

Formal Approaches to Modality

*Stefan Kaufmann, Cleo Condoravdi
and Valentina Harizanov*

1. Modal logic

Modal notions are pervasive in the meaning of a wide range of expressions from grammatical categories, such as tenses, to the lexical semantics of particular words, such as modal adverbials (*probably, necessarily*) and modal auxiliaries (*must, may, can*). The best known modalities are the alethic¹ modalities *necessary* and *possible*. Other modalities include temporal, deontic, epistemic, and doxastic ones, and modalities pertaining to disposition, ability, provability, mood, aspect, and so on. Temporal modalities deal with time. Deontic² modalities deal with obligation and permission. Epistemic³ modalities deal with knowledge, and doxastic⁴ modalities deal with belief. Although modals have been thoroughly studied since Aristotle, the formal theory of modality was revolutionized in the 1960's with the introduction of the possible world semantics by Hintikka and Kripke.

In our formal approach to the semantic analysis of modal expressions, we will employ the model-theoretic apparatus of modal logic. In modal logic, modals correspond to sentential operators, whose semantic role is to qualify the truth of the sentences in their scope. For example, the sentence *He is possibly right* is represented by a formula that is roughly equivalent to *It is possibly true that he is right*. Likewise, the sentence *Nature must obey necessity* is represented as *It is necessarily true that nature obeys necessity*. Given a sentential symbol p , we use the square notation $\Box p$ for the statement *p is necessarily true*, and the diamond notation $\Diamond p$ for the statement *p is possibly true*. Thus, *necessarily* and *possibly* are formalized as modal operators that act on sentences. Most modal expressions, though by no means all, are treated similarly.

The modal operators \Box and \Diamond are interdefinable. Even Aristotle in his *De Interpretatione* established that the negation of *It is necessary that p* is *It is not possible that not-p*. Similarly, *It is possible that p* is equivalent to *It is*

not necessary that not-p. It is contingent that p is equivalent to *It is possible that p, but not necessary that p*.

We will formalize this as follows. Let the symbol \neg be used for sentential negation. That is, $\neg p$ stands for *It is not the case that p*. The *Law of Double Negation* implies that the negation of this sentence is equivalent to p . Formally:

$$(1) \quad \neg\neg\varphi \Leftrightarrow \varphi,$$

where \Leftrightarrow is the symbol for equivalence and φ is a variable over arbitrary sentences. While \Box and \Diamond are unary sentential operators (combining with single sentences to produce new sentences), \Leftrightarrow is a binary relational symbol that is not part of the standard language of modal logic, but rather part of the metalanguage – i.e., used to make statements *about* the sentences of the formal language. If we have $\varphi \Leftrightarrow \psi$, then we also have $\neg\varphi \Leftrightarrow \neg\psi$. The statement *It is not necessary that p* is equivalent to *It is possible that not-p*. Thus:

$$(2) \quad \neg\Box p \Leftrightarrow \Diamond\neg p, \text{ and so}$$

$$(3) \quad \neg\neg\Box p \Leftrightarrow \neg\Diamond\neg p$$

Therefore, by the Law of Double Negation, we could define \Box in terms of \Diamond by:

$$(4) \quad \Box p \Leftrightarrow \neg\Diamond\neg p$$

Similarly, we could define \Diamond in terms of \Box and \neg :

$$(5) \quad \Diamond p \Leftrightarrow \neg\Box\neg p$$

The above are general properties that modal operators have under all interpretations. Within particular modalities (deontic, epistemic, etc.) the same operators take on different flavors and support additional inference patterns, which are quite distinct from case to case. In deontic logic, $\Box p$ usually stands for *It must be the case/is required, that p*. The sentence $\Diamond p$ is equivalent to $\neg\Box\neg p$, which means *It is not required that not-p*, or equivalently, *It may be the case/is permissible, that p*. The sentences *It is required that p* and *It is required that not-p* cannot both be true – they are *contraries* in the square of negation (cf. Horn 2001). On the other hand, *It is permissible that*

p and *It is permissible that not- p* cannot both be false – they are *subcontraries*. There is also general agreement that in deontic logic, we have that $\Box p \Rightarrow \Diamond p$ (if p is required, it is allowed), but not $\Box p \Rightarrow p$ (just because p is required, it does not follow that it is true).

In epistemic logic, $\Box p$ stands for *It is known that p* , where *known* here has to mean ‘known to someone.’ Let us assume that the possessor of the knowledge is the speaker. This is one common interpretation, although others are certainly possible. Thus $\Box p$ is interpreted as *The speaker knows that p* . Accordingly, the dual sentence, $\Diamond p$, is equivalent to *It is not true that the speaker knows that not- p* , which is equivalent to *It is possible, for all the speaker knows, that p* . Similarly, in doxastic logic, $\Box p$ stands for *The speaker believes that p* . Thus the dual $\Diamond p$ is equivalent to *The speaker does not believe that not- p* , which is again equivalent to *It is compatible with the speaker’s beliefs that not- p* . Both in epistemic and doxastic logic, the interpretation of $\Box p$ relative to some particular individual i is sometimes indicated explicitly by an index on the operator. Thus, for instance, $\Box_i p$ stands for *i knows that p* or *i believes that p* .

1.1. Propositional modal logic: syntactic approach

For now, we will limit our discussion to sentences, ignoring the words and phrases they are composed of. Formally, we will use a language $\mathcal{L}_{\mathcal{A}}$ based on a set \mathcal{A} of atomic propositional letters p, q, r, \dots , the sentential connectives $\wedge, \vee, \rightarrow$ and \neg for conjunction (*and*), disjunction (*or*), material conditional (*if-then*), and negation (*it is not the case that*), respectively, parentheses (shown only where needed to avoid ambiguity), and modal operators \Box and \Diamond . Using Greek letters, such as φ and ψ , as variables over strings built from these symbols, we define the language $\mathcal{L}_{\mathcal{A}}$ as follows:

- (6) a. All atomic propositional letters are sentences in $\mathcal{L}_{\mathcal{A}}$:
 $\mathcal{A} \subseteq \mathcal{L}_{\mathcal{A}}$
- b. $\mathcal{L}_{\mathcal{A}}$ is closed under the truth-functional connectives:
 If $\varphi, \psi \in \mathcal{L}_{\mathcal{A}}$, then $\varphi \wedge \psi, \varphi \vee \psi, \varphi \rightarrow \psi, \neg \varphi \in \mathcal{L}_{\mathcal{A}}$
- c. $\mathcal{L}_{\mathcal{A}}$ is closed under the unary modal operators:
 If $\varphi \in \mathcal{L}_{\mathcal{A}}$, then $\Box \varphi, \Diamond \varphi \in \mathcal{L}_{\mathcal{A}}$
- d. Nothing else is in $\mathcal{L}_{\mathcal{A}}$.

The standard language of (non-modal) propositional logic is obtained by omitting clause (6c).

Classical propositional logic can be axiomatized by several axiom schemata⁵ and one rule of inference, *Modus Ponens*:

$$(MP) \frac{\varphi, \varphi \rightarrow \psi}{\psi}$$

This rule asserts that from φ and $\varphi \rightarrow \psi$ as hypotheses, we can derive ψ as a conclusion. Equivalently, if φ and $\varphi \rightarrow \psi$ are provable, then ψ is provable as well. The following set of axioms is one way to characterize classical logic, jointly with the rule of Modus Ponens.

- (7) a. $\varphi \rightarrow (\theta \rightarrow \varphi)$
(truth follows from anything: if φ is true, then so is $(\theta \rightarrow \varphi)$ for any θ).
- b. $(\varphi \rightarrow (\theta \rightarrow \psi)) \rightarrow ((\varphi \rightarrow \theta) \rightarrow (\varphi \rightarrow \psi))$
(distributivity of implication: if $\theta \rightarrow \psi$ follows from φ , then $\varphi \rightarrow \psi$ follows from $\varphi \rightarrow \theta$).
- c. $(\neg\psi \rightarrow \neg\varphi) \rightarrow ((\neg\psi \rightarrow \varphi) \rightarrow \psi)$
(proof by contradiction: if the falsehood of φ follows from the falsehood of ψ , then showing that the truth of φ follows from the falsehood of ψ establishes the truth of ψ).

We now define the following notions, which are at the center of the syntactic approach to modal logic, as well as other axiomatic systems. The key notion is that of a *derivation*: a finite sequence of sentences, each of which is either an axiom or obtained from axioms and sentences already in the sequence by applying an inference rule of the system (Modus Ponens in our case). In addition, a derivation from *hypotheses* allows the use of hypotheses in the derivation sequence. Although the set of hypotheses can be infinite, every derivation must be finite. A *derived rule* is a rule whose conclusion has a derivation from its hypotheses. A *proof* for a sentence φ is a derivation sequence whose last member is φ . A *theorem* is a sentence that has a proof. Hence axioms are (trivially) theorems.

In addition to the above axioms, all propositional modal systems have the following axiom schema:

- (K) $\Box(\varphi \rightarrow \psi) \rightarrow (\Box\varphi \rightarrow \Box\psi)$
 (distributivity of \Box over \rightarrow : if ψ necessarily follows from φ , then the necessity of ψ follows from the necessity of φ).

