

7 Temporal Logic

7.1 Reasoning about time

It is currently raining in Edinburgh. But it wasn't raining yesterday, and perhaps it won't rain tomorrow. Let's introduce some operators to formalize reasoning about the unfolding of events through time.

If we read r as 'it is raining', we will use Fr to express that it will be raining at some point in the future. We will use Pr to express that it has been raining at some point in the past. In general:

FA is true at a time t iff A is true at some time after t .

PA is true at a time t iff A is true at some time before t .

The operators F and P can be nested. We can use FPr to express that at some point it will have rained, PFr to say that it was once going to rain, PPr to say that there was a time before which it rained, and FFr to say that there will come a time after which it will rain.

Unlike \Box and \Diamond , F and P are not duals of each other: $\neg PA$ is not equivalent to $F\neg A$, and $\neg FA$ is not equivalent to $P\neg A$. But it is useful to have duals of F and P . We therefore introduce two more operators. G will be the dual of F , and H the dual of P .

Intuitively, GA means that A is always *going* to be the case. (Hence the symbol 'G'.) If it is not the case that at some point in the future it will not rain ($\neg F\neg r$), then it is always going to be the case that it will rain (Gr). Similarly, HA means that A *has* always been the case. If it is not the case that at some point in the past it was not raining ($\neg P\neg r$), then it has always been raining (Hr).

We can state the truth-conditions of GA and HA in parallel to the above truth-conditions for FA and PA :

GA is true at a time t iff A is true at all times after t .

HA is true at a time t iff A is true at all times before t .

The language of standard propositional logic, extended by the four operators F, P, G, H is known as the **language of basic temporal logic**. We will sometimes call it \mathcal{L}_t .

Exercise 7.1

Translate the following sentences into the language of basic temporal logic.

- (a) It has never been warm.
- (b) There will be a sea battle.
- (c) There will not have been a sea battle.
- (d) At some point, it will be warm or it will have been warm.
- (e) If you haven't studied, you won't pass the exam.
- (f) I was having tea when the door bell rang.

7.2 Temporal models

A complete scenario for temporal logic needs to tell us what times there are, how they are ordered, and what is going on at each of them. We can represent such a scenario, together with an interpretation of \mathcal{L}_t 's non-logical vocabulary, by a structure that settles (a) what times there are, (b) which times come before or after which others, and (c) which sentence letters are true at which times. This is enough to determine, for every \mathcal{L}_t -sentence and every time, whether the sentence is true at that time.

Definition 7.1: Temporal Model

A **temporal model** consists of

- a non-empty set T (of “times”),
- a binary relation $<$ on T (the **precedence relation**),
- a function V that assigns to each sentence letter of \mathcal{L}_T a subset of T .

We use ' $M, t \models A$ ' as a short-hand notation to express that sentence A is true at time t in model M . The following definition formally specifies the truth-value of every \mathcal{L}_T -sentence at every time in every model.

Definition 7.2: Standard Temporal Semantics

If $M = \langle T, <, V \rangle$ is a temporal model, t is a member of T , P is any sentence letter, and A, B are any \mathcal{L}_T -sentences, then

- (a) $M, t \models P$ iff t is in $V(P)$.
- (b) $M, t \models \neg A$ iff $M, t \not\models A$.
- (c) $M, t \models A \wedge B$ iff $M, t \models A$ and $M, t \models B$.
- (d) $M, t \models A \vee B$ iff $M, t \models A$ or $M, t \models B$.
- (e) $M, t \models A \rightarrow B$ iff $M, t \not\models A$ or $M, t \models B$.
- (f) $M, t \models A \leftrightarrow B$ iff $M, t \models (A \rightarrow B)$ and $M, t \models (B \rightarrow A)$.
- (g) $M, t \models F A$ iff $M, s \models A$ for some $s \in T$ such that $t < s$.
- (h) $M, t \models G A$ iff $M, s \models A$ for all $s \in T$ such that $t < s$.
- (i) $M, t \models P A$ iff $M, s \models A$ for some $s \in T$ such that $s < t$.
- (j) $M, t \models H A$ iff $M, s \models A$ for all $s \in T$ such that $s < t$.

Clause (a) says that a sentence letter is true at a time in a model iff the model's interpretation function specifies that the sentence letter is true at that time. Clauses (b)–(f) say that the truth-functional connectives have their normal truth-table meaning at each time. Clauses (g)–(j) formalize the truth-conditions for temporal sentences from the previous section.

All this should remind you of our Kripke semantics for \mathcal{L}_M in chapter 3. In fact, temporal models *are* Kripke models, as defined on page 52. I have merely relabelled the set ' W ' as ' T ', and the relation ' R ' as ' $<$ '. Definition 7.2 resembles definition 3.2 from page 53, except that we have two box-like operators G and H , and two diamond-like operators F and P . The language of basic temporal logic is bi-modal, with forward-looking operators (F and G) and backward-looking operators (P and H). Unlike ordinary models for multi-modal languages (definition 5.1), temporal models have only a single accessibility relation. That's because the accessibility relation for P and H is definable from the accessibility relation for F and G : a time s is earlier than a time t iff t is later than s .

