0 INTRODUCTION

Knowledge and belief play an important role in everyday life. In fact, most of what we do has to do with the things we know or believe. Likewise, it is not so strange that when we have to specify the behaviour of artificial agents in order to program or implement them in some particular way, it is thought to be important to be interested in the 'knowledge' and 'belief' of such an agent. In many areas of computer science and artificial intelligence one is concerned with the description or representation of knowledge of users or even the systems themselves. For example, in database theory one tries to model knowledge about parts of reality in certain formal ways to render it implementable and accessible to users. In AI one tries to design knowledge-based decision-support systems that are intended to assist professional users in some specialistic field when making decisions by providing pieces of knowledge and preferably some deductions from the input data by means of some inference mechanism. The representation and manipulation of knowledge of some sort is ubiquitous in the information sciences.

This book is not about knowledge representation in general, but rather concentrates on the logic of knowledge and belief. What (logical) properties do knowledge and belief have? What is the difference between knowledge and belief? We do not intend to answer these questions in a deep philosophical discussion of these notions. (For this the reader is referred to the philosophical literature, e.g. [Pol86], [Hin86], or, more technically, [Voo92].) Rather, we will look at a basic logical approach of epistemic logic based on modal logic, and show how this can be applied to problems that concern computer science and AI, such as the characterisation of a state of knowledge of an artificially intelligent agent or a processor in a computer network.

In particular, we will study various modal operators that have to do with (various kinds of) knowledge and belief, and with (certain or uncertain) information available to an agent more in general. Sometimes, we shall also briefly encounter other modalities, like time. In fact, we shall often have occasion to use more than one modality at the same time. There are several justifications for using multiple modalities: apart from the need to express the knowledge of *several agents* (denoting persons or processors in a distributed system), it might be interesting to combine *different interpretations* of several boxes in one system, like that of *knowledge and belief*, or *knowledge and time.*

What kind of logical properties may we expect knowledge and belief to have? Let us consider an example. Let us assume the following situation: When John lies to me, his heart-beat doubles. And when his heart-beat doubles, I know that he is lying to me. Now I can conclude the following: whenever John lies to me, I know that he does so. Formally (letting $K\varphi$ stand for "I know that φ"):

$$\text{from } p,\ p \to q,\ q \to Kp \text{ conclude } Kp. \tag{1}$$

Even from the point of view of propositional logic this scheme is completely sound. If we think that this is an unpleasant situation in the case of our interpretation, we must have a closer look at the premises of (1). Indeed, the last premise ($q \to Kp$: "if John's heart-beat doubles, I know he is lying") does not seem to make sense: the doubling of John's heart-beat does not affect my knowledge as long I am not aware of this change in his physical condition. In other words, instead of the premise mentioned ($q \to Kp$), we probably prefer a premise with a shift of *scope* of the knowledge operator: I know that, if John's heart-beat doubles, he is lying to me. This would give the unsound scheme:

$$\text{from } p,\ p \to q,\ K(q \to p) \text{ conclude } Kp. \tag{2}$$

It is interesting to observe that the bare formal representation of the arguments, and the explicit use of modal operators, already enlightens matters tremendously. Moreover, formalisations like (1) and (2) offer a further analysis of the problem. For instance, under what circumstances *would* it be acceptable to conclude Kp? One way to think about this is to try another instantiation of (some of) the primitives. For example,

p: John lies to me (3)
q: John stammers while talking to me.

The difference between (1) and (3) is that, in (3), the evidence for John's lying will be clear to me, so it seems reasonable to add the premise that I know that John stammers to me (whenever he does), giving

$$\text{from } p,\ p \to q,\ q \to Kq,\ K(q \to p), \text{ conclude } Kp. \tag{4}$$

The reader may ask whether the argument in (4) is valid: the answer in the 'standard' modal epistemic logic that we shall treat in Chapter 1 would be positive, since it would use the following accepted subargument:

$$\text{from } K(q \to p) \text{ conclude } Kq \to Kp. \tag{5}$$

In fact, the argument (5) is considered a *characteristic scheme* for modal logic. It is valid in any approach in which the K-operator is considered as a modal (necessity-type) operator. However, we shall see in Chapter 2 that this validity is sometimes questionable.

