2 Basic Modal Logic

2.1 Introduction

2.1.1 In this chapter, we look at the basic technique – possible-world semantics – variations on which will occupy us for most of the following chapters. (We will return to the subject of the conditional in chapter 4.)

2.1.2 This will take us into an area called *modal logic*. This chapter concerns the most basic modal logic, *K* (after Kripke).

2.2 Necessity and Possibility

2.2.1 Modal logic concerns itself with the *modes* in which things may be true/false, particularly their possibility, necessity and impossibility. These notions are highly ambiguous, a subject to which we will return in the next chapter.

2.2.2 The modal semantics that we will examine employ the notion of a *possible world*. Exactly what possible worlds are, we will return to later in this chapter. For the present, the following will suffice. We can all imagine that things might have been different. For example, you can imagine that things are exactly the same, except that you are a centimetre taller. What you are imagining here is a different situation, or possible world. Of course, the actual world is a possible world too, and there are indefinitely many others as well, where you are two centimetres taller, three centimetres taller, where you have a different colour hair, where you were born in another country, and so on.

2.2.3 The other intuitive notion that the semantics employs is that of *relative possibility*. Given how things are now, it is possible for me to be in New York in a week's time, 26 January. Given how things will be in six days and twenty-three hours, it will no longer be possible. (I am writing in Brisbane.) Or, even if one countenances the possibility of some futuristic and exceptionally fast form of travel, assuming that I do not leave Brisbane in the next eight days, it will then be impossible for me to be in New York on 26 January. Hence, certain states of affairs are possible relative to some situations (worlds), but not others.

2.3 Modal Semantics

2.3.1 A propositional modal language augments the language of the propositional calculus with two monadic operators, $\Box$ and $\Diamond$.[1] Intuitively, $\Box A$ is read as 'It is necessarily the case that A'; $\Diamond A$ as 'It is possibly the case that A'.

2.3.2 Thus, the grammar of 1.2.2 is augmented with the rule:

If A is a formula, so are $\Box A$ and $\Diamond A$.

2.3.3 An *interpretation* for this language is a triple $\langle W, R, \nu \rangle$. W is a non-empty set. Formally, W is an arbitrary set of objects. Intuitively, its members are possible worlds. R is a binary relation on W (so that, technically, $R \subseteq W \times W$). Thus, if u and v are in W, R may or may not relate them to each other. If it does, we will write uRv, and say that v is *accessible* from u. Intuitively, R is a relation of relative possibility, so that uRv means that, relative to u, situation v is possible. ν is a function that assigns a truth value (1 or 0) to each pair comprising a world, w, and a propositional parameter, p. We write this as $\nu_w(p) = 1$ (or $\nu_w(p) = 0$). Intuitively, this is read as 'at world w, p is true (or false)'.

2.3.4 Given an interpretation, ν, this is extended to assign a truth value to every formula at every world by a recursive set of conditions. The conditions for the truth functions ($\neg, \wedge, \vee$, etc.) are the same as those for propositional logic (1.3.2), except that things are relativised to worlds. Thus, for $\neg$, $\wedge$ and $\vee$, the conditions go as follows. For any world $w \in W$:

$\nu_w(\neg A) = 1$ if $\nu_w(A) = 0$, and 0 otherwise.
$\nu_w(A \wedge B) = 1$ if $\nu_w(A) = \nu_w(B) = 1$, and 0 otherwise.
$\nu_w(A \vee B) = 1$ if $\nu_w(A) = 1$ or $\nu_w(B) = 1$, and 0 otherwise.

[1] Some logicians use L and M, respectively.

In other words, worlds play no essential role in the truth conditions for the non-modal operators.

2.3.5 They play an essential role in the truth conditions for the modal operators. For any world $w \in W$:

$\nu_w(\Diamond A) = 1$ if, for some $w' \in W$ such that wRw', $\nu_{w'}(A) = 1$; and 0 otherwise.
$\nu_w(\Box A) = 1$ if, for all $w' \in W$ such that wRw', $\nu_{w'}(A) = 1$; and 0 otherwise.