Furthermore, all propositional modal systems have the additional inference rule, the *Necessitation* rule (introduced by Gödel):

$$(N) \frac{\varphi}{\Box\varphi}$$

This rule says that if φ is provable, then φ is necessarily true. It does *not* say that φ implies $\Box\varphi$. The sentence $\varphi \rightarrow \Box\varphi$ is not a theorem in all systems of modal logic.

The most basic propositional modal system is called *K* (named after Kripke). It contains, in addition to the axioms of classical propositional logic, only the axiom (K) and, in addition to Modus Ponens, the Necessitation rule.

So, for example, in system K, we can derive $\Box p \rightarrow \Box q$ from $p \rightarrow q$. Here is a derivation sequence beginning with the hypothesis and ending with the conclusion:

- | | | |
|-------|---|---------------------------------------|
| (i) | $p \rightarrow q$ | (hypothesis) |
| (ii) | $\Box(p \rightarrow q) \rightarrow (\Box p \rightarrow \Box q)$ | (axiom (K)) |
| (iii) | $\Box(p \rightarrow q)$ | (from (i) by Necessitation) |
| (iv) | $\Box p \rightarrow \Box q$ | (from (ii) and (iii) by Modus Ponens) |

Based upon the central notion of derivation, we define the following notions, which are at the heart of all logical systems:

- (8) A set Φ of sentences is syntactically *consistent* if and only if there is no derivation of a sentence of the form $\varphi \wedge \neg\varphi$ (a contradiction) from Φ .

Furthermore, we say that a sentence φ is syntactically consistent with a set Φ of sentences iff $\Phi \cup \{\varphi\}$ is consistent (i.e., φ can be added to Φ consistently). A consistent set Φ is *maximally consistent* if no sentence outside Φ can be added to Φ consistently. Every consistent set of sentences can be extended to a maximally consistent one.

- (9) A sentence φ is a syntactic *consequence* of a set Φ of sentences if and only if there is a derivation of φ from Φ .

1.2. Propositional modal logic: semantic approach

The central notion in the semantic interpretation of modal logic is that of *possible worlds*. The history of this concept can be traced back to Leibniz, who believed that we live in a world which is one of infinitely many possible worlds created by God (and, fortunately for us, the best one among them!). In philosophical logic, the notion was introduced only in the 1960's, independently by Hintikka (1961) and Kripke (1963). Its ontological status continues to be controversial among philosophers. These debates will not concern us here, however. The utility of possible worlds as a methodological tool in semantic analysis has been amply demonstrated in recent decades and does not depend on any particular stance on metaphysical questions, such as whether worlds other than ours “exist” in any real sense. For our purposes, they are nothing but abstract entities which help us in modeling certain semantic relations among linguistic expressions.

Towards this end, we only need to assume that possible worlds fix the denotations of the relevant expressions – truth values for sentences, properties for verb phrases, and so on. For now we will limit the discussion to sentences, and we will continue to use the formal language we introduced.

Possible worlds play a central role in defining the denotations of sentences. The meaning of a sentence is analyzed in terms of its role in distinguishing between possible worlds. Assuming that every sentence is guaranteed to be either true or false (the *Law of Excluded Middle*),⁶ but not both true and false (the *Law of Non-Contradiction*), each possible world determines the truth values of all atomic sentences, as well as, via the interpretation of the logical connectives, those of their Boolean combinations.

To be more precise, we define a *model* as a pair $M = \langle W, V \rangle$, consisting of a nonempty set W (the set of possible worlds) and a function V which, for each world w in W , assigns truth values to the atomic sentences in the language. We write 1 and 0 for the values ‘true’ and ‘false,’ respectively. The truth values of complex sentences are defined recursively by the following clauses.⁷

(10) For atomic p , $V_w(p)$ is either 1 or 0.

$$V_w(\neg \varphi) = \begin{cases} 1 & \text{if } V_w(\varphi) = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$V_w(\varphi \wedge \psi) = \begin{cases} 1 & \text{if } V_w(\varphi) = 1 \text{ and } V_w(\psi) = 1 \\ 0 & \text{otherwise} \end{cases}$$

$$V_w(\varphi \vee \psi) = \begin{cases} 1 & \text{if } V_w(\varphi) = 1 \text{ or } V_w(\psi) = 1 \\ 0 & \text{otherwise} \end{cases}$$

$$V_w(\varphi \rightarrow \psi) = \begin{cases} 1 & \text{if } V_w(\varphi) = 0 \text{ or } V_w(\psi) = 1 \\ 0 & \text{otherwise} \end{cases}$$

Given the assignment function, each sentence in the language distinguishes between those worlds in which it is true and those in which it is false. We may, therefore, associate each sentence φ with the set of those worlds in which it is true. Thus we introduce a function $\llbracket \cdot \rrbracket^M$ for our model M , mapping sentences to sets of worlds:

$$(11) \llbracket \varphi \rrbracket^M =_{\text{def}} \{w \in W \mid V_w(\varphi) = 1\}$$

We call $\llbracket \varphi \rrbracket^M$ the *denotation* of sentence φ in M . We will omit the superscript M when the choice of model makes no difference. The term *proposition*, in its technical use in this context, is reserved for sets of worlds. Hence the denotation of the sentence *It is raining* is the proposition ‘that it is raining,’ the set of just those worlds in which it is raining.

We can now characterize standard logical properties of sentences set-theoretically in terms of the propositions they denote:

- (12) A sentence φ is:
- a. *tautologous* iff $W \subseteq \llbracket \varphi \rrbracket$
 - b. *contradictory* iff $\llbracket \varphi \rrbracket \subseteq \emptyset$
 - c. *contingent* otherwise (i.e., iff $\emptyset \subset \llbracket \varphi \rrbracket \subset W$)

Likewise, semantic relationships between sentences can be defined in terms of the propositions they denote. For instance, the interpretation of the logical connectives gives rise to the following relations between the denotations of complex sentences and the parts they are composed of (\cap , \cup , and \setminus stand for intersection, union, and relative complement or set subtraction, respectively):

$$\begin{aligned}
(13) \quad & \llbracket \neg\varphi \rrbracket = W \setminus \llbracket \varphi \rrbracket \\
& \llbracket \varphi \wedge \psi \rrbracket = \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket \\
& \llbracket \varphi \vee \psi \rrbracket = \llbracket \varphi \rrbracket \cup \llbracket \psi \rrbracket \\
& \llbracket \varphi \rightarrow \psi \rrbracket = W \setminus (\llbracket \varphi \rrbracket \setminus \llbracket \psi \rrbracket) = (W \setminus \llbracket \varphi \rrbracket) \cup \llbracket \psi \rrbracket
\end{aligned}$$

The clause for conjunction states that the denotation of *It is raining and it is cold* is the intersection of the denotations of *It is raining* and *It is cold*, i.e., the set of worlds in which both of these sentences are true. This idea can be generalized to arbitrary sets Φ of propositions, writing $\llbracket \Phi \rrbracket$ for the set of worlds in which all sentences in Φ are true – i.e., the intersection of their respective denotations:⁸

$$(14) \quad \llbracket \Phi \rrbracket^M =_{\text{def}} \bigcap \{ \llbracket \varphi \rrbracket^M \mid \varphi \in \Phi \}$$

Using this notation, we can now define the central notions of modal logic semantically.

(15) Consistency

- a. A set Φ of sentences is semantically *consistent* iff there is some world in which all sentences in Φ are true; i.e., iff $\llbracket \Phi \rrbracket \neq \emptyset$.
- b. A sentence φ is semantically consistent with a set Φ of sentences iff $\Phi \cup \{\varphi\}$ is consistent (φ can be added to Φ consistently); i.e., iff $\llbracket \Phi \cup \{\varphi\} \rrbracket \neq \emptyset$.

(16) Consequence

A sentence φ is a semantic *consequence* of a set Φ of sentences if and only if φ is true in all possible worlds in which all sentences in Φ are true; i.e., iff $\llbracket \Phi \rrbracket \subseteq \llbracket \varphi \rrbracket$.

We often use these terms relative to Φ consisting of just a single sentence. In a slight abuse of terminology, we will speak of a sentence being consistent with, or a consequence of, such a single proposition, rather than the singleton set containing it. This is a safe move to make since $\llbracket \{\psi\} \rrbracket = \llbracket \psi \rrbracket$.

Much of the work in modal logic explores the relationship between axiomatic systems and the semantic relationships they give rise to. For instance, it is easy to check that all the axioms in (7) above are tautologies, and that Modus Ponens preserves the property of being a tautology. Hence all theorems in this system are tautologies. This property of a formal system

is called *soundness*. What is remarkable is that the converse also holds. Every tautology can be derived as a theorem in this system, i.e., the classical propositional logic is *complete* with respect to this axiomatic system (for a proof by Kalmár see Mendelson 1987). The completeness theorem establishes that being provable in a sound formal system is equivalent to being true in every model. As a corollary, it follows that any given set Φ of sentences is syntactically consistent if and only if Φ has a model (i.e., is semantically consistent).

By adding more axioms, we obtain systems which may be sound and complete with respect to some restricted class of models, rather than all models. We will see a few illustrations of this in later sections. In general, the task of identifying the class of models relative to which a given axiomatic system is sound and complete can be difficult. We will not discuss it in more detail in this chapter.