How do we know whether inferences like (2) are unsound for the (modal) epistemic logic under consideration? The answer is that modal logic is provided with a very appealing *model theory*, in which exactly those conclusions that are derivable in the logic are valid. In this model theory, or *semantics*, the notion of *possible world* plays a vital role. This notion, for which the term *situation* is also used occasionally, goes back to Leibniz, and was further developed by philosophers like Carnap and Kanger. However, it was not until the formalisation by Kripke ([Kri63]) that this notion became fully recognised by modal logicians. To put it more strongly, Kripke's work was the start of a prosperous period in which all kinds of new modal logics were developed — not only dealing with knowledge and belief, but also with such notions as time, obligation and action — each of them equipped with its *possible world semantics* or *Kripke semantics*.

To get an idea of this Kripke-style semantics using possible worlds: consider a student in a classroom without a view to the outside of the building. As he is more interested in plans for the afternoon than in the lecture today, he is wondering whether it is raining outside (denoted by the proposition r). As he cannot see outside, he considers two situations (worlds) as possible, one in which r holds, and one in which ¬r holds. So he does not know (for sure) whether r holds. On the other hand, in both situations it holds that the lecture is boring (b) and (having completed primary school successfully) that $2 \times 2 = 4$. So, assuming that these two situations are the only ones that are considered possible by the student, he knows that b and that $2 \times 2 = 4$. In general, knowing an assertion φ is modelled by the property that all possible worlds (i.e. worlds that are considered possible on the basis of one's knowledge) satisfy φ.

In the example we formalized in (2), we had a situation, or world (say w), in which the premises express that the following hold: John lies to me (p), If John lies to me, his heart-beat doubles ($p \to q$) and I know that, if his heart-beat doubles, John lies to me: $K(q \to p)$. Assuming that p and q are the only propositions that are relevant, we are led to the following model of (epistemically) possible worlds:

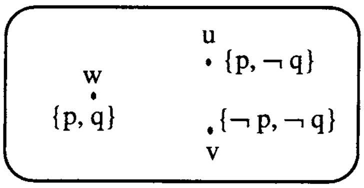

Figure 0.1

Note that, since the student knows (in world w) that q implies p, the only situations that are considered possible (epistemically) by him are worlds in which $(q \rightarrow p)$ is true (worlds u, v, and w itself). Now it is clear that the argument of (2) is not valid: in w, the formulas p, $p \rightarrow q$ and $K(q \rightarrow p)$ are true, but Kp is not; for, not in all worlds that were considered possible, p is true (it is not in v). We are deliberately slightly vague about what the the picture of Figure 0.1 denotes exactly (precise definitions will follow in Chapter 1): for the moment, let us assume it to be the set of worlds *that are epistemically compatible with the agent's knowledge at w*. In Chapter 1, we will make this idea of basing knowledge on possible worlds precise, and we shall see what consequences this approach has for the K-operator. In fact, here we shall lay a modal logic foundation of knowledge, treating both semantics and proof-theory. In Appendix A1 and A2 we shall also treat some alternative (non-modal) bases for epistemic logic that can be found in the literature. Furthermore in Chapter 1, we pay special attention to the logic of knowledge for one agent, where some simplifications can be made both semantically and syntactically (the latter pertaining to certain normal forms).

