COSC462 Lecture 8: Modal languages

Willem Labuschagne
University of Otago

Abstract

Finite ranked interpretations are generalised to possible worlds interpretations and a modal operator $\Box$ is introduced to the object language. We briefly examine how to interpret $\Box$ in different ways to get the logic of necessity, temporal logic, and epistemic logic.

1 Preamble

Soon Hans will be talking to you about *epistemic logic*, and today's lecture is intended to build a bridge for you, from nonmonotonic logic to epistemic logic. Basically, it's not a long bridge. In nonmonotonic logic we associated with the language L_A a thing of the form $(S, \preccurlyeq, V)$, which we called by a fancy name — a finite ranked interpretation. In epistemic logic you will associate with your language a very similar thing of the form $(S, \sim, V)$, called by a different fancy name — a Kripke model.

There are some things to look out for, if you are not to become confused by changes of notation and terminology. In the remainder of this section we give a brief summary of what to watch out for. The second section gives a more extended account of how epistemic logic fits into modal logic, and what motivated the development of modal logic.

A (single-agent) Kripke model, at its most general, is a triple (S, R, V) where R may be any binary relation on S. So the ranked interpretations we used in lectures 6 and 7 are examples of Kripke models, with R a total preorder usually named $\preccurlyeq$ or something similar. But in the case of Kripke models for epistemic logic:

- we do **not** assume that S is finite
- we often call the members of S *possible worlds* instead of states

- the relation R will be an equivalence relation called something like $\sim$, and
- the labelling function V is usually called a valuation and is often described in a different way.

You should watch out for the following potentially confusing differences between what you're used to and what is taken for granted in epistemic logic.

Difference 1: The labelling function $V : S \longrightarrow W_A$ does a simple and specific job — it associates with every $s \in S$ an assignment $w \in W_A$ of truth values to atoms (i.e. $w : A \longrightarrow \{0, 1\}$). Basically, V sees to it that we can take any $s \in S$ and any atom p and work out whether p is satisfied by s. Now, this job can be done in other ways. For example, instead of using 0 and 1 as truth values, we might use `false` and `true` instead. More radically, we could let V be a function that eats atoms and spits out the states satisfying the atom, i.e. $V : A \longrightarrow \wp(S)$. For example, if we consider the Light-Fan System with $S = W_A = \{11, 10, 01, 00\}$ and $A = \{p, q\}$, then we could let V be defined as follows: $V(p) = \{11, 10\}$ and $V(q) = \{11, 01\}$. This tells us everything we need to know about which states satisfy which atoms. And in epistemic logic most of the literature uses this approach to V.

Difference 2: When we have spoken of satisfaction (i.e. truth), we have tended to say "s satisfies φ" and let the context determine which ontology (S, V) or which ranked interpretation $(S, \preccurlyeq V)$ was relevant. In epistemic logic, the fashion is to mention the relevant Kripke model explicitly, and because it is very cumbersome to write "s satisfies φ relative to the Kripke model $(S, \sim, V)$", a concise notation for satisfaction is typically used. Some symbol is chosen to abbreviate "satisfies", for example we might choose the symbol $\Vdash$. Some shorter name is chosen for the Kripke model, for example we might decide that $M = (S, \sim, V)$. And now we may write

"$M, s \Vdash \varphi$"

as a short way to say

"s satisfies φ relative to the Kripke model $M = (S, \sim, V)$".

While several modern authors use the symbol $\Vdash$, often the symbol for entailment, $\models$, is used instead.

Difference 3: Apart from Kripke models, epistemic logic also uses a propositional language containing, in addition to the connectives you know and love, some new connectives that you haven't seen before. These new connectives stand for modal operators. Specifically, there is an operator K that stands for "Knows", or more comprehensibly "the agent knows that", so that a sentence of the form $K\varphi$ is read "the agent knows that φ". And there is another operator M (yes, the same symbol as the often use for the Kripke model, grrr, spit) and M stands for "Maybe", or more comprehensibly "the agent considers it possible that", so sentences of the form $M\varphi$ are read "the agent considers it possible that φ".

Difference 4: Satisfaction of sentences containing the new connectives is a bit more complicated than for the connectives you're accustomed to:

- $M, s \Vdash K\varphi$ if and only if $M, s' \Vdash \varphi$ for all s' such that $s \sim s'$

  In other words, s satisfies $K\varphi$ in a Kripke model $M = (S, R, V)$ iff **all** the little s' that are related to s by R satisfy the shorter sentence φ. For example, if the relation R were a total preorder $\preccurlyeq$, then s would satisfy $K\varphi$ iff all the s' such that $s \preccurlyeq s'$ satisfied φ.