In other words, 'It is possibly the case that *A*' is true at a world, *w*, if *A* is true at *some* world, possible relative to *w*. And 'It is necessarily the case that *A*' is true at a world, *w*, if *A* is true at *every* world, possible relative to *w*.

2.3.6 Note that if *w* accesses no worlds, everything of the form $\Diamond A$ is false at *w* – if *w* accesses no worlds, it accesses no worlds at which *A* is true. And if *w* accesses no worlds, everything of the form $\Box A$ is true at *w* – if *w* accesses no worlds, then (vacuously) at all worlds that *w* accesses *A* is true.[2]

2.3.7 A finite interpretation (that is, where *W* is a finite set) can be perspicuously represented diagrammatically. For example, let $W = \{w_1, w_2, w_3\}$; w_1Rw_2, w_1Rw_3, w_3Rw_3 (and no other worlds are related by *R*); $\nu_{w_1}(p) = 0$, $\nu_{w_1}(q) = 0$; $\nu_{w_2}(p) = 1$, $\nu_{w_2}(q) = 1$; $\nu_{w_3}(p) = 1$, $\nu_{w_3}(q) = 0$. This interpretation can be represented as follows:

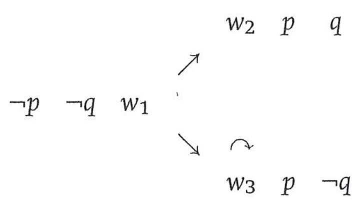

The arrows represent accessibility. In particular,

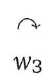

means that w_3 accesses itself.

[2] Recall that 'all *X*s are *Y*s' is logically equivalent to 'there are no *X*s that are nnot *Y*s'.

2.3.8 The truth conditions of 2.3.4 and 2.3.5 can be used to work out the truth values of compound sentences, and these can be marked on the diagram in the same way. For example, since p and q are true at w_2, so is $p \wedge q$. But $w_1 R w_2$; hence, $\Diamond(p \wedge q)$ is true at w_1. At the only world that w_3 accesses (namely itself), p is true. Hence, $\Box p$ is true at w_3. But w_1 accesses w_3, hence, $\Diamond\Box p$ is true at w_1. w_2 accesses no world; hence, $\Diamond q$ is false at w_2, so $\neg\Diamond q$ is true there. We can add these facts to the diagram in the obvious way:

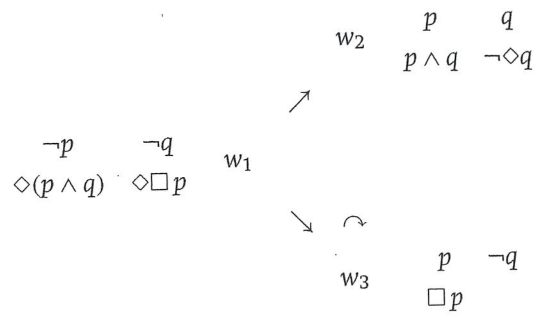

2.3.9 Observe that the truth value of $\neg\Diamond A$ at any world, w, is the same as that of $\Box\neg A$. For:

$$\begin{aligned} \nu_w(\neg\Diamond A) = 1 \quad & \text{iff} \quad \nu_w(\Diamond A) = 0 \\ & \text{iff} \quad \text{for all } w' \text{ such that } wRw',\ \nu_{w'}(A) = 0 \\ & \text{iff} \quad \text{for all } w' \text{ such that } wRw',\ \nu_{w'}(\neg A) = 1 \\ & \text{iff} \quad \nu_w(\Box\neg A) = 1 \end{aligned}$$

2.3.10 Similarly, the truth value of $\neg\Box A$ at a world is the same as that of $\Diamond\neg A$. The proof is left as an exercise.

2.3.11 An inference is valid if it is truth-preserving at all worlds of all interpretations. Thus, if Σ is a set of formulas and A is a formula, then semantic consequence and logical truth are defined as follows:

$\Sigma \models A$ iff for all interpretations $\langle W, R, \nu\rangle$ and all $w \in W$: if $\nu_w(B) = 1$ for all $B \in \Sigma$, then $\nu_w(A) = 1$.