Specific applications of epistemic logic in the areas of computer science and artificial intelligence show up in a still growing variety. At the end of Chapter 1 and in Chapter 2, we shall discuss the use of epistemic logic for reasoning about distributed systems and intelligent systems. We shall see how epistemic logic is used for reasoning about protocols for communication (where processor A keeps sending a message until it knows that B knows the message — K_AK_Bm; and B, on its turn, has to know whether A knows this — $K_BK_AK_Bm$, in order to decide that a new message can be expected instead of the repetition of the old one). Furthermore, it is shown how modal epistemic logic can be fruitfully used to study other epistemic notions (like *explicit, implicit* (or *distributed*), and *common knowledge*) in the area of distributed systems. Finally, in this chapter attention is paid to other epistemic notions, such as various forms of *belief*, with some special focus on the problem of so-called logical omniscience: in standard epistemic (doxastic) logic, belief is closed under a number of idealized properties that are not always very realistic. We show several ways out of this problem at the expense of complicating the clear-cut modal semantics of the basic approach in some way or another. We also look briefly at the role of knowledge and belief in connection with other modalities such as *time* and *action*. Here we investigate expressions such as "believing that tomorrow I will believe that φ" and "I know that after the execution of the action α it will hold that ψ". These issues are especially important in the realm of AI research, for example in the area of *planning*. Chapter 2 ends with a treatment of a numerical version of standard epistemic logic, so-called *graded* modal logic, in which degrees of knowledge of a formula can be expressed by counting exceptions (in terms of worlds) to that formula.

Epistemic logic is especially useful as a 'meta'-tool to specify the contents and behaviour of, for instance, a knowledge base (KB). We may consider the facts that are stored in a KB as a set of formulas that are believed by that KB (we deliberately use the term 'belief' here, because a fact p in such a knowledge base need not be true, i.e. $Bp \wedge \neg p$ is consistent, which we do not allow for *known* formulas — cf. Chapter 2; $B\varphi$ stands for "I believe that φ"), and the modal language provides us (or the knowledge base itself) with a powerful tool to reason about its (non-) beliefs. We will discuss this in depth in Chapter 3, where we will give a number of theories to describe knowledge vs ignorance of an agent (or KB), such as *autoepistemic logic* (AEL), Levesque's logic of "all I know" and Halpern & Moses' theory of honest formulas.

However, as we shall see, the role of epistemic (and doxastic) logic in computer science and artificial intelligence goes even further. As we have said already, computer science is interlarded with *representation of* and *reasoning with* knowledge — where the reasoner may be *someone* or *something*, and the knowledge may pertain to some specified domain (the *Universe of Discourse* or *UoD*). A systematic description of "someone knows" is also relevant for the area of cognitive psychology and artificial intelligence, where we try to understand and represent *common-sense knowledge* and describe and simulate/implement *common-sense reasoning*, i.e. knowledge and reasoning about everyday things (in some UoD) in an everyday (human) way. This is less simple than it might appear at first glance. Contrary to what we are led to believe from lectures on logic, humans seldom employ just plain propositional or predicate logic while reasoning, but often (mostly?) use quite different reasoning methods as well. Most of these are even logically unsound. On the other hand, there are good reasons for a rational agent to use certain forms of these in practice. (It is not just reasoning in a wrong way.) A good example is *default reasoning*: unless it is known to the contrary, we assume (expect) something to hold. Reasoning by default is commonly used in daily life, e.g. when one is going to attend a(n announced) meeting: by default, i.e. unless something extraordinary is *known* to be the case (e.g. that the meeting is cancelled because of the illness of a speaker, etc.), we presume (expect) that the meeting will take place, and act accordingly (e.g. by taking the car to go to the place of meeting). As is clear from the way we put it, this involves the notion of *knowledge* again, and we shall treat the subject of default reasoning from an epistemic perspective in Chapter 4, along with a related topic, namely that of *counterfactual reasoning* ('reasoning against the facts'), that is to say, reasoning assuming some assertion φ while (you *know* that) φ is not actually the case. In this chapter we also encounter other default-related issues that are of great importance to AI, such as problems of representing the dynamics of systems like the infamous frame problem, and discuss these from an epistemic point of view. Furthermore, (in Chapter 3) we digress into the related topic of *non-monotonic* or *defeasible reasoning* in general, in which it is possible to lose previously inferred conclusions when more information becomes available. We shall discuss this issue, which has been a prominent AI research topic for the last 15 years, at some length as well.

