This thesis contains three logical investigations into dynamic semantics. The subjects of these three investigations are:

- An application of dynamic semantics to the Problem of the Liar Paradox and other circular propositions (Chapter 2).
- A *theoretical* investigation of notions of logical consequence in dynamic semantics (Chapter 3).
- An *extension* of dynamic semantics to various Systems of dynamic epistemic logic that deal with changes of higher-Order information (Chapter 4).

These three issues are different, and accordingly, the three chapters in which they are addressed can be read independently. Each of the chapters comes with its own introduction and conclusion. In this general introduction 1 will give a sketch of what dynamic semantics is, and a short overview of the three chapters that follow.

Dynamic Semantics

Dynamic semantics may be characterized by its aim, its motivation, and by the way in which its aim is realized. The aim of dynamic semantics is to raise the level of semantic analysis from the level of the sentence to the level of the text. This is motivated by various semantic phenomena that can only be explained by taking account of the order in which sentences in a text occur. There are two paradigmatic examples of this order sensitivity. First there is the phenomenon of coreference between an anaphor and its antecedent, as in

1. A man walks in the park. He whistles.
2. He whistles. A man walks in the park.

In a proper semantic description of 1, the indefinite ‘a man’ and the pronoun ‘he’ must refer to the same person, but in 2, we do not want an automatic coreference. This problem, of finding a semantic mechanism that establishes coreference between an anaphor and its antecedent, was one of the key motives for the development of Discourse Representation Theory *(DRT; see* [Kam81, Hei82, KR93]), and of Dynamic Predicate Logic *(DPL; see (GS91]).*

The second standard example of order sensitivity pertains to epistemic modals, such as ‘might’.

3. It might be raining.(. . .) It isn’t raining.
4. It isn’t raining.(. . .) It might be raining.

Here the idea is that a possibility that is open at one stage can be ruled out in a later stage of the discourse, as in 3. Thus we can regard the interpretation process of 3 as a successive elimination of possibilities. Under this perspective, 4 is an incoherent discourse, since a possibility that is ruled out cannot be reintroduced at a later stage, at least not without revision. These types of phenomena were an important motivation for the development of Update Semantics (US; see [Vel9l], where also a dynamic analysis of the expressions ‘normally’ and ‘presumably’ is given).

Both kinds of phenomena as illustrated in these examples call for a type of incremental semantics of texts, where the interpretation of the text is obtained by a step by step interpretation of the sentences that constitute the text, in the order in which they occur. In dynamic semantics this is achieved by a special view on the meaning of sentences. In dynamic semantics, sentence meanings are ***information change potentials.*** This can be explained as follows. A system of dynamic semantics can be seen as a triple $\mathcal{L}, \mathcal{K}, [\cdot].$, where $\mathcal{L}$ is some language, K a class of models, and $[\cdot].$ is an operation that assigns to each sentence ϕ of $\mathcal{L}$ and each model $M \in K$ a binary relation $[\phi]_{\mathcal{M}}$. The idea is that each model $M \in K$ comes with a set of ***information states*** $\mathcal{S}_{\mathcal{M}}$, and that the relations $[\phi]_{\mathcal{M}}$ are relations over $\mathcal{S}_{\mathcal{M}}$. These relations are the information change potentials: if (s, t) $\in [\phi]_{\mathcal{M}}$ then processing ϕ in the input state s can result in the output state t. The interpretation of texts is then obtained by taking the relational composition of the relations that are the meanings of the sentences of the text in the order in which they occur. Since relational composition is not a commutative operation, changing the order of the sentences of a text can change the meaning of the text.

In this brief description I have tried to isolate the most central ideas of dynamic semantics. It does not do justice to the many differences between the systems of dynamic semantics that can be found in the literature. For example, in US the relations actually are functions. Also, there is variation in the notion of information that is used: in ***DPL*** this is roughly information about the values of variables, but in US it is information about the world. The interested reader is referred to the cited papers. For an overview of dynamic semantics see [vBMV95]; for ‘state of the art’ work in dynamic semantics, see [Dek93, Ver94, GSV94a, GSV94b, Bea95].