2. Modal bases

Above, we defined the notions of consistency and consequence relative to the proposition $[\Phi]$, the set of those worlds at which all sentences in Φ are true. Following Kratzer (1977, 1981, 1991) we call this set of worlds the *modal base*. Kratzer calls the set of propositions whose intersection forms the modal base *conversational background*. In this section we give a more direct definition of modal bases, sidestepping conversational backgrounds.

Modal bases constitute a central parameter in all formal treatments of modality. They serve to distinguish between the different readings of modals (epistemic, deontic, etc.) we identified earlier. Thus it is customary to speak of epistemic, deontic, and other types of modal bases. The intuition is that on each occasion of use the modal base is the set of just those worlds that are compatible with all of the speaker's beliefs or desires, the applicable laws, and so on. The contents of speakers' beliefs, of laws, etc. are of course themselves contingent and, therefore, may vary from world to world. Thus the contents of, say, an epistemic modal base do not remain constant across different worlds. Just like the truth values of the sentences in our language, the modal base depends on the world of evaluation.

Formally, a modal base is given by a function from possible worlds to propositions. We will reserve the letter R as a general symbol for such functions, using superscripts to distinguish between them, such as R^{epist} and R^{deont} for epistemic and deontic modal bases, respectively. Given a world w

of evaluation, we will write R_w^{epist} for the (speaker's) epistemic modal base at w , i.e., the set of exactly those worlds that are compatible with the speaker's beliefs in w . We will use the Greek letter ρ as a variable over modal bases.

The second major parameter in the interpretation of modality is the *modal force*. Unlike the modal base, which is usually left implicit and contextually given, the modal force is an integral part of the lexical meaning of all modals. For instance, both (17b) and (17c) below can be used as assertions about either the speaker's beliefs (*given what I know...*) or John's options and obligations (*given the rules and John's age...*), but the logical relations invoked in the two sentences are unambiguously possibility and necessity, respectively. With respect to the intended modal base, this difference between (17b) and (17c) corresponds to consistency and consequence.

- (17) a. John is at the party.
 b. John may be at the party.
 c. John must be at the party.

We can now give the first version of the truth conditions for modalized sentences such as (17b,c). We assume that in both cases the modal auxiliary functions as a sentential operator which takes (17a) as its argument. For simplicity, we assume that the intended modal base ρ is given, and extend the valuation function V as follows:

$$(18) \quad V_w(\diamond_\rho \varphi) = \begin{cases} 1 & \text{if } V_{w'}(\varphi) = 1 \text{ for some } w' \in \rho_w \\ 0 & \text{otherwise} \end{cases}$$

$$(19) \quad V_w(\square_\rho \varphi) = \begin{cases} 1 & \text{if } V_{w'}(\varphi) = 1 \text{ for all } w' \in \rho_w \\ 0 & \text{otherwise} \end{cases}$$

Modal operators are thus quantifiers over possible worlds, with the modal base providing the domain of quantification. Relative to modal base ρ , \diamond_ρ and \square_ρ are duals.

We can now identify the meaning of necessity modals with the interpretation of \square , and that of possibility modals with the interpretation of \diamond . The truth-conditional content of a modal is given relative to a context of use that fixes the modal base, that is, the appropriate index ρ for the operators \square and \diamond . For instance, supposing the modal base of an utterance of (17b) is determined to be epistemic, (17b) is true in a world w iff there exists $w' \in R_w^{epist}$

such that *John is at the party* is true in w' . If instead the modal base of an utterance of (17b) is determined to be deontic, the only thing that would change in determining whether (17b) is true in w is the domain of quantification. In general, R^{epist} and R^{deont} are distinct so an utterance of (17b) expresses a different proposition, depending on whether the modal is construed with an epistemic or a deontic modal base.

As we mentioned earlier, our definition of modal bases is a shortcut compared to Kratzer's treatment. There the modal base is defined indirectly via a conversational background, the latter being a function from worlds to sets of propositions. Our more direct definition is not meant to deny the utility of the notion of conversational background (the examples to which we apply it here do not illustrate its specific advantages), nor is it merely for the sake of simplicity. Rather, it takes us to the area of contact between the Kratzer-style theory and another, equally influential approach in terms of *accessibility relations* between possible worlds. For any function ρ from worlds to sets of worlds, there is a relation R_ρ which pairs up each world w with all and only the worlds in ρ_w :

$$(20) R_\rho =_{\text{def}} \{\langle w, w' \rangle \mid w' \in \rho_w\}$$

Many authors, especially in the area of philosophical logic, take such relations between possible worlds as basic, rather than define them in terms of modal bases. In technical discussions of modal logic, they are given as part of the model. Assuming for simplicity that we are only interested in one accessibility relation R , we can define a model to be a triple $\langle W, R, V \rangle$, where W and V are as before. The pair $\langle W, R \rangle$, the set of worlds and the accessibility relation, is called a *frame*. For a fixed frame, certain axioms may be guaranteed to hold solely in virtue of the properties of the accessibility relation, regardless of what truth values V assigns to the sentences of the language at individual worlds. We will discuss some of them, and the structural properties they correspond to, in the following subsection. In the course of the discussion, it will be useful to switch back and forth between modal bases as functions from worlds to propositions and as accessibility relations. In light of the close formal affinity between these two perspectives, we will use the same symbol R for both.

2.1. Properties of modal bases

The modal base is not only a useful parameter to capture the variability and context dependence of modal expressions. It is also the right place to state generalizations about the properties of particular readings of modals. The investigation of such general properties and their logical consequences is the topic of modal logic; the reader is referred to introductions into this field, such as Hughes and Cresswell (1996), for discussions going far beyond what we can cover here. We will only mention a few such conditions which have turned out important in linguistic analysis. For concreteness, unless otherwise stated, we will restrict the discussion in this subsection to the special case of a *doxastic* modal base – i.e., one that represents the speaker’s beliefs. We will write R instead of R^{dox} for simplicity.

Consistency. The first general condition one can impose on modal bases is that they be *consistent*. Formally, this corresponds to the requirement that for all worlds w , the set R_w be nonempty. Alternatively, in relational terms, the requirement is that R be *serial*: For every world w of evaluation, there is at least one world w' that is accessible from w via R .

In linguistic theory, consistency is generally taken to be a requirement for all modal bases. An inconsistent modal base may result in the interpretation of sentences by quantification over an empty set of worlds, leading to presupposition failure.

Realism. Next, one may impose the condition that a modal base be *realistic*, in the sense that none of the sentences supported by the modal base for a world w (i.e., true at all worlds in R_w) are false at w . Formally, this condition means that each world w must itself be a member of R_w . In relational terms, the analog of this condition is that R is reflexive – i.e., each world w is related to itself by R .

Realism is a sensible condition for some modalities, but not for others. For doxastic modality, the condition means that all of the speaker’s beliefs are true. This property is often taken to distinguish *knowledge* from *belief* or, equivalently, *epistemic* from *doxastic* accessibility relations. Deontic modality is generally not realistic. Assuming otherwise would amount to the claim that all obligations are fulfilled.⁹

Introspection. Under a doxastic interpretation, R_w is the set of those worlds that are compatible with what the speaker believes at w . These worlds will differ from each other with respect to the truth values of non-modal sen-

tences, such as *It is raining*. In addition, they may also differ with respect to the speaker's beliefs. Each world w' in R_w is one in which the speaker has certain beliefs, represented in the set $R_{w'}$. What this means for w is that the modal base does not only encode the speaker's beliefs about the facts, but also her beliefs about her own beliefs.

Little reflection is required to see that some conditions ought to be imposed on these beliefs about one's own beliefs. Suppose the speaker believes in w that it is raining (i.e., all worlds in R_w are such that it is raining), and suppose further that there are two worlds w', w'' in R_w , such that in w' the speaker believes that it is raining, and in w'' she believes that it is not. While there is nothing wrong with this scenario from a formal point of view, it is very peculiar indeed as a representation of a speaker's beliefs. It amounts to the claim that the speaker has a very definite opinion on the question of whether it is raining, but does not know what that opinion is.

Such outcomes are avoided by imposing constraints governing the relationship between speakers' actual beliefs and the beliefs they have at the worlds compatible with their beliefs, i.e., the worlds in the corresponding modal bases. The two most commonly encountered conditions of this kind are those of positive and negative introspection.