Finally, a word on notation and some notions of classical logic that we shall use in this book. We assume the reader to be familiar with classical propositional logic (see for this any of the numerous introductions to logic, e.g. the 'classic' [Men64] or the more recent [Gal90] and [Gam91]). Throughout the book we shall use the following notions and notations regarding classical propositional logic: we use $\mathfrak{t}$ and $\mathfrak{f}$ for the semantic truth values 'true' and 'false', respectively. For a classical valuation w (i.e., a function that assigns truth ($\mathfrak{t}$) or falsehood ($\mathfrak{f}$) to the atomic propositions) we denote the fact that w satisfies a (propositional) formula φ by $w \models \varphi$. (This is defined inductively as usual.) Classical logical entailment is denoted $\models_{prop}$; $\varphi \models_{prop} \psi$ is defined formally as: for all valuations w: $w \models \varphi$ implies $w \models \psi$. A (classical) tautology φ, denoted $\models_{prop}\varphi$ (or simply $\models \varphi$, when no confusion can arise), is a formula that is satisfied by all classical valuations, that is to say, $w \models \varphi$ for all valuations w. A formula φ is (classically) satisfiable if there is a valuation w such that $w \models \varphi$. Satisfaction and logical entailment are lifted to sets of formulas: $w \models \Phi$ if $w \models \varphi$ for all $\varphi \in \Phi$; $\Phi \models_{prop} \Psi$ is defined as $w \models \Phi$ implies $w \models \Psi$, for all w. The (classically) propositional closure (or theory) of a set Φ of formulas is denoted $Th_{prop}(\Phi)$ or simply $<\Phi>$: $Th_{prop}(\Phi) = <\Phi> = \{\psi \mid \Phi \models_{prop} \psi\}$. We also use this notation for formulas: $Th_{prop}(\varphi) = <\varphi> = \{\psi \mid \varphi \models_{prop} \psi\} = Th_{prop}(\{\varphi\}) = <\{\varphi\}>$. On the syntactic level we use $\Phi \vdash_{PC} \psi$ to denote the derivability of ψ from the premises in the set Φ in (some sound and complete axiomatisation of) propositional calculus, and $\vdash_{PC} \psi$ (or $PC \vdash \psi$) that ψ is a theorem in propositional calculus. (We also extend this notation to $\Phi \vdash_{PC} \Psi$ for sets Φ and Ψ, in the obvious way.) We shall use the usual notion of propositional (in)consistency: φ is propositionally inconsistent if $\vdash_{PC} \neg\varphi$, and φ is consistent otherwise. Often we shall have occasion to use the well-known completeness result of classical propositional logic: $\vdash_{PC} \psi$ iff $\models_{prop} \psi$ or, more generally, $\Phi \vdash_{PC} \psi$ iff $\Phi \models_{prop} \psi$, so that $Th_{prop}(\Phi)$ can be alternatively given by $Th_{prop}(\Phi) = <\Phi> = \{\psi \mid \Phi \vdash_{PC} \psi\}$. Furthermore, by this completeness we also have that (classical) satisfiability and consistency are equivalent notions. Finally, we use occasionally $\Rightarrow$, $\Leftrightarrow$ and the 'outlined' symbols $\forall$ and $\exists$ as abbreviations for the meta-logical expressions 'implies', 'is equivalent to', 'for all' and 'there exist(s)', respectively. We also use the non-English but very convenient word 'iff' for 'if and only if'.

