. As such, they should be confined in our representational language. Kamp's recently proposed theory of the attitudes [Kam97] eschews these difficulties. Like Asher, Kamp's concern is to give a holistic representation of the total cognitive state of an agent. For this reason, he includes in his definition of intensional models an interpretation function which assigns to each individual x in each world w a cognitive state $AS(x, w)$:

Definition 11 An intensional model for a DRS language L is a quintuple $\langle W_M, U_M, Name_M, Pred_M, AS_M \rangle$ such that:


- W_M is a non empty set of worlds
- U_M is a function giving the universe of each world, and which singles out a set of agents in each world
- $Name_M$ is a function from the set of names in L into U_M
- $Pred_M$ is a function giving the extension of each predicate in each world
- AS_M is a function assigning a total cognitive state to each agent in each world

However, contrary to Asher, Kamp does not assign DRS-like representations to the agents in order to represent their attitudes. Like Geurts, he chooses to individuate the attitudes by their contents, not by their syntactic forms. Furthermore, Kamp considerably improves on Geurts' theory by taking the interrelationships of the attitudes into account.

4.3.1 Interrelationships between the attitudes in Kamp's theory

The paradigmatic situation he describes should be quoted extensively:

Suppose you see - or seem to see - a gold coin lying on the pavement; you form the belief that there is something x lying there on the pavement in front of you and that this x is a gold coin; being like most people, you form, more or less instantly, the desire to be in possession of x ; and so you conceive of the action of going into the middle of the road and picking x up, thereby realizing the desired state of the coin being yours. Here we have in short succession, a belief, a desire, and an intention to act, all "focused" on the object (represented as) x ; that is, their respective contents are referentially connected through sharing x :



The following definition formalises this intuitive idea of a bunch of inter-connected attitudes:


Definition 12 An articulated DRS is a set of pairs $\langle MOD, K \rangle$ where:

- MOD is one of the mode indicators BEL, DES, INT
- K is a DRS

As it stands, this definition cannot capture the idea of a direct thought, of a thought formed because of a causal relation to a certain particular object. It is straightforward, however, to incorporate Asher's notions of internal and external anchors into the theory. Kamp chooses to do so by defining a new kind of attitude, which he calls anchor ($ANCH$).

Since they contribute to the description of the internal cognitive state of an agent, anchors are first construed as *internal*. As in Asher's theory, they represent the set of descriptive conditions an agent associates with an individual he is familiar with - the mental file the agent associates with the object, to use a metaphor. To attribute an internal anchor to an agent, however, does not commit one to the *existence* of any object. Therefore, an anchor first models an *internal* aspect of someone's mind. Consider thus the following situation:

Because of sustained observations of the sky, John, an astronomer, is led to postulate the existence of a new planet, W . Note that this belief is formed because of causal interactions with the environment, even though there is in fact no planet matching the familiarity conditions John entertains about W . John has a mental file about W , and therefore we should attribute him an internal anchor for this imaginary object, as well as attitudes connected to this anchor:



A special characteristic of internal anchors is that they are intimately associated with a mental referent: whereas a belief may concern many different individuals, an internal anchor necessarily expresses familiarity conditions associated with *one* distinguished object. This remark motivates the following definition:

Definition 13 An internally anchored articulated DRS is a set consisting of pairs of the form $\langle MOD, K \rangle$, where:

- MOD is one of $BEL, DES, INT, [ANCH, x]$ (with x a distinguished discourse referent)
- K is a DRS

An internal anchor can be externally anchored or not, depending on the circumstances. For instance, John's internal anchor is not externally anchored in the actual world, since there is no planet he is causally related with in the way described by his familiarity conditions for x , but it may be externally anchored in another possible world. To formalise this idea, Kamp defines external anchors as functions from the set of internally anchored discourse referents of an articulated DRS onto elements of the domain of discourse.


We should note that, in this framework, the traditional *de dicto/de re* distinction is seriously blurred. A belief can be purely *de dicto*, when it is not based upon an anchor, or purely *de re*, when it belongs to an externally anchored articulated DRS. But there is a middle term, as our example about the non-existent planet shows: a belief can be based upon an internal anchor which, because of the connexions between the agent's cognitive state and the world, lacks an external anchor.

It is intuitively clear that an anchored articulated DRS will be well-suited to characterise the total cognitive state of an agent - with the inclusion of the dependency relations between attitudes. It is necessary, however, to investigate the semantics of these states in a more detailed manner.