$\models A$ iff $\phi \models A$, i.e., for all interpretations $\langle W, R, \nu\rangle$ and all $w \in W$, $\nu_w(A) = 1$.

2.4 Modal Tableaux

2.4.1 Tableaux for modal logic are similar to those for propositional logic (1.4), except for the following modifications. At every node of the tree there is either a formula and a natural number (0, 1, 2,. . .), thus: A, i; or something of the form irj, where i and j are natural numbers. Intuitively, different numbers indicate different possible worlds; A, i means that A is true at world i; and irj means that world i accesses world j.[3]

2.4.2 Second, the initial list for the tableau comprises $A, 0$, for every premise, A (if there are any), and $\neg B, 0$, where B is the conclusion.

2.4.3 Third, the rules for the truth-functional connectives are the same as in non-modal logic, except that the number associated with any formula is also associated with its immediate descendant(s). Thus, the rule for disjunction, for example, is:

$$\begin{array}{ccc} & A \vee B, i & \\ \swarrow & & \searrow \\ A, i & & B, i \end{array}$$

2.4.4 There are four new rules for the modal operators:

$$\begin{array}{cc} \neg\Box A, i & \neg\Diamond A, i \\ \downarrow & \downarrow \\ \Diamond\neg A, i & \Box\neg A, i \end{array}$$

$$\begin{array}{cc} \Box A, i & \Diamond A, i \\ irj & \downarrow \\ \downarrow & irj \\ A, j & A, j \end{array}$$

In the rule for $\Box$ (bottom left), both of the lines above the arrow must be present for the rule to be triggered (the lines do not have to occur in the order shown, and they do not have to be consecutive), and it is applied to *every* such j. In the rule for $\Diamond$ (bottom right), the number j must be *new*. That is, it must not occur on the branch anywhere above.

[3] I will avoid using r as a propositional parameter where this might lead to confusion.

2.4.5 Finally, a branch is closed iff for some formula, A, and number, i, A, i and $\neg A, i$ both occur on the branch. (It must be the same i in both cases.)[4]

2.4.6 Here are some examples of tableaux:

(i) $\Box(A \supset B) \wedge \Box(B \supset C) \vdash \Box(A \supset C)$.

$$\Box(A \supset B) \wedge \Box(B \supset C), 0$$
$$\neg\Box(A \supset C), 0$$
$$\Box(A \supset B), 0$$
$$\Box(B \supset C), 0$$
$$\Diamond\neg(A \supset C), 0$$
$$0r1 \qquad (1)$$
$$\neg(A \supset C), 1 \qquad (1)$$
$$A, 1$$
$$\neg C, 1$$
$$A \supset B, 1 \qquad (2)$$
$$B \supset C, 1 \qquad (2)$$
$$\swarrow \qquad \searrow$$
$$\neg A, 1 \qquad\qquad B, 1$$
$$\times \qquad\qquad \downarrow \qquad \searrow$$
$$\neg B, 1 \qquad C, 1$$
$$\times \qquad \times$$

The lines marked (1) are obtained by applying the rule for $\Diamond$ to the line immediately above them. Note that in applying the rule for $\Diamond$, a number new to the branch must be chosen. The lines marked (2) are the results of two applications of the rule for $\Box$ to the conjuncts of the premise. Note that the rule for $\Box$ is applied to numbers already on the branch.

(ii) $\vdash \Diamond(A \wedge B) \supset (\Diamond A \wedge \Diamond B)$. The arrow at the bottom of a branch indicates that it continues on the next page.

$$\neg(\Diamond(A \wedge B) \supset (\Diamond A \wedge \Diamond B)), 0$$
$$\Diamond(A \wedge B), 0$$
$$\neg(\Diamond A \wedge \Diamond B), 0$$
$$\swarrow \qquad \searrow$$
$$\neg\Diamond A, 0 \qquad\qquad \neg\Diamond B, 0$$
$$\Box\neg A, 0 \qquad\qquad \Box\neg B, 0$$
$$\downarrow \qquad\qquad \downarrow$$

[4] It is not obvious, but, as in the propositional case, every tableau of the kind we are dealing with here is finite.