1 BASICS: THE MODAL APPROACH TO KNOWLEDGE

1.1. The Language: Epistemic Formulas

In this subsection we introduce the language of knowledge that we shall consider in the first instance. Let **P** be a set of propositional constants (atoms); $\mathbf{P} = \{p_n \mid n \in \mathbb{N}\}$ or $\mathbf{P} = \{p_0, ..., p_{n-1}\}$ for some $n \in \mathbb{N}$. We also refer to propositional constants as 'atomic' or 'primitive' propositions. Furthermore, let **A** be a set of m 'agents'. For notational convenience, we shall take $\mathbf{A} = \{1, ..., m\}$. The set $\mathcal{L}^m_{\mathbf{K}}(\mathbf{P})$ of epistemic formulas $\varphi, \psi, ...$ over **A** is the smallest set closed under:

(i) If $p \in \mathbf{P}$ then $p \in \mathcal{L}^m_{\mathbf{K}}(\mathbf{P})$.

(ii) If $\varphi, \psi \in \mathcal{L}^m_{\mathbf{K}}(\mathbf{P})$ then $(\varphi \wedge \psi), \neg\varphi \in \mathcal{L}^m_{\mathbf{K}}(\mathbf{P})$.

(iii) If $\varphi \in \mathcal{L}^m_{\mathbf{K}}(\mathbf{P})$ then $K_i\varphi \in \mathcal{L}^m_{\mathbf{K}}(\mathbf{P})$, for all $i \in \mathbf{A}$.

Here $K_i\varphi$ is read as: "*agent* i *knows that* φ". From here on we omit the outermost parentheses, and write e.g. $p_1 \wedge \neg p_4$ instead of $(p_1 \wedge \neg p_4)$. Moreover we introduce $\varphi \vee \psi$, $\varphi \rightarrow \psi$ and $\varphi \leftrightarrow \psi$ as the usual abbreviations for $\neg(\neg\varphi \wedge \neg\psi)$, $\neg\varphi \vee \psi$ and $(\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$, respectively, together with $\bot = p_0 \wedge \neg p_0$ (*falsum*). Occasionally we shall also use the abbreviation **true** for $\neg\bot$. When the set **P** of primitive propositions is understood, we omit it and write $\mathcal{L}^m_{\mathbf{K}}$ rather than $\mathcal{L}^m_{\mathbf{K}}(\mathbf{P})$. We may furthermore use the abbreviation $M_i\varphi$ for $\neg K_i \neg\varphi$, read as: "*agent* i *considers* φ *as possible* (on the basis of his / her knowledge)". When we only have one agent (i.e. **A** is a singleton set), we often omit subscripts, and just write $K\varphi$ and $M\varphi$. Formulas without occurrences of the modal operators K and M are called *purely propositional* or *objective*. The set of objective formulas is denoted $\mathcal{L}_0(\mathbf{P})$ or simply $\mathcal{L}_0$.

1.1.1. EXAMPLE. The following are formulas in $\mathcal{L}^4_{\mathbf{K}}(\mathbf{P})$ (let $p, q \in \mathbf{P}$): p, K_1q, $\neg K_2p$, K_1K_2q, $K_2\neg K_1p$, $\neg K_1\neg M_1(p \vee q)$, $K_2(K_1q \vee K_1\neg q)$, $K_1(\neg K_3(K_4q \wedge M_1p) \vee K_2\neg q)$.

1.2. Kripke Structures

1.2.1. DEFINITION. A *Kripke structure* (or Kripke model) $\mathbb{M}$ is a tuple $\langle S, \pi, R_1, ..., R_m\rangle$ where:

(i) S is a non-empty set of *states*,

(ii) $\pi: S \to (\mathbf{P} \to \{\mathfrak{t}, \mathfrak{f}\})$ is a truth assignment to the propositional atoms per state,

(iii) $R_i \subseteq S \times S$ $(i = 1, ..., m)$ are the so-called *possibility/accessibility relations.*

A (Kripke) *world* w consists of a Kripke model $\mathbb{M}$ together with a distinguished state $s \in S$: $(\mathbb{M}, s)$. The interpretation of $(s, t) \in R_i$ is: in world $(\mathbb{M}, s)$ agent i considers world $(\mathbb{M}, t)$ as a *possible world*, that is to say a world that he considers as possible. In the context of epistemic logic these possible worlds are also called *epistemic alternatives*. This term emphasises that agent i considers such an alternative world as possible on the basis of his knowledge. For convenience, we may also write $R_i(s, t)$ instead of $(s, t) \in R_i$.