4.3.2 Kamp's semantics for mental states

To begin with, we should note that the definition of an open proposition stated above can be useful here as well. According to Kamp, an open proposition based on a set X of discourse referents is a set of pairs $\langle f, w \rangle$, where w is a member of W , and f is a function from X onto the domain of discourse. Intuitively, we should conceive an open proposition as the semantic correlate of a DRS having X for universe.

It is important, however, to notice that the notion of an open proposition is not enough to model a whole cognitive state. Take, for instance, the following situation (due again to Kamp), in which my intention to pick up a coin lying in front of me is motivated by the charitable desire to give it to a child:



The content of the desire, in this situation, presupposes the content of the belief. Does it mean, however, that we should decide to take the result of filtering the open proposition expressed by the

belief with the condition associated with the desire as its content? According to Kamp, this would not do, because the content of the desire could *itself* serve as a new propositional basis for another desire concerning a new aspect of the situation. This would happen if I formed the desire that the child give the gold coin to his mother: the DRS representing this new desire would introduce a new discourse referent for the child's mother. Therefore, we must also take into account the fact that some new discourse referents can have been introduced. That is the aim of the following definition:

Definition 14 *The open proposition P_2 based on the open proposition P_1 and determined by the DRS K is the set:*

$$\{\langle w, f \rangle \mid (\exists g)(\langle w, g \rangle \in P_1 \wedge g \subseteq f \wedge f \text{ satisfies } K \text{ in } w)\}$$

This definition being relative to a DRS K , we should generalise it to get a notion of an open proposition based on another open proposition:

Definition 15 *An open proposition P_2 based on the open proposition P_1 is any set:*

$$\{\langle w, f \rangle \mid (\exists g)(\langle w, g \rangle \in P_1 \wedge g \subseteq f)\}$$

Note that given an open proposition, there is always another one on which it is based. This is due to the fact that any open proposition can be based on the tautological open proposition, which conveys no information. We may therefore consider the notion of 'being based on' as an order relation.


Theorem 1 *The relation defined between two open propositions by 'to be based on' is an order relation*

Let us show first that the relation is transitive. Suppose that P_3 is based on P_2 , and that P_2 is based on P_1 . Let $\langle f, w \rangle$ be a member of P_3 . Since P_3 is based on P_2 , there is a pair $\langle g, w \rangle$ in P_2 such that $g \subseteq f$. Now, since P_2 is based on P_1 , there is a pair $\langle h, w \rangle$ in P_1 such that $h \subseteq g$. Since $g \subseteq f$, $h \subseteq f$. So for each pair $\langle w, f \rangle$ in P_3 , there is a pair $\langle w, h \rangle$ in P_1 such that $h \subseteq f$. Hence P_3 is based on P_1 , which means that the relation is transitive.

Let us show now that the relation is anti-symmetric. Suppose that P_2 is based on P_1 , and that P_1 is based on P_2 . The fact that P_2 is based on P_1 implies that, for each member $\langle w, f \rangle$ of P_2 , there must be a member $\langle w, g \rangle$ of P_1 such that $g \subseteq f$. Now, since P_1 is based on P_2 , there must be a member $\langle w, h \rangle$ of P_2 such that $h \subseteq g$. But then, either $g = f$, or $h \subset g$. Since every function in an open proposition must have the same domain, we cannot have $h \subset g$. Therefore, $f = g$, which means that $P_2 = P_1$.

To be of real use, this notion of an order relation between open propositions should be supplemented with an ordering of the different attitudinal modes, in order to find an adequate structure for interpreting anchored structured DRSs.

This order relation should reflect the presuppositions holding between different kinds of attitudes. Consider again, for instance, the following example:



(21) $\langle [ANCH, x] \rangle$

$\langle BEL_x \rangle$

$\langle DES_x \rangle$

It seems clear that the desire presupposes the belief, which itself presupposes an anchoring of the discourse referent x . In general, we may recognise the following presuppositional relations:

1. A belief can presuppose another belief, or an anchor, when it has a singular aspect
2. A desire can presuppose a belief, or simply an anchor, when the desire is singular, or else another desire
3. An intention can presuppose a belief, an anchor, or another intention

We are now ready to characterise adequate denotations for structured DRS.