- $M, s \Vdash M\varphi$ if and only if $M, s' \Vdash \varphi$ for some s' such that $s \sim s'$.

  In other words, s satisfies $M\varphi$ in a Kripke model $M = (S, R, V)$ iff **at least one of** the little s' that are related to s by R satisfies the shorter sentence φ. For example, if the relation R were a total preorder $\preccurlyeq$, then s would satisfy $K\varphi$ iff at least one of the s' such that $s \preccurlyeq s'$ satisfied φ.

2 Necessity

The father of modal logic, CI Lewis, began in 1912 by exploring the difference between assertions such as the following:

- α = Either Matilda does not love me, or she does love me.
- β = Either Caesar died, or the moon is made of green cheese.

Assertion β happens to be locally true in 'the real world', because it is an historical fact that there was such a person as Julius Caesar who was stabbed to death on the Ides of March in the year 44 BC. But we can imagine a 'possible world' in which both *Caesar died* and *The moon is made of green cheese* are false — for example, our world at the time of Caesar's invasion of Britain in 54 BC, or a fictional world in which Caesar rescued a fairy princess from a wicked witch and lived happily ever after. Formalised, β gives a sentence of L_A having the form $p \vee q$, and this sentence is false in states that make both p and q false.

Assertion α, in contrast, is not just locally true but holds in all possible worlds, i.e. globally. Formalised, α gives a sentence of L_A having the form $p \vee \neg p$, and this is a tautology. No state can make a tautology false.

The distinction between α and β poses a challenge for the logician who wishes to represent this difference inside the object language itself. Syntactically, we would like to have some way to say, in the object language, that α is, and β is not, supposed to hold globally, at all states $s \in S$. To do this, we may enrich the language L_A by adding to it a new 'modal' operator $\Box$, which is a kind of connective. The intention is that $\Box\gamma$ will be read 'It is necessarily the case that γ'. (There are other readings that may be more interesting, but this was the original idea behind modal logic.) However, this brings us face to face with a tricky question: how are truth values to be allocated to sentences of the new richer language?

The simplest way to look at it is the following. We start with an ordinary propositional language L_A for knowledge representation. We associate with the language L_A a set S of states or possible worlds. Intuitively, one of the states in S is the actual state of the world, in which Caesar existed and died. Other members of S correspond to imagined states of the world. In some of the imagined worlds Caesar never existed and in others he discovered a elixir of immortality and did not die. Thus β is satisfied in some states, notably in the actual state, but not in others. Since β is not globally true over all the possible states of the world, the modal sentence $\Box\beta$ should be false. On the other hand, we would expect the tautology α to be globally true over all states of the world because of the way connectives interact with truth values. This should be reflected by our semantics making the modal sentence $\Box\alpha$ true Of course, these ideas are very informal and we will need to check whether it is possible to give a precise definition of satisfaction.

Let's start by saying precisely what the language is that we get when we add the modal operator.

Definition 1 *Suppose A is some nonempty set of atoms. The modal language $L_A^\Box$ is the set of all sentences α, where α is a sentence iff one of the following is the case:*

- *$\alpha \in A$ (i.e. α is an atom)*
- *$\alpha = \neg\beta$ for some previously constructed sentence β*
- *$\alpha = (\beta * \gamma)$ with $* \in \{\wedge, \vee, \rightarrow, \leftrightarrow\}$ and β, γ previously constructed sentences*
- *$\alpha = \Box\beta$ for some previously constructed sentence β.*

The underlying nonmodal language of $L_A^\Box$ is the set L_A of all sentences in which $\Box$ does not occur.

As an example, let us take the modal propositional language having a set of two atoms, $A = \{p, q\}$.

The underlying **nonmodal** language L_A is just the propositional language built up from $A = \{p, q\}$ by means of the usual connectives $\wedge$, $\vee$, $\rightarrow$, and $\leftrightarrow$. This familiar language has the set W_A of valuations $v : A \longrightarrow \{0, 1\}$, written as 11, 10, 01, and 00.

The sentences of the **modal** language $L_A^\Box$ include such additional strings as $\Box(p \vee q)$, with the intended reading 'It is necessarily the case that p or q', and $\Box p \vee \Box q$, read 'Necessarily p or necessarily q'. The reader may find it interesting to pause and think about whether these two modal sentences should be equivalent or not.