Positive introspection is the requirement that at each world compatible with what the speaker believes (i.e., each world in R_w), she has all the beliefs that she actually has at w (and possibly more). Formally, this means that for each such belief-world w' , the speaker's doxastic modal base $R_{w'}$ is a subset of the actual modal base R_w . The corresponding condition on the accessibility relation is that it be *transitive*. The following two statements are equivalent, and each imposes the requirement of positive introspection, provided it is true at every world w :

- (21) a. For all $w' \in R_w$, we have $R_{w'} \subseteq R_w$
 b. If wRw' and $w'Rw''$, then wRw''

It is helpful to visualize this formal constraint by illustrating the kind of case it rules out. Such a case is shown in Figure 1. The modal bases R_w and $R_{w'}$ are indicated as partial spheres (whether the worlds are in their respective modal bases, and whether there is any overlap between R_w and $R_{w'}$ is not relevant here). Positive introspection fails because $R_{w'}$ is not fully contained in R_w : There is at least one world, w'' , which is in $R_{w'}$, but not in R_w . Equivalently, in terms of the accessibility relation, there is a path leading from w to w'' , but w'' is not directly accessible from w . Now, suppose some sen-

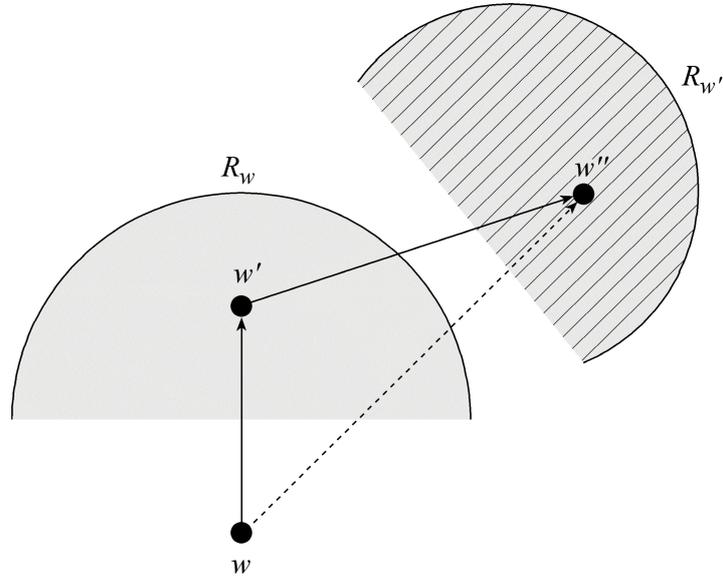


Figure 1. A violation of transitivity. No direct accessibility link leads from w to w'' ; $R_{w'}$ is not a subset of R_w .

tence φ is true at all worlds in R_w , but false at w'' . Thus there is a sentence, φ , which the speaker believes at w but not at w' . Since w' is compatible with what the speaker believes at w , the picture shows a speaker who at w thinks she may not believe that φ , even though she actually believes φ at w . To rule out such cases, positive introspection is usually imposed as a condition on epistemic and doxastic modal bases.

Negative introspection is the corresponding requirement that there be no world compatible with what the speaker believes at which she holds any beliefs that she does not actually hold. In relational terms, this requires R to be *euclidean*. Formally:

- (22) a. For all $w' \in R_w$, we have $R_w \subseteq R_{w'}$
 b. If wRw' and wRw'' , then $w'Rw''$

As before, we can illustrate the effect of this condition by giving an example in which it is violated, as illustrated in Figure 2:

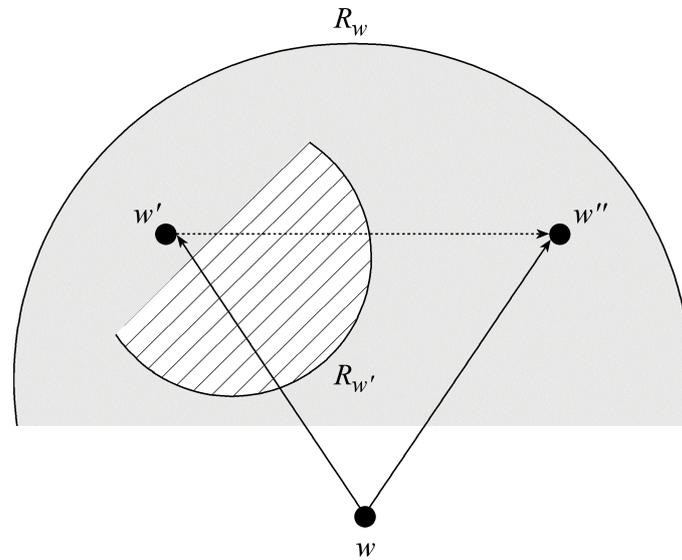


Figure 2. A violation of euclidity. No direct accessibility leads from w' to w'' (or vice versa); R_w is not a subset of $R_{w'}$.

The accessibility relation in Figure 2 is not euclidean, and the modal base R_w lacks the property of negative introspection. Even though w' and w'' are both accessible from w , the world w'' is not directly accessible from w' . Suppose φ is true at w'' and false at worlds in R_w . Then the situation depicted is one in which the speaker does not believe that φ is false at w , but does believe that it is false at w' . In other words, the speaker thinks she may believe that φ is false, but is not sure if she does. To rule out such counter-intuitive models, negative introspection (i.e., euclidity) is usually required of epistemic and doxastic modal bases.

Table 1 summarizes these properties of modal bases and the corresponding properties of accessibility relations, along with the axioms of modal logic that are guaranteed to hold for any modal base with the respective properties.

Table 1. Correspondences between some properties of modal bases and accessibility relations, and their characteristic axioms. (Free variables w , w' , w'' are universally quantified over.)

<i>Modal Base</i>	<i>Accessibility Relation</i>	<i>Axiom</i>	
consistency $R_w \neq \emptyset$	seriality $(\exists w') wRw'$	$\Box_R \varphi \rightarrow \Diamond_R \varphi$	(D)
realism $w \in R_w$	reflexivity wRw	$\Box_R \varphi \rightarrow \varphi$	(T)
total realism $R_w = \{w\}$	identity $wRw' \Leftrightarrow w = w'$	$\Box_R \varphi \leftrightarrow \varphi$	
positive introspection $w' \in R_w \Rightarrow R_w \subseteq R_{w'}$	transitivity $wRw' \wedge w'Rw'' \Rightarrow wRw''$	$\Box_R \varphi \rightarrow \Box_R \Box_R \varphi$	(4)
negative introspection $w' \in R_w \Rightarrow R_w \subseteq R_{w'}$	euclidity $wRw' \wedge wRw'' \Rightarrow w'Rw''$	$\Diamond_R \varphi \rightarrow \Box_R \Diamond_R \varphi$	(5)

In most applications, it is some combination of these conditions that in its totality determines the properties of a particular modal base and distinguishes it from others. For instance, *doxastic* modal bases are generally taken to be consistent and fully introspective, whereas *epistemic* ones (modeling knowledge, rather than mere belief) are, in addition to these properties, realistic. Positive or negative introspection are rarely imposed on modal bases other than epistemic or doxastic ones.

2.2. Ordering sources

Our discussion so far has shown how the great variability in the readings of modal expressions can be reduced to the specification of their modal force in their lexical meaning and the contextual determination of a modal base on a particular occasion of use. Natural language though abounds in modal expressions that exhibit a wider variety of modal forces than plain necessity and possibility, appropriately relativized to a modal base. Expressions of graded modality, such as *may well*, *barely possible*, and *more likely than*, for instance, show that our semantics had better incorporate a gradable notion of possibility, built around the absolute notion of consistency.

As argued by Kratzer, the interpretation of modal expressions requires an additional parameter, which she terms *ordering source*. This parameter enables us to specify a wider range of modal forces, such as weaker necessity than in (19) and stronger possibility than in (18), as well as to cover expressions of gradable modality. In this subsection we motivate ordering sources by discussing first the need for weaker necessity in the epistemic domain and then the use of ordering sources in avoiding potential inconsistencies.

Consider the necessity modals *should* and *must* in (23b) and (24b) below, supposing they are construed with an epistemic (hence realistic) modal base. If they quantified universally over R_w^{epist} , for any w , then (23b) and (24b) would entail (23a) and (24a), respectively, but no such entailment is in fact present:

- (23) a. John has reached Athens by now.
 b. John should have reached Athens by now.
- (24) a. It rained overnight.
 b. It must have rained overnight.

Rather, (23b) and (24b) carry an implication of uncertainty that (23a) and (24a) do not. A speaker who chooses to assert (23b) or (24b), over (23a) or (24a), indicates that she does not have direct knowledge about John's whereabouts, or the reason for the wet ground, but can only infer where John is, or what explains the wet ground, based on other facts she knows and certain reasonable assumptions. Conversely, a speaker who has direct knowledge of John's whereabouts or witnessed the rain directly would choose (23a) and (24a), over (23b) or (24b).

The conclusion to be drawn is that the universal quantification of epistemic necessity modals needs to be further relativized. The sentences (23b) and (24b) claim that (23a) and (24a) are true *provided* some additional assumptions are brought to bear. In the case of (23b), such an implicit assumption can be that John's trip is following the planned itinerary; in the case of (24b), that rain is the most plausible explanation for the overt evidence at hand, the wet ground.

The meaning of epistemic necessity modals thus makes reference to a set of contextually determined assumptions, construed as a set of propositions. This set of propositions, which for epistemic modals must be nonempty, is the ordering source parameter. By using an epistemic necessity modal, a speaker is signaling uncertainty as to whether these implicit assumptions are true, that is, whether the actual world is one conforming to the planned

course of events, or whether the actual world is one to which the most plausible explanation for the evidence at hand applies.

Next we consider potential inconsistencies arising when facts of the matter are not appropriately distinguished from rules, laws, aims, desires, etc. Suppose the sentences in (25) below are uttered in a situation in which John was caught speeding and the law both prohibits speeding and requires that anyone caught speeding pays a fine. In such a situation, with the modals construed deontically, (25a) is true and (25b) false:

- (25) a. John must pay a fine for speeding.
 b. John need not pay any fine for speeding.