Definition 16 *An attitudinal state is a set of pairs $\langle ATT, p \rangle$, where $ATT \in \{ANCH, BEL, DES, INT\}$. Furthermore, a partial order $\leq$ fulfilling the following conditions must be definable on the set:*

- Suppose that $\langle ATT, p \rangle \leq \langle ATT', p' \rangle$. Then, one of the following must hold:
 - ATT' is BEL and ATT is BEL or ANCH
 - ATT' is ANCH and ATT is ANCH
 - ATT' is DES and ATT is BEL or ANCH or DES
 - ATT' is INT and ATT is ANCH or BEL or INT
- Suppose that $\langle ATT, p \rangle \leq \langle ATT', p' \rangle$. Then p' must be based on p .

We can supplement this notion by putting external anchors into the story:

Definition 17 *An externally anchored attitudinal state is a pair consisting in an attitudinal state and a function An from some set U of distinguished discourse referents of the state into the domain D of the model*

As in Asher's theory, an external anchor constrains the value that internal anchors may have by associating them with individuals. Thus, Kamp's theory can provide adequate models of singular beliefs.

4.3.3 The logic of anchors

Since the introduction of anchors as separate attitudes is one of the main innovations in Kamp's theory, we must state precisely their relation with the other attitudes – beliefs, in particular. One thing which seems clear is that anchors entail beliefs: if one ascribe a set of familiarity conditions about an object to an agent, one is also entitled to ascribe a singular belief to this agent. Therefore, we should have:

$$(22) \text{ Anch}_x(P) \rightarrow Bel_x(P)$$

The converse is not true. Suppose that Pierre believes Kim is a politician, and that he believes that every politician is honest. He will believe, in particular, that Kim is honest. But intuitively, I see no reason why we should attribute to Pierre any familiarity with Kim as an honest man. Furthermore, anchors do not share the logical properties of beliefs. For instance, we have

$$(23) \text{ Bel}_x(P \wedge Q) \rightarrow Bel_x(P)$$

But does the corresponding axiom for anchors,

$$(24) \text{ Anch}_x(P \wedge Q) \rightarrow \text{Anch}_x(P)$$

hold? This is less clear. Maybe *all* the information contained in an anchor is essential for an agent to be able to entertain particular thoughts about an object. More importantly, the logic of beliefs validates

$$(25) (\text{Bel}_x(P) \wedge \text{Bel}_x(Q)) \rightarrow \text{Bel}_x(P \wedge Q)$$


But this axiom has no counterpart with anchors. It is very important that an agent may have two anchors containing some descriptive information about an object without having the conjunction of these anchors.

Some other axioms which hold for beliefs also hold for anchors, however. Thus, it seems the following statements should be valid:

- $$(26) \text{ Anch}_x(p \wedge q) \rightarrow \text{Anch}_x(q \wedge p)$$
- $$(27) \text{ Anch}_x(p \wedge (q \wedge r)) \rightarrow \text{Anch}_x((p \wedge q) \wedge r)$$

4.3.4 How does Kamp's semantics deal with the puzzles?

Let's consider first Pierre's puzzle, this paradigm of a Frege style case. A relevant portion of Pierre's attitudinal state, according to our Kampian framework, can be described by the following externally anchored DRS:



According to this representation, it is clear that the belief attribution

- (29) Pierre believes that London is pretty

is ambiguous. Indeed, it can either attribute to Pierre a belief anchored on a file concerning Londres, the beautiful city he had heard of during his childhood, or a belief anchored on a file concerning London, this awfully slum he lives in at the moment. According on whether the first anchoring or the second one is chosen, the attribution is right or wrong. Twins cases are handled as elegantly and easily as Frege's cases, since a given DRS may be externally anchored, in different circumstances, on different objects, expressing different singular propositions.