Remark 2 *In sentences of the modal language, a useful pattern is formed when negation and the box operator are combined in the sequence* $\neg\Box\neg$. *Since we would read* $\neg\Box\neg p$ *as 'It is not the case that necessarily not* p*', and this appears to be saying 'It is possibly the case that* p*', we introduce a new symbol* $\Diamond$ *(the diamond operator) as an abbreviation for* $\neg\Box\neg$. *Thus* $\Diamond\alpha$ *is an alternative way to write* $\neg\Box\neg\alpha$*, and is read 'Possibly* α*'.*

How should truth values be allocated to sentences of $L_A^\Box$? Our experience with nonmonotonic logic and ranked interpretations offers guidance.

Recall that a ranked interpretation associated three things with the language: a set S of states, a function V connecting states to valuations so that the process of allocating truth values could get off the ground, and a relation (total preorder $\preccurlyeq$) on S which allowed us to focus on certain sets of states (the minimal models of sentences). We generalise ranked interpretations to allow relations other than total preorders to be used.

Definition 3 *A* ***Kripke model*** *of the modal propositional language* $L_A^\Box$ *is a triple* $\mathcal{M} = (S, R, V)$ *such that*

- S *is any nonempty set, the members of which are called states or possible worlds*
- $R \subseteq S \times S$*, called the accessibility relation*
- $V : S \longrightarrow W_A$*, called the labelling function or the valuation.*

  The pair (S, R) *is called a* ***frame****.*

Remark 4 *Clearly there is a potential confusion whenever we use the word 'valuation'. We might mean one of the familiar little things in* W_A*, or we might mean* V*. Life is hard.*

The term 'Kripke model' is used in honour of Saul Kripke, a famous logician who, at the age of sixteen, published in 1959 a very good paper on modal logic in which he used triples (S, R, V). However, other logicians had used the relevant ideas before that, starting with Rudolf Carnap in 1947, Alfred Tarski in 1951, the New Zealand logician and inventor of temporal logic Arthur Prior in 1957, the Finnish logician Jaakko Hintikka in 1957, and Stig Kanger in 1957. We should call them Carnap models, or Carnap-Tarski-Prior-Hintikka-Kanger-Kripke models :-)

Example 5 *Here is a Kripke model. Consider the Light-Fan System, and take $A = \{p, q\}$ so that $W_A = \{11, 10, 01, 00\}$. Let $S = \{11, 10, 01, 00\} = W_A$ and $R = S \times S$. Of course, V is just the identity function on S. We shall see that this Kripke model allows us to read $\Box$ as 'it is necessarily the case that'. For this we need a precise definition of satisfaction.*

Definition 6 *If $\mathcal{M} = (S, R, V)$ is a Kripke model for the modal propositional language $L_A^\Box$ then a sentence α is satisfied at (or by) $s \in S$ in $\mathcal{I}$ iff one of the following is the case:*

- *$\alpha \in A$ and the valuation $V(s)$ makes α true, i.e. $(V(s))(\alpha) = 1$*
- *$\alpha = \neg\beta$ and s does not satisfy β*
- *$\alpha = \beta \wedge \gamma$ and s satisfies both β and γ*
- *$\alpha = \beta \vee \gamma$ and s satisfies at least one of β and γ*
- *$\alpha = \beta \rightarrow \gamma$ and s satisfies γ or fails to satisfy β*
- *$\alpha = \beta \leftrightarrow \gamma$ and s satisfies both β and γ or v satisfies neither*
- *$\alpha = \Box\beta$ and for every s' such that $(s, s') \in R$, s' satisfies β. (Thus $\Box\beta$ is satisfied at s iff β is satisfied at every s' accessible from s via the relation R.)*

Example 7 ***Necessity:***

Take $L_A^\Box$ to be the modal language with $A = \{p, q\}$, so that $W_A = \{11, 10, 01, 00\}$. The simplest Kripke model of $L_A^\Box$ is $\mathcal{M} = (S, R, V)$ with $S = W_A$, $R = W_A \times W_A$, and V the identity function on W_A. (Why is this so simple? Because we are in effect getting rid of the relation. If R consists of all pairs, then it doesn't constrain us at all.)

Intuitively we expect a sentence of the form $\Box\beta$ to be satisfied at s if β is a tautology of the underlying nonmodal language L_A. Let us test this.

Consider the tautology $p \vee \neg p$.