So far, we have identified deontic modals as those whose modal base is deontic, a mapping from any world w to the set of those worlds in which the dictates of the law in w are adhered to. This leads to two interrelated problems arising from the fact that the proposition that John was speeding has to be true in all worlds in the modal base. This has to be the case if (25a) is to come out true. The first problem is an unavoidable inconsistency. Given the prohibition against speeding and John's speeding, the modal base ends up being inconsistent and, in that case, both (25a) and (25b) are verified. The second problem is an unwanted entailment of (25a), namely (26).¹⁰

- (26) John was required to be speeding.

In order for these two problems to be avoided, the propositions corresponding to the relevant actual facts and the propositions corresponding to the content of the law must constitute distinct parameters. The modal base of deontic modals is not itself deontic, but determined by the relevant actual facts, for instance, that John was speeding. The deontic parameter in the interpretation of deontic modals is the ordering source, determined by the contents of the law in every world. Different mappings from a world w to a set of propositions determine a moral ideal, a normal course of events, a likely scenario, etc., for w . Moreover, each such mapping can be used to rank the worlds in ρ_w . For instance, the relevant ranking for the interpretation of (25a) is one where worlds in which what is prohibited in w is never committed are ranked higher than worlds in which violations occur; worlds in which violations occur but are punished are in turn ranked higher than worlds with violations but no punishment.

More formally, we say that, for any world w , a set of propositions \mathcal{P} induces a *preorder*¹¹ \leq_w on a set of worlds as follows: $w'' \leq_w w'$ if and only if any proposition in \mathcal{P} that is true at w' is also true at w'' . We can also define ordering sources more directly, just as we did for modal bases, as functions from worlds to preorders between worlds. We will follow this more direct method here and use o as a variable over ordering sources. For any function o from worlds to preorders, there is a ternary relation relating each world w with the members of all and only the pairs in o_w :

$$(27) \leq_o =_{\text{def}} \{ \langle w, w', w'' \rangle \mid \langle w', w'' \rangle \in o_w \}$$

We can now give the semantics of a doubly modal necessity operator, making reference to both modal bases and ordering sources:¹²

$$(28) V_w(\Box_{\rho, o} \varphi) = \begin{cases} 1 & \text{if for all } w' \in \rho_w \text{ there is some } w'' \in \rho_w \\ & \text{such that } \langle w'', w' \rangle \in o_w \text{ and} \\ & \text{for all } u \in \rho_w, V_u(\varphi) = 1 \text{ if } \langle u, w'' \rangle \in o_w \\ 0 & \text{otherwise} \end{cases}$$

The reader is referred to Kratzer (1981, 1991) for definitions of other doubly relative modal notions, such as ‘good possibility,’ ‘slight possibility,’ or ‘better possibility.’ Formally, the order induced by the ordering source is similar to the relation of comparative similarity between possible worlds, which is central to the Stalnaker/Lewis theory of counterfactuals (see Lewis (1981) for a comparison).

2.3. Interim summary

So far we outlined some of the basic apparatus in which formal accounts of modal expressions are usually framed. We saw that modals are treated as quantifiers over possible worlds, whose domains of quantification at any given world of evaluation are determined by accessibility relations with certain properties. These basic ingredients were introduced more than twenty years ago and have formed the mainstay of the formal semantics of modality ever since.

3. Modality and time

In this section, we turn to modal aspects of expressions of tense and temporal reference. The semantic interactions between the temporal domain and modality are so pervasive that no analysis of one can be complete without an account of the other. We will be focusing on ways in which temporal reference determines and constrains modal interpretations. As part of this discussion, we will further develop and refine the formal tools we have so far introduced.¹³

3.1. The English present tense

The present tense in English can be used with both present and future reference times. Thus, as far as temporal reference is concerned, the Present complements the Past. Consider the following examples:

- (29) a. Megan was in her office yesterday.
 b. Megan is in her office (now).
 c. Megan is in her office tomorrow.

Based on (29), it would appear that there are two tenses in English, Past and Non-Past, illustrated by (29a) on the one hand, and (29b,c), on the other. However, the formal semantic literature on tense in English presents a different picture. While there is general agreement that (29a,b) differ only in tense and temporal reference, most authors maintain that the futurate Present in (29c) differs rather dramatically from both. Specifically, for the truth of (29c) it is not sufficient that a certain state holds in the future, but this occurrence has to be *predetermined*, in some sense, at speech time. No such connotation is observed in either (29a) or (29b).

Notice that the sentences in (29) are *stative*. A slightly different pattern is observed with non-stative predicates:

- (30) a. Megan came to her office yesterday.
 b. Megan comes to her office (now).
 c. Megan comes to her office tomorrow.

Here, both (30b) and (30c) call for a *scheduling* reading. The contrast between (29b) and (30b) is accompanied by a difference in temporal reference: (29b) asserts that a state holds at speech time, whereas the event in (30b) is

asserted to occur in the (near) future. A full analysis of these differences would lead us far beyond the scope of this section and into the area of aspectual classes. We merely note that the reference time follows the speech time in all and only the sentences in (29c) and (30b,c). Our goal in this subsection is to show that, with the right model-theoretic setup, the presence or absence of the scheduling reading falls out from this fact about temporal reference alone, with a simple and uniform analysis of the tenses and no further stipulations. This is a modest goal, but first we must further develop the technical apparatus.

Some background. Like many interesting phenomena at the intersection of modality and temporality, the connection between temporal reference and the scheduling reading can be traced to a fundamental and, it appears, virtually universally shared background assumption speakers make about the nature of time: that there is an important difference between a *fixed* past and an *open* future. The past up to and including the present time has (now) no chance of being different from what it actually has been. Consequently, sentences whose truth or falsehood depends solely on times no later than the speech time are either unequivocally true or unequivocally false, regardless of whether their truth values are known or not.

The situation is different with respect to the future. To use a well-known example from Aristotle (*De Interpretatione* 9), consider the claim that there will be a sea battle tomorrow. It is intuitively clear that the truth or falsehood of this claim will, in time, be fully determined by the course of events. There either will or will not be a sea battle, and there can be no two ways about it. However, most speakers share with Aristotle the intuition that, regardless of which way history eventually settles the question, at present it is still possible for things to turn out otherwise. It is precisely with regard to this intrinsic non-determinacy that the future differs from past and present.

The rigorous treatment of this old idea in present-day formal logic starts with the work of Prior (1967). In our brief discussion of this topic, we will focus on two issues: The notion of truth in time, and temporal constraints on the interaction between two kinds of uncertainty: objective (metaphysical) and subjective (epistemic/doxastic).

Two notions of truth. In Aristotle's example, the statement that there will be a sea battle tomorrow and its negation, the statement that there will not be a sea battle, have a similar semantic status at the present time. Both have a certain chance of being true, and neither is verified or falsified by the facts accumulated through history up until now. Aristotle was concerned

with the implications of this judgment for the validity of the logical principles of Excluded Middle and Non-Contradiction. When applied to the sea battle example, these two general principles entail that one of the sentences is true and the other is false, contrary to intuition.

In formal two-valued logic, there are two ways of reconciling the intuition with the logical principles, depending on which of two notions of truth one considers appropriate for statements about the future. The first, known as the *Ockhamist* notion of truth, upholds Excluded Middle and Non-Contradiction for such statements and asserts that they are indeed *already* either true or false, even though it may be impossible, even in principle, to *know* their truth values ahead of time. This was Aristotle's own answer to the problem.

The second notion of truth, sometimes called *Peircean*, maintains, in contrast, that neither of the statements about tomorrow's sea battle is true until the facts actually settle the question. Until such time, both are false. Does this mean that Excluded Middle is not applicable to future statements under this view? Not necessarily, if the notion of truth for such statements is properly construed. The idea is to treat them as *modal* statements, involving the modal force of necessity. It is a straightforward consequence of the interpretation of the modal operators that relative to any given modal base ρ , two statements $\Box_{\rho} \varphi$ and $\Box_{\rho} \neg\varphi$ may both be false (or, for that matter, vacuously true) without any violation of logical principles.¹⁴

3.2. The temporal dimension

Let us consider in some detail a formal implementation in which these distinctions can be clarified. To integrate time into our model-theoretic apparatus, we add a temporal dimension to the possible worlds, which in the previous sections had no internal structure at all. Technically, we can achieve this by introducing an ordered set $(T, <)$ of times, where $<$ is the *earlier-than* relation. We assume that $<$ has the following properties for all $t, t', t'' \in T$:

- (31) a. *irreflexivity*: not $(t < t)$
 b. *transitivity*: if $t < t'$ and $t' < t''$, then $t < t''$
 c. *linearity*: $t < t'$ or $t' < t$ or $t = t'$

The inverse $<^{-1}$ of the earlier-than relation is the *later-than* relation, defined as:

$$(32) \quad t' <^{-1} t \text{ iff } t < t'.$$

The properties of $<$ ensure that its inverse is also reflexive, transitive and linear (these properties are ‘preserved under inverse’). We will refer to both relations below, writing $<$ for *earlier-than* and $>$ for *later-than*.