$0r1$	$0r1$	(1)
$A \wedge B, 1$	$A \wedge B, 1$	(1)
$A, 1$	$A, 1$	
$B, 1$	$B, 1$	
$\neg A, 1$	$\neg B, 1$	(2)
$\times$	$\times$	

The lines marked (1) result from an application of the rule for $\Diamond$ to the formula at the second node of the tableau. The line marked (2) results from applications of the rule for $\Box$ to $\Box\neg A, 0$ (left branch) and $\Box\neg B, 0$ (right branch).

(iii) $\nvdash (\Diamond p \wedge \Diamond\neg q) \supset \Diamond\Box\Diamond p$

$\neg((\Diamond p \wedge \Diamond\neg q) \supset \Diamond\Box\Diamond p), 0$

$\Diamond p \wedge \Diamond\neg q, 0$

$\neg\Diamond\Box\Diamond p, 0$

$\Diamond p, 0$

$\Diamond\neg q, 0$

$\Box\neg\Box\Diamond p, 0$ (1)

$0r1$ (2)

$p, 1$ (2)

$\neg\Box\Diamond p, 1$ (3)

$\Diamond\neg\Diamond p, 1$

$1r2$

$\neg\Diamond p, 2$

$\Box\neg p, 2$

$0r3$ (4)

$\neg q, 3$ (4)

$\neg\Box\Diamond p, 3$ (5)

$\Diamond\neg\Diamond p, 3$

$3r4$

$\neg\Diamond p, 4$

$\Box\neg p, 4$

The lines marked (2) result from an application of the rule for $\Diamond$ to the fourth line of the tableau. The lines marked (4) result from an application of the same rule to the fifth line of the tableau. Note that, as the example shows, when we apply the rule for $\Diamond$, we may have to go back and apply the rule for $\Box$ again, to the new world (number) that has been introduced. Thus, the line marked (3) results from a first application of the rule to line (1). Line (5) results from a second application. For this reason, if one is ticking nodes to show that one has finished with them, one should never tick a node of the form $\Box A$, since one may have to come back and use it again.

2.4.7 Counter-models can be read off from an open branch of a tableau in a natural way. For each number, i, that occurs on the branch, there is a world, w_i; $w_i R w_j$ iff irj occurs on the branch; for every propositional parameter, p, if p, i occurs on the branch, $v_{w_i}(p) = 1$, if $\neg p, i$ occurs on the branch, $v_{w_i}(p) = 0$ (and if neither, $v_{w_i}(p)$ can be anything one wishes).

2.4.8 Thus, the counter-model given by the open (and only) branch of the third example of 2.4.6 is as follows: $W = \{w_0, w_1, w_2, w_3, w_4\}$. $w_0 R w_1$, $w_1 R w_2$, $w_0 R w_3$, $w_3 R w_4$. There are no other worlds related by R. $v_{w_1}(p) = 1$, $v_{w_3}(q) = 0$; otherwise, v is arbitrary. The interpretation can be depicted thus:

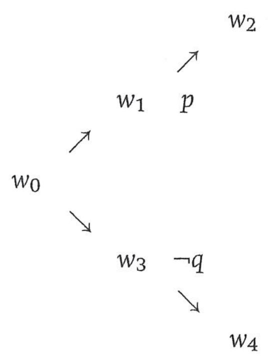

Using the truth conditions, one can check directly that the interpretation works. Since p is true at w_1, $\Diamond p$ is true at w_0. Similarly, $\Diamond\neg q$ is true at w_0. Hence, the antecedent is true at w_0. w_2 accesses no worlds; so $\Diamond p$ is false at w_2, and $\Box\Diamond p$ is false at w_1. Similarly, $\Box\Diamond p$ is false at w_3. Hence, there is no world which w_0 can access at which $\Box\Diamond p$ is true. Thus, $\Diamond\Box\Diamond p$ is false at w_0. It follows, then, that $(\Diamond p \wedge \Diamond\neg q) \supset \Diamond\Box\Diamond p$ is false at w_0.