Is $\Box(p \vee \neg p)$ *is satisfied at the state* $s = 11$*? Well, what are the states accessible from* 11*? The accessibility relation consists of all possible pairs of states, and so every state is accessible from every state, and thus the states* 11, 10, 01, *and* 00 *are accessible from* 11*. Now let us chop off the operator at the front of* $\Box(p \vee \neg p)$*. Is the shorter sentence* $p \vee \neg p$ *that we get, satisfied at all the states accessible from* 11, *namely* 11, 10, 01, *and* 00*? Yes indeed, since* $p \vee \neg p$ *is a tautology.*

In a similar fashion we can check that $\Box(p \vee \neg p)$ *is satisfied at the remaining states of* S *too.*

Is every sentence having the form $\Box\beta$ *satisfied at* $s = 11$*? No. Take* $\beta = q$*. Then we can find some state that fails to satisfy* β*, in this case* 10*. And now* $s = 11$ *cannot satisfy* $\Box q$ *because* 10 *is accessible from* 11 *but does not satisfy* q*.*

So suppose $\Box\beta$ *is satisfied at* $s = 11$*. What can we say about* β*? Well,* β *must be satisfied at every state in* $S = W_A$*, and so* β *has to be a tautology. Doesn't it? Well, we need to be a bit careful when we use the word 'tautology'. When we use the word 'tautology', we have in mind a nonmodal sentence like* $p \vee \neg p$ *or like* $p \to p$ *which, by virtue of the way the usual propositional connectives behave, is satisfied by every valuation of the nonmodal language* L_A*. We have seen that if* β *is a tautology,, say* $\beta = p \vee \neg p$*, then* $\Box(p \vee \neg p)$ *is satisfied by all states in* S*. Thus* $\Box\Box(p \vee \neg p)$ *is satisfied at* $s = 11$*, and is of the form* $\Box\beta$ *where* $\beta = \Box(p \vee \neg p)$*, but normally* $\Box(p \vee \neg p)$ *would not be called a tautology, because it isn't an ordinary nonmodal sentence.*

The moral of the story is that in modal logics there will always be some notation introduced to allow us to say neatly and concisely that a sentence (possibly containing modal operators inside it) is satisfied at all states — we can't use the convenient phrase "is a tautology" except for nonmodal sentences.

The example shows that the Kripke model $\mathcal{M}$ allows us to read $\Box$ as 'it is necessarily the case that', because

- by definition a sentence of the form $\Box\alpha$ is satisfied at state s iff α is satisfied at each of the states accessible from s
- our choice of $R = W_A \times W_A$ for the accessibility relation means that every state is accessible from s, no matter which s we look at
- and thus a sentence $\Box\alpha$ is satisfied at a state s iff α is satisfied at every state in S.

Remark 8 *Bearing in mind that every* $\Box$ *is accompanied by a diamond operator* $\Diamond$*, where* $\Diamond$ *is an abbreviation for* $\neg\Box\neg$*, two obvious questions* *arise: what does it take to satisfy a sentence of the form $\Diamond\beta$, and what is the intuitive meaning of $\Diamond$?*

A moment's thought should reveal that a state s satisfies $\Diamond\beta$ iff there is at least one state s' such that $(s, s') \in R$ and s' satisfies β. (Why? Because s satisfies $\neg\Box\neg\beta$ iff it is not the case that, at every state s' accessible from s, β is false.)

Intuitively, therefore, if we have a Kripke model that allows us to read $\Box$ as 'it is necessarily the case that', then $\Diamond$ may be read as 'it is possible that' or, more long-windedly, as 'there is a possible world in which'.

The theorem below shows that some non-tautological sentences are 'universally true', i.e. are satisfied by every state in every Kripke model..

Theorem 9 *If $\alpha \in L_A^\Box$ is of the form $\Box(\beta \to \gamma) \to (\Box\beta \to \Box\gamma)$ then s satisfies α for every $s \in S$ in every Kripke model (S, R, V) of $L_A^\Box$.*

Proof. *Choose any Kripke model $\mathcal{M} = (S, R, V)$ of $L_A^\Box$. Choose any sentence α of the form $\Box(\beta \to \gamma) \to (\Box\beta \to \Box\gamma)$. Choose any $s \in S$.*

Assume that α fails to be satisfied at s. Thus $\Box(\beta \to \gamma)$ is satisfied at s but $(\Box\beta \to \Box\gamma)$ is not. Since $(\Box\beta \to \Box\gamma)$ is not satisfied at s, $\Box\beta$ is satisfied at s but $\Box\gamma$ is not. Since $\Box\gamma$ is not satisfied at s, there must be some s' such that $(s, s') \in R$ but γ is not satisfied at s'.