Each single world $w \in W$ is now associated with a set of pairs $\langle w, t \rangle$ for $t \in T$; temporal precedence is extended to these pairs as expected:

$$(33) \quad \langle w, t \rangle < \langle w', t' \rangle \text{ iff } w = w' \text{ and } t < t'$$

The same holds for $>$.¹⁵

Notice that these temporal relations are themselves modal accessibility relations – though ‘modal’ only in the technical sense of modal logic; we have not yet employed them in the interpretation of linguistic expressions of modality. However, we can now use the same formal tools as before to evaluate statements at individual world-time pairs in this structure, analogously to our definitions in (18)–(19) for single worlds. For instance, $>_{\langle w, t \rangle}$ is the set of all world-time pairs $\langle w, t' \rangle$ such that $t > t'$ – i.e., all world-time pairs that precede $\langle w, t \rangle$ in time. Similarly, we define the set $<_{\langle w, t \rangle}$. We also have the set $\leq_{\langle w, t \rangle}$, which differs from $<_{\langle w, t \rangle}$ only in that it includes $\langle w, t \rangle$. The sets $>_{\langle w, t \rangle}$ and $<_{\langle w, t \rangle}$, illustrated in Figure 3, are the analogs of modal bases in the temporal dimension (although the term *temporal base* has never been proposed for them):

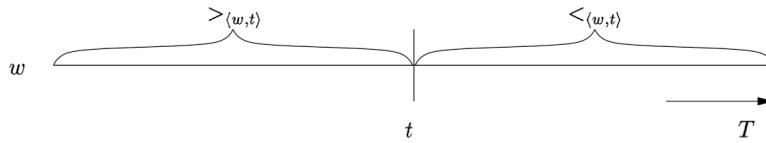


Figure 3. Temporal accessibility. The sets of world-time pairs preceding and following $\langle w, t \rangle$ are labeled ‘ $>_{\langle w, t \rangle}$ ’ and ‘ $<_{\langle w, t \rangle}$ ’, respectively.

Thus:

$$(34) \quad V_{\langle w, t \rangle}(\diamond_{>} \varphi) = \begin{cases} 1 & \text{if } V_{\langle w, t' \rangle}(\varphi) = 1 \text{ for some } \langle w, t' \rangle \in >_{\langle w, t \rangle} \\ 0 & \text{otherwise} \end{cases}$$

Under this definition, the expression $\diamond_{>} \varphi$ states that φ happened at some time in the past. Prior (1957, 1967) used the special symbol P for $\diamond_{>}$, and H for its dual $\square_{>}$, the latter meaning roughly *it has always been the case that...* Notice that the truth conditions for these operators make no reference to worlds other than the world w of evaluation.

There are various ways in which this basic idea can be extended to include present and future reference. For our purposes, the best way to proceed is to ignore for the moment the modal implications of future statements; we will account for them shortly. Regarding temporal reference on its own, we will say that statements about the future are evaluated just like those about the past, but with respect to the relation $<$. The by-now familiar truth definition for modal operators gives us the counterpart to (34) with non-past reference:

$$(35) \quad V_{\langle w, t \rangle}(\diamond_{<} \varphi) = \begin{cases} 1 & \text{if } V_{\langle w, t' \rangle}(\varphi) = 1 \text{ for some } \langle w, t' \rangle \in \langle_{<} \langle w, t \rangle \\ 0 & \text{otherwise} \end{cases}$$

Prior used the symbols F and G for $\diamond_{<}$ and $\square_{<}$, respectively.

In defining the truth conditions in (35), we have followed the Ockhamist tradition. Recall that, in this approach, the truth value of a statement such as $\diamond_{<} \varphi$ is determined at all times, depending solely on the truth values of φ at various times in the world of evaluation. Accordingly, our definition treats past and non-past symmetrically and leaves no room for uncertainty about the future.

We did *not* claim that expressions of the form $\diamond_{<} \varphi$ are suitable logical translations of well-formed English sentences such as *It rains tomorrow* or *It will rain tomorrow*. Before we turn to the interpretation of natural language expressions, let us continue our discussion of the formal apparatus and incorporate the asymmetry between past and future. For this purpose, we will add a modal dimension to our temporal model.

Modal/temporal two-dimensional semantics. Modal notions are introduced in our system as before, by postulating a multitude of possible worlds (or world-lines, as illustrated in Figure 3). We will assume here that all of these worlds are aligned with the same temporal dimension, given by $(T, <)$.¹⁶ One can picture these alternative worlds as lines running in parallel, illustrated in Figure 4:

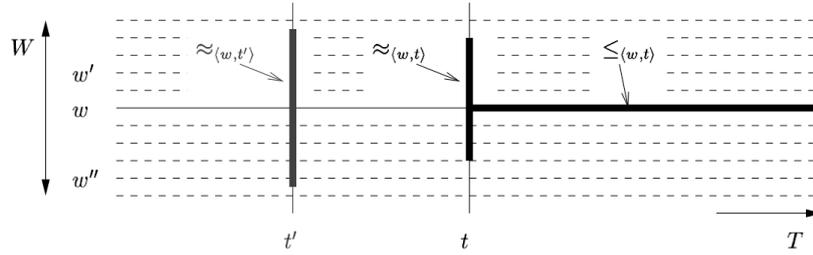


Figure 4. Two-dimensional modal logic. The thick lines represent the sets of indices accessible from $\langle w, t \rangle$ by the modal relation \approx (vertical) and the temporal relation \leq (horizontal). For example, the modal base $\approx_{\langle w', t' \rangle}$ contains $\approx_{\langle w, t \rangle}$, but not *vice versa*.

Modal bases are given by accessibility relations, as before. However, we are now interested in the way they change over time. Our current *tomorrow* is part of the *open* future from today's perspective, but will be part of the *fixed* past from next week's.

Thus modal accessibility relations, such as the temporal ones introduced before, will be time-sensitive, defined between world-time pairs, rather than just worlds. Expressions such as $\langle w, t \rangle \rho \langle w', t' \rangle$ or, alternatively, $\langle w', t' \rangle \in \rho_{\langle w, t \rangle}$ state that $\langle w', t' \rangle$ is accessible from $\langle w, t \rangle$. We will say that *modal* accessibility relations are those which link only world-time pairs whose time coordinate is constant (i.e., if $\langle w, t \rangle \rho \langle w', t' \rangle$, then $t = t'$). *Temporal* relations, on the other hand, link only world-time pairs whose world coordinate is constant (i.e., if $\langle w, t \rangle \rho \langle w', t' \rangle$, then $w = w'$). In Figure 4, modal and temporal relations operate “vertically” and “horizontally,” respectively. Notice that aside from this difference in direction, modal and temporal accessibility relations are the same kind of semantic object. Technically, there is no reason to restrict ourselves to just modal and temporal relations, but we will do so in this chapter for simplicity.¹⁷

With the formal tools developed so far, we are now ready to give a precise formal account of the intuitive difference between a fixed past and an open future. To this end, we employ a special modal accessibility relation \approx , whose intended role is to identify *historical alternatives*. The intention is that the historical alternatives of a world w at time t are worlds that are just like w at all times up to and including t , but may differ from w at times later than t .

We impose certain conditions on the relation \approx to ensure its suitability. Recall that $\langle w, t \rangle \approx \langle w', t \rangle$ is to be read as the statement ‘ w' is just like w up

to time t .' Little reflection is needed to see that this relation of 'being-just-like (up to t)' should have the following properties (some of which we have already encountered before):

- (36) a. *reflexivity*: $\langle w, t \rangle \approx \langle w, t \rangle$
 b. *symmetry*: if $\langle w, t \rangle \approx \langle w', t \rangle$, then $\langle w', t \rangle \approx \langle w, t \rangle$
 c. *transitivity*: if $\langle w, t \rangle \approx \langle w', t \rangle$ and $\langle w', t \rangle \approx \langle w'', t \rangle$,
 then $\langle w, t \rangle \approx \langle w'', t \rangle$

Together, these properties ensure that \approx is an *equivalence relation*, which is our first condition. Next, we want to ensure that two worlds which are each other's historical alternatives at some time t , *have been* historical alternatives at all times up to t . This, of course, is simply part of what it means to share the same past up to t , but we need to make the condition explicit. Formally, it means that the worlds that are accessible from w at a given time must also be accessible at all earlier times. In Figure 4, with \approx for ρ , this condition is respected by the modal bases at t and the earlier t' .

Finally, our last condition concerns the idea that being 'just like w ' at a given time implies being indistinguishable from w by all atomic sentences of the language that are evaluated at that time (and thus also all truth-functional compounds which do not include temporal operators). This is stated as a condition that the truth assignment function V must respect. To summarize, we impose the following conditions on the relation \approx of 'being a historical alternative to':

- (37) a. \approx is modal
 b. \approx is an equivalence relation
 c. If $\langle w, t \rangle \approx \langle w', t \rangle$ and $t' < t$, then $\langle w, t' \rangle \approx \langle w', t' \rangle$
 d. If $\langle w, t \rangle \approx \langle w', t \rangle$, then for all atomic sentences p ,
 $V_{\langle w, t \rangle}(p) = V_{\langle w', t \rangle}(p)$

Within this formal setup, we can easily capture the asymmetry between a fixed past and an open future, as well as shed some light on the two notions of truth introduced above.