2.4.9 The tableaux just described are sound and complete with respect to the semantics. The proof is given in 2.9.

2.5 Possible Worlds: Representation

2.5.1 In the rest of this chapter we look at the major philosophical question that modal semantics generate: what do they *mean*?

2.5.2 One might suggest that they do not mean anything. They are simply a mathematical apparatus – interpretations comprise just bunches of objects (W) furnished with some properties and relations – to be thought of purely instrumentally as delivering an appropriate notion of validity.

2.5.3 But there is something very unsatisfactory about this, as there is about all instrumentalisms. If a mathematical 'black box' gives what seem to be the right answers, one wants to know why. There must be some relationship between how it works and reality, which explains why it gets things right.

2.5.4 The most obvious explanation in this context is that the mathematical structures that are employed in interpretations *represent* something or other which underlies the correctness of the notion of validity.

2.5.5 In the same way, no one supposes that truth is simply the number 1. But that number, and the way that it behaves in truth-functional semantics, are able to represent truth, because the structure of their machinations corresponds to the structure of truth's own machinations. This explains why truth-functional validity works (when it does).

2.5.6 So, the question arises: what exactly, in reality, does the mathematical machinery of possible worlds represent? Possible worlds, of course (what else?). But what are they?

2.6 Modal Realism

2.6.1 The simplest suggestion (usually termed 'modal realism') is that possible worlds are things exactly like the actual world. They are composed of physical objects like people, chairs and stars (if any exist in those worlds), in their own space and time (if there are such things in those worlds). These objects exist just as much as you and I do, just in a different place/time – though not ones in this world.

2.6.2 The thought is, no doubt, a little mind-boggling. But so are many of the developments in modern physics. And why should metaphysics not have the right to boggle the mind just as much as physics?

2.6.3 Many arguments may be put both for and against this proposal – as they may be for all the views that I will mention. Here is one argument against. What makes such a world a *different* possible world, and not simply part of this one? The natural answer is that the space, time and causation of that world are unconnected with the space, time and causation of this world. One cannot travel from here to there in space or time; nor can causal processes from here reach there, or vice versa.

2.6.4 But why should that make it a *different* world? Suppose that because of the spatial geometry of the inside of a black hole, one could travel thence down a worm hole into a part of the cosmos with its own space and time; and suppose, then, that the worm hole closed up. We would not think of that region, now causally isolated from the rest, as a different possible world: merely an inaccessible part of this one.

2.6.5 The point may be put in a different way. Why should we think that something is possible in *this* world merely because it is actually happening at another place/time? I do not, after all, think that it is possible to see kangaroos in Antarctica merely because they are seen in Australia.

2.7 Modal Actualism

2.7.1 Another possibility (frequently termed 'modal actualism') is that, though possible worlds exist, they are not the physical entities that the modal realist takes them to be. They are entities of a different kind: specifically, abstract entities (like numbers, assuming there to be such things).

2.7.2 What kind of abstract entities? There are several possible candidates here. A natural one is to take them to be sets of propositions, or other language-like entities. Crudely, a possible world is individuated by the set of things true at it, which is just the set of propositions it contains.

2.7.3 But a problem arises with this suggestion when one asks which sets are worlds? Clearly not all sets are possible worlds. For example, a set that contains two propositions but not their conjunction could not be a possible world.

2.7.4 For a set of propositions to form a world, it must at least be closed under valid inference. (If a proposition is true at a world, and it entails another, then so is that.) But there's the rub. The machinery of worlds was meant to explain why certain inferences, and not others, are valid. But it now seems that the notion of validity is required to explain the notion of world – not the other way around.

2.7.5 A variation of actualism which avoids this problem is known as 'combinatorialism'. A possible world is merely the set of things in *this* world, rearranged in a different way. So in this world, my house is in Australia, and not China; but rearrange things, and it could be in China, and not Australia.