Chapter 2: Dynamic Semantics and Circular Propositions

In [BE87], Barwise and Etchemendy develop two accounts of paradoxical sentences like the Liar Paradox. One of these is based on Austin’s conception of truth ([Aus50]), and accordingly is called the Austinian account. The authors take view that one of the lessons of this analysis is that the Liar shows that there is a “contextual parameter, one corresponding to Austin’s described situation, a parameter whose value necessarily ***changes*** [my italics] with the utterance of, or reasoning about, a sentence like the Liar.”[1] This observation was the key to the investigations of chapter 2, where I will use a form of Update Semantics ([Vel9l]) to formalize this context-shift.

Technically, the semantics developed in chapter 2 is an extension of Barwise and Etchemendy’s Austinian account: we take over the ontology and the formal language, and define a dynamic semantics for that language. But philosophically, it also amounts to an adjustment of the Austinian account as given by Barwise and Etchemendy. Here the main point is that saying that an utterance changes the described situation (as in the quotation I just gave) comes down to classifying the utterance as a performative speech act. But in the case of the Liar that is wrong. If an utterance of the Liar changes any situation at all, it changes the ***utterance*** situation, or, more precisely, it changes the information of the hearer (or reader). And the latter motivates the use of Update Semantics, which is a dynamic, information oriented semantics, for explaining the ‘dynamic’ aspect of sentences like the Liar.

I will show that in the dynamic semantics, paradoxical sentences like the Liar will have consistent (non-empty) update relations, but that these relations will have strange properties, such as anti-success. These strange properties sharply distinguish them from descriptive sentences, and can be regarded as dynamic notions of paradoxality. One of the theoretical pay-offs of the dynamic semantics will be a semantics for texts with circular cross-references.

Chapter 3: Logical Consequence in Dynamic Semantics

One of the most striking features of dynamic semantics, at least for a logician, is that there is no uniform definition of valid consequence. By contrast, in truth conditional semantics there is such a uniform definition, namely the classical (Tarskian) definition of truth and validity. But in dynamic semantics, there are different definitions around. For example, the notion of consequence in ***DPL*** is radically different from the notion of consequence in US.

---

[1] [BE87], p. 175.

One important aspect in which dynamic notions of consequence differ from classical consequence is in the structural rules of inference they validate. The first part of chapter 3 concentrates on this issue, and is inspired on work by van Benthem in [vB9la, vB9lb]. We take an abstract perspective on dynamic consequence relations, and define a concept of pure structural rule. We will review structural completeness results for six dynamic consequence relations that are found in the literature. Next, we make a case study of Update Test Consequence, the most natural notion of consequence in US. We discuss its structural inference rules, via a representation technique that is based on a strong connection with modal logic. We also try to think about an extension of this method for proving completeness theorems for Update Test Consequence that not only involve structural rules, but also logical rules for connectives.

The abstract discussion of the structural properties of dynamic consequence relations, as well as the use of modal techniques for Update Test Consequence, suggest a tight connection with Propositional Dynamic Logic. This connection is studied in detail, by the development of an extension of *PDL* with a 'loop' modality, in which we can simulate the structural deductions of all dynamic consequence relations we consider.

Finally, the last part of chapter 3 is devoted entirely to the Update Semantics of Veltman. We prove completeness theorems for the *might* systems. Furthermore we show that the *normally* fragment of the *normally -presumably* system is purely statical. We provide an alternative but equivalent truth-conditional semantics for this fragment, and prove completeness for this static semantics.