1.3. Kripke Semantics of Epistemic Formulas

We define the relation $w \models \varphi$ (φ *is true in* w, or w *satisfies* φ) by induction on the structure of the epistemic formula φ:

$$(\mathbb{M}, s) \models p \quad \Leftrightarrow \quad \pi(s)(p) = \mathfrak{t} \text{ for } p \in \mathbf{P}$$

$$(\mathbb{M}, s) \models \varphi \wedge \psi \quad \Leftrightarrow \quad (\mathbb{M}, s) \models \varphi \text{ and } (\mathbb{M}, s) \models \psi$$

$$(\mathbb{M}, s) \models \neg\varphi \quad \Leftrightarrow \quad (\mathbb{M}, s) \not\models \varphi$$

$$(\mathbb{M}, s) \models K_i\varphi \quad \Leftrightarrow \quad (\mathbb{M}, t) \models \varphi \text{ for all t with } (s, t) \in R_i.$$

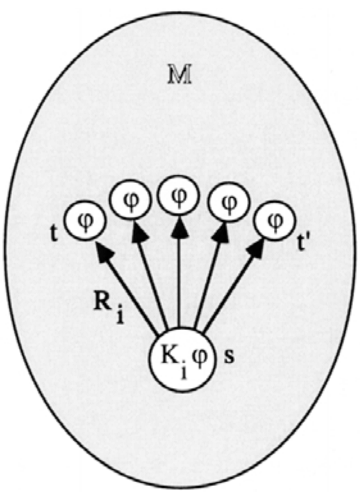

Figure 1.1

The last clause thus states "Agent i knows φ in world $(\mathbb{M}, s)$ iff φ is true in all worlds that i considers possible." (See Figure 1.1.) This can be explained as follows. In state s, the agent i has doubt about the true nature of the real world. So he considers several worlds as possible, namely the worlds t such that $R_i(s, t)$ holds. However, if in all possible worlds t with $R_i(s, t)$ it holds that φ, i has no doubts about the truth of φ: he *knows* φ (for certain).

1.3.0.1. EXERCISE. Show that regarding the abbreviations $\varphi \vee \psi$ and $\varphi \to \psi$ the following hold:

(i) $(\mathbb{M}, s) \models \varphi \vee \psi \Leftrightarrow (\mathbb{M}, s) \models \varphi$ or $(\mathbb{M}, s) \models \psi$,

(ii) $(\mathbb{M}, s) \models \varphi \to \psi \Leftrightarrow ((\mathbb{M}, s) \models \varphi \Rightarrow (\mathbb{M}, s) \models \psi)$.

1.3.0.2. EXERCISE. Show that: $(\mathbb{M}, s) \models M_i\varphi \Leftrightarrow$ there exists a t with $(s, t) \in R_i$ such that $(\mathbb{M}, t) \models \varphi$.

1.3.0.3. EXERCISE. Using the truth definition of formulas as given in this section, give the truth conditions for the formulas of Example 1.1.1.

1.3.1. EXAMPLE. Let $\mathbb{M}$ a Kripke model with $S = \{s_1, s_2, s_3\}$, $\mathbf{A} = \{\text{Alice, Bob}\}$, $\mathbf{P} = \{p\}$, and accessibility relations R_{Alice} (black arrows) and R_{Bob} (grey arrows) as in Figure 1.2.