However, recall that $\Box(\beta \to \gamma)$ is satisfied at s. So $\beta \to \gamma$ is satisfied at every world accessible from s, in particular at s'. Moreover, recall that $\Box\beta$ is satisfied at s, so that β is satisfied at all worlds accessible from s, including s'. Since both β and $\beta \to \gamma$ are satisfied at s', it follows that γ must be satisfied at s'. But it is not.

Our initial assumption, to the effect that α fails to be satisfied at s, has led to a contradiction, and so must be given up. Thus α is satisfied at the arbitrarily chosen s and therefore at every state $s \in S$ in $\mathcal{I}$. And $\mathcal{I}$ too was arbitrarily chosen. ∎

Remark 10 *The schema $\Box(\beta \to \gamma) \to (\Box\beta \to \Box\gamma)$ is called K in honour of Saul Kripke.*

Depending on the kind of relation R that we choose, we get various special kinds of modal logic. Let's briefly look at two examples: temporal logic and epistemic logic.

Example 11 ***Temporal logic:***

Stick to the language $L_A^\Box$ with $A = \{p, q\}$, and keep thinking of the Light-Fan System as the system of interest. But consider the frame (S', R') with $S' = \mathbb{N}$, in other words $S' = \{0, 1, 2, \ldots\}$, and let R' be the usual order relation $<$ on the set of natural numbers. The idea is to think of the states in S' as time instants. This frame forms the basis of a class of temporal Kripke models in which $\Box$ means 'in future'. In other words, I claim that we can build a lot of different Kripke models from the frame (S', R') by choosing various labelling functions $V' : S' \to W_A$, and in each of these Kripke models it will make sense to read $\Box$ as 'henceforth' or 'at all future times'. For example, we could take $V' : S' \to W_A$ to be the function given by $V'(s) = 10$ if s is an even number while $V'(s) = 11$ if s is odd. Intuitively V' says that the system starts at time 0 with the light on and the fan off (because $p \wedge \neg q$ is satisfied at state 0) and then the fan switches on and off at each succeeding instant while the light never goes off (because s satisfies q whenever s is odd but satisfies $\neg q$ whenever s is even, while state 0 satisfies $\Box p$). Clearly we could define other labelling functions each of which portrays a different history for the system.

What does $\Diamond$ mean? At some future time.

Example 12 ***Epistemic logic***

Let's make (S'', R'') a really interesting frame. Let's keep thinking of the Light-Fan System. Take $S'' = S$, in other words $S'' = \{11, 10, 01, 00\}$, and let $R'' = \{(11, 11), (10, 10), (01, 01), (00, 00), (11, 10), (10, 11), (00, 01), (01, 00)\}$. The relation R'' is an equivalence relation, i.e. has the properties of being reflexive on S (since $(s, s) \in R''$ for every $s \in S$), symmetric (because if $(s, t) \in R''$ then $(t, s) \in R''$) and transitive (because if $(s, t) \in R''$ and $(t, u) \in R''$ then also $(s, u) \in R''$). So we could give R'' a name that suggests an equivalence relation, like $\sim$. Now $\sim$ is the relation such that $11 \sim 11$, $10 \sim 10$, $01 \sim 01$, $00 \sim 00$, $11 \sim 10$, $10 \sim 11$, $00 \sim 01$, and $01 \sim 00$. (Do you see that these 'equations' just tell us what ordered pairs belong to the relation $\sim$?)

Here is a diagram depicting the equivalence relation.

01	00
11	10

The equivalence relation $\sim$ divides (partitions) the set S into two equivalence classes, $\{11, 10\}$ and $\{01, 00\}$. Intuitively, $\sim$ represents the point of view of an agent who cannot tell the difference between the states 11 and 10, nor the difference between the states 01 and 00. This might be an agent looking at the Light-Fan system but placed so far from it that, although she can see whether the light is on or off, she cannot detect whether the fan is on or off. Thus the accessibility relation is doing something extremely interesting here — it embodies the constraints on the agent's capacity to extract information from the system of interest. In a sense the frame $(S'', \sim)$ is drawing a map not only of the Light-Fan System but of a broader system that adds an agent to the Light-Fan System. We still think of the states in S as states of the Light-

Fan System. We shall see that, by virtue of $\sim$, this frame supports the interpretation of $\Box$ as 'the agent knows that'.

Because we have chosen $S'' = W_A$, there is really only one labelling function V'' that gives a sensible interpretation, namely the identity function $V''(s) = s$.