One of the conclusions of this chapter will be that there is, at present, no rigorous mathematical property that is distinctive of dynamic consequence. We could hope that the investigation into the structural properties of the dynamic consequence relations would reveal a minimal set of rules that all relations satisfy. Unfortunately, the dynamic inference relations we will consider do not share any non-trivial structural rule. Nevertheless we can give a negative characterization of dynamic consequence, and that is that they all lack the structural rule of Permutation. Since in fact there are infinitely many structural rules that all consequence relations we discuss lack, the special status of Permutation needs some additional motivation. This motivation is simply that the main reason for the development of dynamic semantics is that the order in which sentences occur in a text matters. If the notion of consequence for some system of dynamic semantics is to take account of this order-sensitivity, then the consequence relation will treat the premises of an argument as a sequence rather than as a set. For dynamic consequence relations this is not merely a choice that we could have made otherwise (as in sequential sequent systems for classical consequence), but it reflects the fundamental idea that first doing A and then doing B is something different from first doing B and then doing A.

Chapter 4: Dynamic Epistemic Logic

The objective of chapter 4 is to extend epistemic logic to a system that takes account of changes in knowledge, or belief, or information. There are two main motivations why such an extension of epistemic logic is desirable. First, it is just a pervasive property of knowledge (and of belief and information), that it is not constant, but changes. So an epistemic logic that also covers this dynamic aspect will be an improved epistemic logic. Second, the development of a dynamic epistemic logic is a first step to a logical framework in which we can study *information exchange,* either human or electronic. This intended application will be motivated more extensively in the introduction of the chapter.

The chapter will not reach the stage of the intended application, but will be devoted to the logical problem of defining a reasonable notion of update over higher-order information states. Nevertheless, the intended application to a logic of information exchange does motivate two important decisions about the structure of dynamic epistemic logic. First, the logic will have to be multi-agent. Second, in a proper description of information exchange, the agents higher-order information, in particular information about the information of other actors, will play a prominent part. So dynamic epistemic logic has to be multi-agent and higher-order.

I will approach this problem by extending the language of epistemic logic with formulae of the form of $[\phi]_a\psi$, which have as intended interpretation that after an update of actor a's information with ϕ, ψ is true.[2] I will establish two different ways of providing a semantics for this extended language. These correspond to two different notions of update, which I will call 'eliminative' and 'conscious' update. The eliminative updates are combined with the minimal modal logic K, both for the one agent and the multi-agent case, resulting in the systems Eliminative *K* and Multi-agent Eliminative *K*. In these systems, a minimal kind of static awareness is coupled with a minimal strategy of incorporating new information. We will provide complete axiomatizations for both of these logics via translations to *K*.

Conscious updates are intended to describe a more involved process of 'knowingly changing one's mind', and are developed as a dynamic extension of the modal logic K45. Here, the actors 'static ability to correctly represent their information, which is expressed by the introspection axioms of K45, is combined with a conscious way of incorporating new information, in the sense that the new information is also reflected at higher-order levels. The resulting logic is called Conscious K45, for which I will also present a complete axiomatization.

The existence of these two fundamentally different types of update operations has some important consequences for the meta-theoretical organization of dynamic epistemic logic. Since the two notions of update appear not to be definable in terms of each other, we have to face the existence of (at least) two

[2] The similar approach of [Jas94] is discussed in section 4.6.

different families of dynamic epistemic logics. But even within one such family, the organization is not as smooth as in ordinary modal logic. Most prominently (as I will argue in section 4.3), completeness theorems in dynamic epistemic logic will not be additive in the sense in which they are in standard modal logic, since a restriction on the model class will not necessarily induce an extension of the axiom system.

Let me conclude this introductory section by pointing out a theme common to the last two chapters. In both chapter 3 and 4, the use of standard notions and techniques from modal logic proves very useful. In chapter 3 this concerns the case study of Update Test Consequence (section 3.3), the translations of structural consequence in a variant of ***PDL*** (section 3.4), and the alternative truth-conditional semantics for normally (section 3.5.7). And in chapter 4, the language of dynamic epistemic logic we study is just an extension of a standard multi-modal vocabulary. Moreover, the type of structures we consider in that chapter, Kripke models (in the discussion of section 4.3) and modal structures (section 4.4), are also quite well known. Thus, although dynamic semantics is a quite recent development, in many case we can find proper tools for studying it in the standard repertoire of logic.