2.7.6 Combinatorialism is still a version of actualism, because an arrangement is, in fact, an abstract object. It is a *set* of objects with a certain structure. But it avoids the previous objection, since one may explain what combinations there are without invoking the notion of validity.

2.7.7 But combinatorialism has its own problems. For example, it would seem to be entirely possible that there is an object such that neither it nor any of its parts exist in this world. It is clear, though, that such an object could not exist in any world obtained simply by rearranging the objects in this world. Hence, there are possible worlds which cannot be delivered by combinatorialism.

2.8 Meinongianism

2.8.1 Both realism and actualism take possible worlds and their denizens, whatever they are, to exist, either as concrete objects or as abstract objects. Another possibility is to take them to be non-existent objects. (We know, after all, that such things do not really exist!)

2.8.2 We are all, after all, familiar with the thought that there are non-existent things, like fairies, Father Christmas (sorry) and phlogiston. Possible worlds are things of this kind.

2.8.3 The view that there are non-existent objects was espoused, famously, by Meinong. It had a very bad press for a long time in English-speaking philosophy, but it is fair to say that many of the old arguments against the possibility of there being non-existent objects are not especially cogent.

2.8.4 For example, one argument against such objects is that, since they cannot interact with us causally, we would have no way of knowing anything about them. But exactly the same is true, of course, of possible worlds as both the realist and the actualist conceive them, so this can hardly count to their advantage against Meinongianism about worlds.

2.8.5 Moreover, it is very clear how we know facts about at least some non-existent objects: they are simply stipulated. Holmes lived in Baker Street – and not Oxford Street – because Conan Doyle decided it was so.

2.8.6 The preceding considerations hardly settle the matter of the nature of possible worlds. There are many other suggested answers (most of which are some variation on one or other of the themes that I have mentioned); and there are many objections to the suggestions I have raised, other than the ones that I have given, as well as possible replies to the objections I have raised; philosophers can have hours of fun with possible worlds. This will do for the present, though.

2.9 *Proofs of Theorems

2.9.1 The soundness and completeness proofs for K are essentially variations and extensions of the soundness and completeness proofs for propositional logic. We redefine faithfulness and the induced interpretation. The proofs are then much as in 1.11.

2.9.2 Definition: Let $\mathcal{I} = \langle W, R, \nu \rangle$ be any modal interpretation, and b be any branch of a tableau. Then $\mathcal{I}$ is *faithful* to b iff there is a map, f, from the natural numbers to W such that:

For every node A, i on b, A is true at $f(i)$ in $\mathcal{I}$.
If irj is on b, $f(i)Rf(j)$ in $\mathcal{I}$.

We say that f shows $\mathcal{I}$ to be faithful to b.

2.9.3 Soundness Lemma: Let b be any branch of a tableau, and $\mathcal{I} = \langle W, R, \nu \rangle$ be any interpretation. If $\mathcal{I}$ is faithful to b, and a tableau rule is applied to it, then it produces at least one extension, b', such that $\mathcal{I}$ is faithful to b'.

Proof:
Let f be a function which shows $\mathcal{I}$ to be faithful to b. The proof proceeds by a case-by-case consideration of the tableau rules. The cases for the propositional rules are essentially as in 1.11.2. Suppose, for example, that

$A \wedge B, i$ is on b, and that we apply the rule for conjunction to give an extended branch containing A, i and B, i. Since $\mathcal{I}$ is faithful to b, $A \wedge B$ is true at $f(i)$. Hence, A and B are true at $f(i)$. Hence, $\mathcal{I}$ is faithful to the extension of b. We will therefore consider only the modal rules in detail. Consider the rule for negated $\Diamond$. Suppose that $\neg\Diamond A, i$ occurs on b, and that we apply the rule to extend the branch with $\Box\neg A, i$. Since $\mathcal{I}$ is faithful to b, $\neg\Diamond A$ is true at $f(i)$. Hence, $\Box\neg A$ is true at $f(i)$ (by 2.3.9). Hence, $\mathcal{I}$ is faithful to the extension of b. The rule for negated $\Box$ is similar (invoking 2.3.10).