Truth and settledness. Recall the informal characterization of the two notions of truth: truth at individual worlds (Ockhamist) as opposed to truth at all possible continuations of history (Peircean). In terms of the technical distinctions we introduced, this amounts to *truth simpliciter* on the one

hand, and truth at all historical alternatives, i.e., necessity with respect to the accessibility relation \approx , on the other.¹⁸ This kind of necessity is commonly referred to as *historical necessity* or *settledness*. Notice that it can easily be defined in terms of truth, but not *vice versa*. Peircean truth is Ockhamist settledness, but Ockhamist truth has no analog in the Peircean approach.

Now, in a model in which the accessibility relation \approx and the truth value assignment respect all the conditions we introduced above, these two notions are not independent: they *coincide* for all sentences whose truth value does not depend on times later than time of evaluation. This is guaranteed to be the case for all formulas not containing the operators $\diamond_{<}$ or $\square_{<}$. Historical alternatives at time t are indistinguishable at all times up to and including t , and, therefore, for any given world, sentences whose truth values depend only on the states of affairs at such times are either true at all of its historical alternatives or at none of them. It is only with respect to future reference that historical alternatives may disagree.

3.3. Present tense revisited

With these formal preliminaries in mind, we can now return to the data we cited at the beginning of this section in (29) and (30). As we noted earlier, data such as these have been at the center of much controversy in the literature on English tenses and temporal reference. With the tools now at our disposal, we can give a very simple explanation for the fact that the settledness reading arises with (29c) and (30b,c), and not with (29a,b) or (30a). First, we assume that two tenses are involved in the above sentences, Past and Present, whose temporal interpretation is ‘past’ and ‘non-past,’ respectively. Accordingly, their formal representation will involve the accessibility relations $>$ for the Past and \leq for the Present.

How, then, do we account for the settledness reading arising in (29c) and (30b,c)? Recall the generalization that this reading arises just in case the reference time of the sentence follows the speech time of evaluation. This is the case in (29c) and (30c) due to the adverbial *tomorrow*, and in (30b) because of the aspectual properties of the predicate. Now it is evident that this is precisely the pattern we predict if we assume, against the background of a model as introduced above, that *all* of the six sentences assert that it is not merely true, but *settled* that Megan is in her office at the time in question.

Thus, assuming that our language contains an untensed atomic sentence *Megan be in her office*, we can say that the following sentences in (38a,b) are equivalent to the respective formulas:

- (38) a. Megan was in her office.
 $\Box_{\approx}\Diamond_{>}$ (Megan be in her office)
 b. Megan is in her office.
 $\Box_{\approx}\Diamond_{\leq}$ (Megan be in her office)

The role of the adverbials *now* and *tomorrow* in (29b,c) is to restrict the domain of the operators $\Diamond_{>}$ and \Diamond_{\leq} . The sentences in (30b,c) are treated similarly.

We have shown in this section how the techniques of modal logic are applied in both the modal and the temporal domain, and how due consideration for the interaction between these dimensions can lead to very simple linguistic analyses, here illustrated with the example of the English Present tense. Key to the explanation of the settledness reading was the claim that the interpretation of sentences such as those in (29) and (30) involves the settledness operator \Box_{\approx} .

We stated above that the interpretation of sentences such as (29a,b) is equivalent to the respective formulas with the settledness operator, but we have so far avoided a commitment as to the status of this universal force. Is it part of the truth-conditional meaning of these sentences, or does it enter the interpretation as a pragmatic effect of some kind? On this question, too, there is little agreement. Dowty (1979), for instance, treated it as part of the truth-conditional meaning of sentences of the form *tomorrow* φ (though not of sentences with past and present reference – in this respect, our analysis is more uniform), whereas other authors, notably Steedman (2000), sought a pragmatic explanation. Kaufmann (2002) also argued for a pragmatic explanation, but Kaufmann (2005) took the opposite position in the face of evidence from *if*-clauses and argued that settledness is part of the truth conditions.

While the data behind this distinction are somewhat intricate and beyond the scope of this chapter, it is illuminating to discuss in some detail the theoretical assumptions that made the pragmatic argument compelling in the first place. We will do so in the next subsections. Again, in the course of the discussion we will have occasion to extend and refine our formal framework in ways that will enable us to account for a range of other data as well.

3.4. Objective and subjective uncertainty

Historical necessity, the topic of the last subsection, is a means of dealing formally with the asymmetry between a fixed past and an open future. We saw that settledness is represented as necessity with respect to the set of historical alternatives to the world and time of evaluation. Of central importance was the condition imposed by the relation \approx on admissible truth value assignments in the model: All historical alternatives of a world w at time t are indistinguishable from w at all times up to t . Worlds may part ways only at times later than t .

The kind of uncertainty that this model is designed to capture is *objective* or *metaphysical* uncertainty. The assumption is that the future course of events is literally not determined at present. Accordingly, our uncertainty about the future is not solely due to ignorance of the relevant facts. Rather, it is impossible, even in principle, to know how things will turn out in the future. Even if we could resolve all of our uncertainty about past and present facts, some residual uncertainty about the future would necessarily remain.¹⁹

In contrast to the future, the model mandates that there can be no uncertainty about past or present facts. Now, clearly, this constraint would be too strong if imposed on speakers' beliefs. Ordinary speakers do not know much more about the past than they do about the future. The doxastic accessibility relations encoding speakers' belief states, therefore, must have somewhat different properties from \approx . In particular, they must allow for uncertainty about the past: It should be possible for doxastically accessible worlds at time t to differ from each other with regard to facts at times earlier than t .

However, two assumptions about the interaction between doxastic and metaphysical accessibility are intuitively plausible, and it is these assumptions that will ultimately facilitate our account of certain linguistic data. First, it is reasonable to assume that doxastic states themselves are subject to historical necessity; that is, the relation modeling a speaker's beliefs at a given world-time pair $\langle w, t \rangle$ should be constant across all historical alternatives $\langle w', t \rangle$ such that $\langle w, t \rangle \approx \langle w', t \rangle$. This is because intuitively the contents of a speaker's belief state are facts, just like the *ordinary* facts about the world, and not subject to objective uncertainty at later times.

The second assumption concerns limits on what speakers can rationally believe at a given world and time. Specifically, speakers cannot have full confidence about the truth values of sentences whose truth they also believe

is not yet objectively settled. Having such beliefs would imply that the speaker (believes that she) can “look ahead” in history, an attitude which we will assume (perhaps somewhat optimistically) is not attested.

We will use the symbol \sim to stand for doxastic accessibility relations.²⁰ We assume that \sim , like the metaphysical accessibility relation \approx , is modal (i.e., holds only between indices, or world-time pairs, which share the same temporal coordinate), and that it is serial, transitive and euclidean (i.e., that the corresponding modal base is consistent and fully introspective). The two modal bases are illustrated in Figure 5:

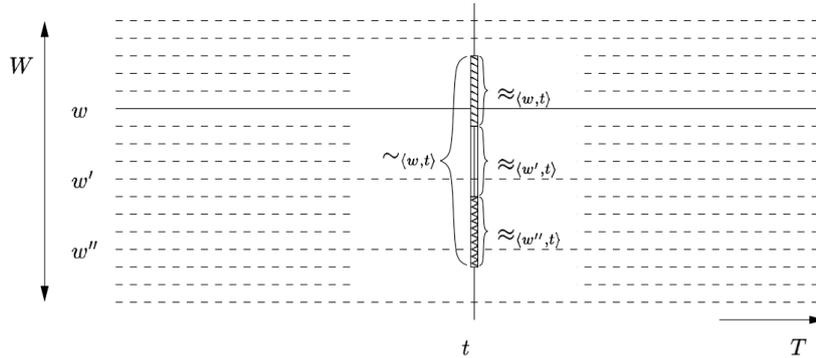


Figure 5. Objective and subjective modal bases, represented by \approx and \sim , respectively. At any given index $\langle w, t \rangle$, the doxastic modal base $\sim_{\langle w, t \rangle}$ must contain $\approx_{\langle w', t \rangle}$ for all $w' \in \sim_{\langle w, t \rangle}$.

Now, the two properties discussed can be formalized as the following constraints on the interaction between the two relations:

- (39) a. *historicity*
 If $\langle w, t \rangle \sim \langle w', t \rangle$ and $\langle w, t \rangle \approx \langle w'', t \rangle$, then $\langle w'', t \rangle \sim \langle w', t \rangle$
 b. *lack of foreknowledge*
 If $\langle w, t \rangle \sim \langle w', t \rangle$ and $\langle w', t \rangle \approx \langle w'', t \rangle$, then $\langle w, t \rangle \sim \langle w'', t \rangle$

Condition (39a) states that two indices cannot be historical alternatives unless they agree on the set of indices that are doxastically accessible. Condition (39b) ensures that if an index $\langle w', t \rangle$ is doxastically accessible, then so are all its historical alternatives. Together, these two constraints guarantee that the two relations interact in the way we discussed above; in particular, in virtue of (39b), objective uncertainty invariably gives rise to subjective

uncertainty (though not *vice versa*): Only what is settled can already be known.

3.5. Settledness and scheduling

Using a model that conforms to the conditions introduced above, it is now quite straightforward to account for the modal connotations of the English Present tense. The generalization was that whenever the reference time lies in the future from the perspective of the evaluation time, the Present is only felicitous on a scheduling reading. Earlier we suggested that scheduling in this context corresponds to settledness, or historical necessity. The sentence is not merely asserted to be true in the world of evaluation, but true in all of its historical alternatives. We also argued that *all* sentences with bare tenses carry this strong reading, not just those in the bare Present.