Having chosen S'', $\sim$, and V'', we have a Kripke model such that, for instance, $s = 11$ satisfies $\Box p$, because p is satisfied at both 11 and 10, the states accessible from 11. Is this what we would expect? Yes — intuitively $\Box p$ should be true at state 11 because, when the system is in state 11, the agent can see that the light is on (i.e. that p is the case). On the other hand, state 11 fails to satisfy $\Box q$, since 10 is accessible from 11 but 10 does not satisfy q. This makes intuitive sense too, because when the system is in state 11 then the agent can see that the light is on but cannot see whether the fan is on, and so cannot know that q is the case.

The moral of the story is that we can use equivalence relations as accessibility relations if we want to represent the knowledge that an agent is able to extract from a system. When the system is in a state s, the agent ***knows*** *α if α is true in all the states that, to the agent, are indistinguishable from s, i.e. in all states that are 'equivalent' to s. And thus to say that the agent knows α relative to state s is the same as to say that $\Box\alpha$ is satisfied at s. And since they want to be reminded of 'knows', epistemic logicians use the symbol K instead of $\Box$.*

What does $\Diamond$ mean? If $\Diamond\alpha$ is satisfied at s, then when the system is in the state s the agent considers α to be possible, because the agent does not know $\neg\alpha$. And since they want to be reminded that 'maybe' α is the case, epistemic logicians use the symbol M instead of $\Diamond$.

Exercise 13 *1. Consider the Light-Fan-Heater System, which has the three obvious components. Let $L_A^{\Box}$ be the modal language having $A = \{p, q, r\}$. The intention is that p should express the claim that the light is on, q the claim that the fan is on, and r that the heater is on. Recall that $W_A = \{111, 110, \ldots, 001, 000\}$.*

Suppose that the blueprint of the system reveals that it is not possible to have all three components on simultaneously, so that the valuation 111 does not represent a 'realisable' state of the system. Assume that the remaining valuations in W_A do represent realisable states of the system.

- *Describe a Kripke model $\mathcal{I} = (S, R, V)$ of $L_A^{\Box}$ in which S is the set of realisable states of the Light-Fan-Heater System.*
- *Is $\Box(\neg p \vee \neg q \vee \neg r)$ satisfied at every $s \in S$?*
- *Is $(\Box\neg p) \vee (\Box\neg q) \vee (\Box\neg r)$ satisfied at every $s \in S$?*

2. *In the theorem of this section, we met a schema K, namely $\Box(\beta \to \gamma) \to (\Box\beta \to \Box\gamma)$. Schema K has the remarkable property that every sentence of this form is universally true (i.e. true at every state in every Kripke model). Examine each of the following schemas of $L_A^{\Box}$ and determine whether all sentences of that form are universally true. If so, give an argument to support your conclusion. If not, give a counterexample consisting of some specific sentence of the correct form and some Kripke model having a state at which the sentence fails to be satisfied.*

   - $\Box\beta \leftrightarrow \neg\Diamond\neg\beta$ — *or with more parentheses,* $(\Box\beta) \leftrightarrow (\neg\Diamond\neg\beta)$
   - $\Diamond(\beta \to \gamma) \to (\Diamond\beta \to \Diamond\gamma)$
   - $\Box(\beta \wedge \gamma) \to (\Box\beta \wedge \Box\gamma)$
   - $\Box(\beta \vee \gamma) \to (\Box\beta \vee \Box\gamma)$
   - $(\Diamond\beta \wedge \Diamond\gamma) \to \Diamond(\beta \wedge \gamma)$

Glossary

- **frame** — a pair (S, R) consisting of a set S of states together with a relation $R \subseteq S \times S$.

- **Kripke model** — what we get when we combine a frame and an ontology to get a triple (S, R, V).

- **modal operator** — either a box operator $\Box$ or a diamond operator $\Diamond$, where $\Diamond$ is an abbreviation for $\neg\Box\neg$; in epistemic logic, it is common for K to be used instead of $\Box$ and M instead of $\Diamond$; think of K as 'knows' and M as 'maybe'.

- **possible world** — a member of S, what we usually call a state.

- **valuation** — either a familiar function $v : A \longrightarrow \{0, 1\}$ sending atoms to truth values or a labelling function $V : S \longrightarrow W_A$ or something equivalent to a labelling function, e.g. a function $V : A \longrightarrow \wp(S)$ associating with each atom the set of states satisfying the atom, or a function $V : A \times S \longrightarrow \{0, 1\}$ associating with every atom and state a truth value.