This leaves the rules for $\Box$ and $\Diamond$. Suppose that $\Box A, i$ is on b, and that we apply the rule for $\Box$. Since $\mathcal{I}$ is faithful to b, $\Box A$ is true at $f(i)$. Moreover, for any i and j such that irj is on b, $f(i)Rf(j)$. Hence, by the truth conditions for $\Box$, A is true at $f(j)$, and so $\mathcal{I}$ is faithful to the extension of the branch. Finally, suppose that $\Diamond A, i$ is on b and we apply the rule for $\Diamond$ to get nodes of the form irj and A, j. Since $\mathcal{I}$ is faithful to b, $\Diamond A$ is true at $f(i)$. Hence, for some $w \in W$, $f(i)Rw$ and A is true at w. Let f' be the same as f except that $f'(j) = w$. Note that f' also shows that $\mathcal{I}$ is faithful to b, since f and f' differ only at j; this does not occur on b. Moreover, by definition, $f'(i)Rf'(j)$, and A is true at $f'(j)$. Hence, f' shows $\mathcal{I}$ to be faithful to the extended branch. ■

2.9.4 Soundness Theorem for K: For finite Σ, if $\Sigma \vdash A$ then $\Sigma \models A$.

Proof:
Suppose that $\Sigma \not\models A$. Then there is an interpretation, $\mathcal{I} = \langle W, R, \nu\rangle$, that makes every premise true, and A false, at some world, w. Let f be any function such that $f(0) = w$. This shows $\mathcal{I}$ to be faithful to the initial list. The proof is now exactly the same as in the non-modal case (1.11.3). ■

2.9.5 Definition: Let b be an open branch of a tableau. The interpretation, $\mathcal{I} = \langle W, R, \nu\rangle$, induced by b, is defined as in 2.4.7. $W = \{w_i : i \text{ occurs on } b\}$. $w_i R w_j$ iff irj occurs on b. If p, i occurs on b, then $\nu_{w_i}(p) = 1$; if $\neg p, i$ occurs on b, then $\nu_{w_i}(p) = 0$ (and otherwise $\nu_{w_i}(p)$ can be anything one likes).

2.9.6 Completeness Lemma: Let b be any open complete branch of a tableau. Let $\mathcal{I} = \langle W, R, \nu\rangle$ be the interpretation induced by b. Then:

if A, i is on b then A is true at w_i
if $\neg A, i$ is on b then A is false at w_i

Proof:
The proof is by recursion on the complexity of A. If A is atomic, the result is true by definition. If A occurs on b, and is of the form $B \vee C$, then the rule for disjunction has been applied to $B \vee C, i$. Thus, either B, i or C, i is on b. By induction hypothesis, either B or C is true at w_i. Hence, $B \vee C$ is true at w_i, as required. The case for $\neg(B \vee C)$ is similar, as are the cases for the other truth functions. Next, suppose that A is of the form $\Box B$. If $\Box B, i$ is on b, then for all j such that irj is on b, B, j is on b. By construction and the induction hypothesis, for all w_j such that $w_i R w_j$, B is true at w_j. Hence, $\Box B$ is true at w_i, as required. If $\neg\Box A, i$ is on b, then $\Diamond\neg A, i$ is on b; so, for some j, irj and $\neg A, j$ are on b. By induction hypothesis, $w_i R w_j$ and A is false at w_j. Hence, $\Box A$ is false at w_i as required. The case for $\Diamond$ is similar. ■

2.9.7 Completeness Theorem: For finite Σ, if $\Sigma \models A$ then $\Sigma \vdash A$.

Proof:
Suppose that $\Sigma \nvdash A$. Given an open branch of the tableau, the interpretation that this induces makes all the premises true at w_0 and A false at w_0 by the Completeness Lemma. Hence, $\Sigma \not\models A$. ■