References

- [Ash86] Nicholas Asher. *Belief in Discourse Representation Theory*. *Journal of Philosophical Logic*, 15:127–189, 1986.
- [Ash87] Nicholas Asher. A typology for attitude verbs and their anaphoric properties. *Linguistics and Philosophy*, 10:125–197, 1987.
- [Ash89] Nicholas Asher. Belief, acceptance and belief reports. *Canadian Journal of Philosophy*, 19(3):327–362, September b1989.
- [Ash93] Nicholas Asher. *Reference to Abstract Objects in Discourse*. Number 50 in Studies in Linguistics and Philosophy. Kluwer, Dordrecht, 1993.
- [AL94] Nicholas Asher and Alex Lascarides. Intentions and information in discourse. In *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics*, pages 34–41, Las Cruces, New Mexico, 1994. New Mexico State University.
- [AL96] Nicholas Asher and Alex Lascarides. Bridging. In *Proceedings of the International Workshop on Semantic Underspecification*, Max Planck Institut, Berlin, 1996.
- [AS93] Nicholas Asher and Munindar Singh. A logic of intentions and beliefs. *Journal of Philosophical Logic*, 22:513–544, 1993.
- [BP81] Jon Barwise and John Perry. Situations and attitudes. *Journal of Philosophy*, 128(11):668–691, 1981.
- [BP83] Jon Barwise and John Perry. *Situations and Attitudes*. MIT Press, Cambridge, Mass., 1983.
- [Fod78] Jerry A. Fodor. Propositional attitudes. *The Monist*, 61:501–523, 1978.
- [Fod85] Jerry Alan Fodor. Fodor's guide to mental representation: the intelligent auntie's vademecum. *Mind*, 94:76–100, March 1985.
- [Fod94] Jerry Fodor. *The Elm and the Expert*. Jean Nicod Series. Bradford Books. MIT Press, Cambridge, Mass., 1994.
- [Fre77] Gottlob Frege. On sense and reference. In Peter Geach and Max Black, editors, *Translations from the Philosophical Writings of Gottlob Frege, 3rd Edition*, pages 56–78. Basil Blackwell, Oxford, 1977.
- [Geu95] Bart Geurts. *Presupposing* PhD Dissertation, University of Stuttgart, 1995.
- [GS91] Jeroen Groenendijk and Martin Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14(1):39–100, February 1991.
- [Hei82] Irene Heim. *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, Amherst, Mass., 1982.
- [Hei83] Irene Heim. On the projection problem for presuppositions. *Proceedings of the West Coast Conference on Formal Semantics 2*: 114–126, 1983.
- [Hei92] Irene Heim. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9: 183–221.
- [Hin62] Jaacko Hintikka *Knowledge and Belief* Cornell University Press, Ithaca, N.Y., 1962.
- [Hin69] Jaacko Hintikka Semantics for propositional attitudes. In: J.W. Davis, D.J. Hockney, and W.K. Wilson (eds.), *Philosophical Logic*. Redel, Dordrecht, pp. 21–45, 1969.
- [Kam85] Hans Kamp. Context, thought and communication. *Proceedings of the Aristotelian Society*, 85:239–261, 1984/1985.
- [Kam90] Hans Kamp. Prolegomena to a structural theory of belief and other attitudes. In C. Anthony Anderson and Joseph Owens, editors, *Propositional Attitudes: The Role of Content in Logic, Language, and Mind*, number 20 in CSLI Lecture Notes, chapter 2, pages 27–90. Center for the Study of Language and Information, Stanford, Ca., 1990.
- [Kam97] Hans Kamp. *Manuscript on Propositional Attitudes*. 1997.
- [KR93] Hans Kamp and Uwe Reyle. *From Discourse to Logic*. Dordrecht: Kluwer, 1993.

- [Kon86] Kurt Konolige. *A deduction model of belief*. Research Notes in Artificial Intelligence, Pitman, London, 1986.
- [Kri79] Saul A. Kripke. A puzzle about belief. In Avishai Margalit, editor, *Meaning and Use: Papers Presented at the Second Jerusalem Philosophical Encounter, April 1976*, number 3 in Synthese Language Library, pages 239–283. D. Reidel, London, 1979.
- [Per93] John Perry. *The Problem of the Essential Indexical and Other Essays*. Oxford University Press, New York, 1993.
- [Qui60] Willard Van Orman Quine. *Word and Object*. MIT Press, Cambridge, Mass., 1960.
- [Rec93] Francois Recanati. *Direct Reference: from Language to Thought*. Blackwell, Oxford, 1993.
- [Sta84] Robert C. Stalnaker. *Inquiry*. MIT Press, Cambridge, Mass., 1984.
- [Sch87] Steden Schiffer. *Remnants of Meaning*. Bradford Books. MIT Press, Cambridge, Mass., 1987.
- [Zee87] Henk Zeevat. A treatment of belief sentences in discourse representation theory. In Jeroen Groenendijk, Dick de Jongh, and Martin Stokhof, editors, *Studies in Discourse Representation Theory and the Theory of Generalized Quantifiers*, number 8 in Groningen-Amsterdam Studies in Semantics, pages 189–215. Foris Publications, Dordrecht, 1987.