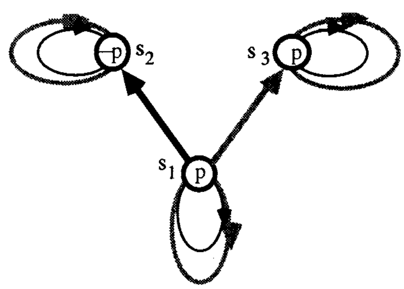

Figure 1.2

We now have (see Figure 1.2):

(i) $(\mathbb{M}, s_1) \models p$
(*in* s_1 p *holds*)

(ii) $(\mathbb{M}, s_1) \models \neg K_{Alice}p$
(*Alice does not know in* s_1 *that* p)

(iii) $(\mathbb{M}, s_1) \models \neg K_{Alice}p \wedge \neg K_{Alice}\neg p$
(*Alice does not know in* s_1 *whether* p *holds*)

(iv) $(\mathbb{M}, s_1) \models K_{Bob}p$
(*Bob knows in* s_1 *that* p, *since in* s_1 *and* s_3 p *holds*)

(v) $(\mathbb{M}, s_1) \models K_{Alice}(K_{Bob}p \vee K_{Bob}\neg p)$
(Alice knows in s_1 *that Bob knows whether* p *holds, since* $K_{Bob}p$ *holds in* s_1 *and* $K_{Bob}\neg p$ *in* s_2.)

(vi) $(\mathbb{M}, s_1) \models \neg K_{Bob}\neg K_{Alice}p$
(Bob does not know in s_1 *that Alice does not know that* p, *since in* s_1: $\neg K_{Alice}p$ *holds and in* s_3: $K_{Alice}p$ *holds.).*

1.3.1.1. EXERCISE. Check the above assertions.

1.3.2. EXERCISE. Let $\mathbf{P} = \{p, q\}$. Consider Kripke structure $\mathbb{M} = \langle S, \pi, R_A, R_B\rangle$ with $S = \{s_1, s_2, s_3, s_4\}$ and π, R_A, R_B as indicated in Figure 1.3 where black and grey arrows denote relations R_A and R_B, respectively.

Show that the following statements hold.

(i)	$(\mathbb{M}, s_1) \models p$	(vi)	$(\mathbb{M}, s_1) \models \neg K_B q$	(x)	$(\mathbb{M}, s_1) \models K_A\neg K_A q$
(ii)	$(\mathbb{M}, s_1) \models q$	(vii)	$(\mathbb{M}, s_1) \models K_A K_A p$	(xi)	$(\mathbb{M}, s_1) \models \neg K_B\neg K_A q$
(iii)	$(\mathbb{M}, s_1) \models K_A p$	(viii)	$(\mathbb{M}, s_1) \models \neg K_B K_A p$	(xii)	$(\mathbb{M}, s_1) \models \neg K_A\neg K_B q$
(iv)	$(\mathbb{M}, s_1) \models \neg K_B p$	(ix)	$(\mathbb{M}, s_1) \models \neg K_A\neg K_B p$	(xiii)	$(\mathbb{M}, s_1) \models K_B\neg K_B q$
(v)	$(\mathbb{M}, s_1) \models \neg K_A q$				

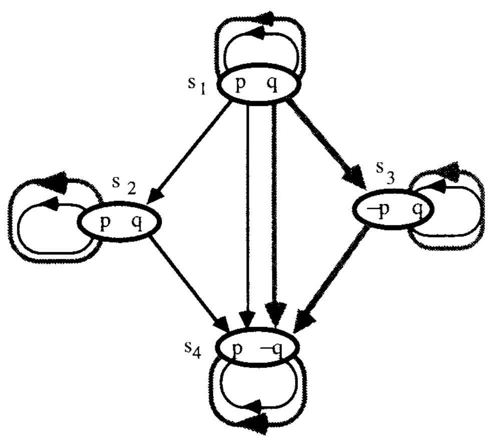

Figure 1.3

1.3.3. EXERCISE. Let $\mathbf{P} = \{p, q\}$. Construct Kripke structures $\mathbb{M} = \langle S, \pi, R_1, R_2\rangle$ such that for some $s \in S$:

(i) $(\mathbb{M}, s) \models K_1(K_1p \wedge K_2q) \wedge \neg K_1(p \wedge q)$