2.10 History

Modal logic is as old as logic. Aristotle himself gave an account of which modal syllogisms he took to be valid (see Kneale and Kneale, 1975, ch. 2, sect. 8). Modal logic and semantics were also discussed widely in the Middle Ages (see Knuuttila, 1982). In the modern period, the subject of modal logic was initiated by C. I. Lewis just before the First World War (see Lewis and Langford, 1931). Initially, it received a bad press, largely as a result of the criticisms of Quine – whose work also produced much of the unpopularity of Meinongianism. (On both, see the papers in Quine, 1963.) Things changed with the invention of possible-world semantics in the early 1960s. These are due to the work of a number of people, most notably that of Kripke (1963a). (For a history, see Copeland, 1996, pp. 8–15.)

The notion of a possible world is to be found in Leibniz (e.g., *Monadology*, sect. 53). Modal realism has been espoused most famously by D. Lewis (1986). Notable proponents of actualism include Plantinga and Stalnaker. Combinatorialism is espoused by Cresswell. See the papers by all three in Loux (1979). The idea that worlds are non-existent objects is proposed in

Routley (1980a) and defended in Priest (2005c). Kripke's own views on the nature of possible worlds can be found in Kripke (1977).

2.11 Further Reading

Perhaps the best introduction to modal logic is still Hughes and Cresswell (1996). The semantics of *K* are given in chapter 2. (Hughes and Cresswell use axiom systems rather than tableaux for their proof theory.) Chellas (1980) is also excellent, though a little more demanding mathematically. Tableaux for modal propositional logics can be found in chapters 2 and 3 of Girle (2000). A somewhat different form can be found in chapter 2 of Fitting and Mendelsohn (1999). A useful collection of essays on the nature of possible worlds is Loux (1979); chapter 15, 'The Trouble with Possible Worlds', by Lycan, is a good orientational survey. Read (1994, ch. 4) is also an excellent discussion.

2.12 Problems

1. Check the details of 2.3.10.
2. Show the following. Where the tableau does not close, use it to define a counter-model, and draw this, as in 2.4.8.
   (a) $\vdash (\Box A \wedge \Box B) \supset \Box(A \wedge B)$
   (b) $\vdash (\Box A \vee \Box B) \supset \Box(A \vee B)$
   (c) $\vdash \Box A \equiv \neg\Diamond\neg A$
   (d) $\vdash \Diamond A \equiv \neg\Box\neg A$
   (e) $\vdash \Diamond(A \wedge B) \supset (\Diamond A \wedge \Diamond B)$
   (f) $\vdash \Diamond(A \vee B) \supset (\Diamond A \vee \Diamond B)$
   (g) $\Box(A \supset B) \vdash \Diamond A \supset \Diamond B$
   (h) $\Box A, \Diamond B \vdash \Diamond(A \wedge B)$
   (i) $\vdash \Box A \equiv \Box(\neg A \supset A)$
   (j) $\vdash \Box A \supset \Box(B \supset A)$
   (k) $\vdash \neg\Diamond B \supset \Box(B \supset A)$
   (l) $\nvdash \Box(p \vee q) \supset (\Box p \vee \Box q)$
   (m) $\Box p, \Box\neg q \nvdash \Box(p \supset q)$
   (n) $\Diamond p, \Diamond q \nvdash \Diamond(p \wedge q)$
   (o) $\nvdash \Box p \supset p$
   (p) $\nvdash \Box p \supset \Diamond p$

(q) $p \nvdash \Box p$
(r) $\nvdash \Box p \supset \Box\Box p$
(s) $\nvdash \Diamond p \supset \Diamond\Diamond p$
(t) $\nvdash p \supset \Box\Diamond p$
(u) $\nvdash \Diamond p \supset \Box\Diamond p$
(v) $\nvdash \Diamond(p \vee \neg p)$

3. How might one reply to the objections of 2.5–2.8, and what other objections are there to the views on the nature of possible worlds explained there? What other views could there be?
4. *Check the details omitted in 2.9.3 and 2.9.6.