This assumption is in line with the fact that only future reference gives rise to a settledness condition that is felt as an additional semantic element over and above the mere condition that the sentence be true. Recall now that settledness and truth coincide for all sentences whose truth depends solely on past and present facts. Indeed, by attributing settledness to all such sentences, we can immediately account for the fact that the Present sometimes does and sometimes does not carry the settledness reading. It does so if the reference time lies in the future (as it must with non-statives), but not if the evaluation is co-temporal with the evaluation time (as it can be with statives).

However, why should it be that sentences carry this strong modal force? While the pragmatic account stops short of including the modal element into the truth conditions, it nevertheless gives a simple and compelling answer to this question: Because they are *asserted*.

Two assumptions about the assertion of sentences are sufficient to derive this explanation:

- (40) a. In asserting a sentence, the speaker signals that she believes that it is true.
- b. If the linguistic transaction is successful, the listener will end up believing that the sentence is true.

While these statements gloss over a number of important fine details about the way communication comes to succeed, they are widely accepted and

part of the standard pragmatic account of what goes on in standard communicative situations.²¹

Now, relative to a doxastic accessibility relation \sim , a sentence p is believed to be true at world w and time t if and only if it is true at all world-time pairs $\langle w', t \rangle$ such that $\langle w, t \rangle \sim \langle w', t \rangle$. Furthermore, given the conditions imposed on the interaction between the relations \sim and \approx , p is believed to be true if and only if it is believed to be settled – recall that for any world-time pair $\langle w', t \rangle$ accessible via \sim , the historical alternatives of $\langle w', t \rangle$ are also accessible via \sim . Thus the speaker, in asserting p , signals that she believes that p is not only true, but settled.

Consider, on the other hand, the update to the listener's belief state that results from her accepting the speaker's assertion. It is standardly assumed that this update proceeds by elimination of all those world-time pairs in which the sentence is false. Given the constraints we introduced, this update may lead to an inadmissible belief state in the case of future reference – i.e., a doxastic relation which accesses some world-time pairs, but fails to access all of their historical alternatives. Of course, the listener also knows, from the very fact that the speaker asserted the sentence, that she believes it to be settled. In reaction to this evidence, and to ensure that the conversation proceeds smoothly, she may also eliminate from her belief state those indices in which the truth or falsehood of the sentence is not yet determined. This amounts to an *accommodation* of the information that the truth value of p is objectively settled.

4. Conclusion

In this chapter, we covered fundamental concepts and approaches to the formal treatment of modality in semantic theory. The introduction of formal modal logic into linguistic theory was a significant step forward. Among its main early achievements was the ability to bring order to the class of modal auxiliaries and adverbials, and to analyze what had been a bewildering polysemy, in terms of a few basic modal parameters. We described some of the standard formal modal logic tools that are widely employed in semantic theory, and illustrated their use with examples from expressions that are traditionally considered *modal*. At the same time, the appeal to modality alone is not sufficient for the proper analysis of even the most prototypical modal expressions. Invariably, other notions interact with modality in intricate ways. Hence we also discussed and analyzed some more subtle exam-

ples showing the interaction between modality and other grammatical categories. These interactions have been at the center of much recent theoretical work (see, for example, Condoravdi 2002, Ippolito 2003, Kaufmann 2005). The area of modality continues to be exciting and rife with new ideas and further questions.

Acknowledgement

The authors would like to thank William Frawley for continuous support and encouragement, as well as assistance with the manuscript.

Notes

1. From Greek *aletheia* ‘truth.’
2. From Greek *deon* ‘obligation.’
3. From Greek *episteme* ‘knowledge.’
4. From Greek *doxa* ‘belief.’
5. An axiom is a sentence whose provability is guaranteed by assumption. An axiom schema is a sentence of a certain form, for example, $\varphi \rightarrow (\theta \rightarrow \varphi)$, where φ and θ stand for arbitrary propositional sentences.
6. This is a simplification. Sentences carrying semantic presuppositions are commonly assumed to be neither true nor false at worlds at which their presuppositions are not satisfied.
7. To be sure, none of the logical connectives adequately captures the semantic richness of the English words that are customarily associated with it. However, as our primary goal in this chapter is a discussion of modals, we will adopt the connectives without further discussion.
8. One could think of $[\Phi]^M$ as the denotation of the conjunction of all members of Φ , but this analogy breaks down if Φ has infinitely many members. In this case, $[\Phi]^M$ is still defined, but our language does not include the corresponding infinite conjunction.
9. Notice that the absence of the condition does not mean that w cannot be a member of R_w – it merely does not have to be.
10. This kind of problem is known as the *Samaritan paradox*.
11. A binary relation on a set is a preorder if it is reflexive and transitive.
12. The semantics in (28) could be simplified if it were guaranteed that there were minimal elements in the preorder.

13. The formal apparatus introduced in this section is partly based on Thomason (1984).
14. Another compelling possibility, which we will not discuss here, is to consider multi-valued logic, or to say that sentences whose truth value is not yet determined are not false, but truth-valueless (e.g., see Thomason 1970 for discussion).
15. World-time pairs of this sort are also behind Montague's (1973) treatment of intensionality. It is fair to say, though, that the analysis of modality is not among the areas in which Montague himself made very substantive contributions.
16. This assumption has nontrivial logical consequences, which are immaterial for our purpose.
17. Kaufmann (2005) argues that hybrid modal-temporal relations that operate "diagonally" are required for the analysis of certain conditionals.
18. *Truth simpliciter* is equivalent to necessity with respect to the identity relation.
19. It is immaterial in this connection whether physicists tell us that the world is in fact deterministic or non-deterministic. What matters for our purpose is that we talk *as if* certain things could not be known in advance.
20. Belief states are subjective, tied to individual agents. We assume here that we are speaking about the beliefs of a particular speaker, but, for simplicity, we do not indicate this formally.
21. See Stalnaker (1978) for an early exposition.

References

- Condoravdi, C.
 2002 Temporal interpretation of modals: Modals for the Present and for the Past. In *The Construction of Meaning*, D. I. Beaver, L. D. Casillas Martínez, B. Z. Clark and S. Kaufmann (eds.), 59–87. Stanford: CSLI Publications.
- Dowty, D. R.
 1979 *Word Meaning and Montague Grammar*. Berlin: Springer.
- Hintikka, J.
 1961 Modality and quantification. *Theoria* 27: 119–128.
- Horn, L. R.
 2001 *A Natural History of Negation*. Stanford: CSLI Publications (originally Chicago: University of Chicago Press).
- Hughes, G. E. and M. J. Cresswell
 1996 *A New Introduction to Modal Logic*. London: Routledge.
- Ippolito, M.
 2003 Quantification over times in subjunctive conditionals. In *Proceedings of Semantics and Linguistic Theory XIII*, R. B. Young and Y. Zhou (eds.), 145–186. Cornell University: CLC Publications.

- Kaufmann, S.
 2002 The presumption of settledness. *CLS* 38: 313–328.
 2005 Conditional truth and future reference. *Journal of Semantics* 22: 231–280.
- Kratzer, A.
 1977 What ‘must’ and ‘can’ must and can mean. *Linguistics and Philosophy* 1: 337–355.
 1981 The notional category of modality. In *Words, Worlds, and Contexts: New Approaches in Word Semantics*, H. J. Eikmeyer and H. Rieser (eds.), 38–74. Berlin/New York: Mouton de Gruyter.
 1991 Modality. In *Semantik: Ein Internationales Handbuch der Zeitgenössischen Forschung*, A. von Stechow and D. Wunderlich (eds.), 639–650. Berlin/New York: Mouton de Gruyter.
- Kripke, S.
 1963 Semantical analysis of modal logic, I. Normal modal propositional calculi. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik* 9: 67–96.
- Lewis, D.
 1981 Ordering semantics and premise semantics for counterfactuals. *Journal of Philosophical Logic* 10: 217–234.
- Mendelson, E.
 1987 *Introduction to Mathematical Logic*. Pacific Grove, CA: Wadsworth.
- Montague, R.
 1973 The proper treatment of quantification in ordinary English. In *Approaches to Natural Language*, K. J. J. Hintikka, J. M. E. Moravcsik and P. Suppes (eds.), 221–242. Dordrecht: Reidel.
- Prior, A. N.
 1957 *Time and Modality*. Oxford: Clarendon Press.
 1967 *Past, Present and Future*. Oxford: Clarendon Press.
- Stalnaker, R.
 1975 Indicative conditionals. *Philosophia* 5: 269–286.
 1978 Assertion. In *Syntax and Semantics*, vol. 9, P. Cole (ed.), 315–332. New York: Academic Press.
- Steedman, M.
 2000 *The Syntactic Process*. Cambridge, Mass.: MIT Press.
- Thomason, R. H.
 1970 Indeterminist time and truth-value gaps. *Theoria* 36: 264–281.
 1984 Combinations of tense and modality. In *Handbook of Philosophical Logic: Extensions of Classical Logic*, vol. II, D. M. Gabbay and F. Guenther (eds.), 135–165. Dordrecht, Reidel.