Let's look at an example of a temporal model. For the set of times T , we use the set of natural numbers 0, 1, 2, etc. Let's say that the precedence relation $<$ holds between t and s iff t is smaller than s . So $0 < 1$ and $1 < 25$. (We could just as well have stipulated that $<$ holds between t and s iff t is greater than s ; we would then

have $1 < 0$ and $25 < 1$. In temporal logic, the symbol ' $<$ ' means 'earlier than', not 'smaller than'.) Finally, let's say that the interpretation function assigns to p the set of all even numbers.

Let's call this model M . By definition 7.2, we can figure out the following facts, among others.

- $M, 0 \models p$ (because 0 is even);
- $M, 0 \models Fp$ (because there are even numbers greater than 0);
- $M, 0 \models GFp$ (because for every number there is a greater number that is even);
- $M, 0 \models \neg FGp$ (because there is no number for which all greater numbers are even).

Exercise 7.2

Now let M be the following model. As before, T is the set of natural numbers $\{0, 1, 2, \dots\}$, and $t < s$ iff t is smaller than s . This time, $V(p)$ is the set of numbers smaller than 10. Which of the following statements are true?

- (a) $M, 0 \models Fp \wedge F\neg p$
- (b) $M, 0 \models G\neg p$
- (c) $M, 0 \models FG\neg p$
- (d) $M, 0 \models GFp$
- (e) $M, 0 \models G(Fp \rightarrow FFp)$
- (f) $M, 0 \models FHp$
- (g) $M, 0 \models \neg P(p \vee \neg p)$
- (h) $M, 0 \models Hp$

Real times are, of course, not numbers. When I say that 'it is raining' is true now, I don't mean that the sentence is true at a number. It isn't obvious what kinds of things times are. Fortunately, this doesn't matter for us, just as the nature of possible worlds doesn't matter for the logic of possibility and necessity. As long as the formal structure of the times in a scenario matches the structure of the natural numbers, it does no harm to use numbers as times in a model of the scenario.

The formal structure of time in a temporal model is captured by the relevant frame: the pair $\langle T, < \rangle$ of the set of times and the precedence relation. Frames in temporal logic are also called **flows of time**. Different applications of temporal logic often come with different assumptions about the flow of time.

In computer science, for example, the “times” T are often understood as possible states of a computational process; the precedence relation holds between states t and s iff the computation can lead from t to s . If the computation is indeterministic, so that a given state can have different successors, the relevant flow of time will involve forks towards the future: we can have different “times” s and r such that $t < s$ and $t < r$ but neither $s < r$ nor $r < s$. Here the precedence relation cannot be modelled by the less-than relation on the natural numbers, because the structure of the less-than relation does not include forks.

In other applications, we may be interested in how the weather changes from day to day. Here we might identify the relevant times with days and the precedence relation with the earlier-relation between days – even though intuitively a day is not a single time, but an interval comprising many times. For this application, the natural numbers might have the right formal structure.

For yet other applications, we may want to assume that time is **dense**, meaning that whenever $t < s$ then there is another point of time lying in between t and s . This assumption is common in physics. The natural numbers, by contrast, have a **discrete** structure. There is no natural number in between 2 and 3. For dense models, we could use real or rational numbers (fractions) instead of natural numbers.


If we want to take seriously what physics tells us about time, it is not enough to assume that time is dense. We also need to reconceptualize the set T . According to the theory of special relativity, whether a point in time is earlier or later than another is relative to a spatial frame of reference. An adequate model of relativistic time must therefore include a representation of space. In these **spacetime models** (or *Minkowski models*), the set T consists of spacetime points $\langle x_1, x_2, x_3, t \rangle$ with three spatial and one temporal coordinate; $\langle x_1, x_2, x_3, t \rangle < \langle y_1, y_2, y_3, s \rangle$ holds iff the second point can be reached from the first without travelling faster than the speed of light.

7.3 Logics of time

Let’s define the minimal temporal logic K_t as the set of \mathcal{L}_t -sentences that are true at all times in all temporal models. Since temporal models are just Kripke models, proof methods for the minimal modal logic K are easily adapted to K_t . The main novelty is that the rules for the box and the diamond can be used twice over, once for

the forward-looking operators F and G, and once for the backward-looking P and H.

In the tree method for K_t , we have all the K-rules, with G as the box and F as the diamond. In addition, we have rules for H as the box and P as the diamond with a reversed perspective on the accessibility (or precedence) relation:



In the axiomatic approach, we have two versions of the (K) schema, one for the forward-looking box G and one for the backward-looking box H:

$$(GK) \quad G(A \rightarrow B) \rightarrow (GA \rightarrow GB)$$

$$(HK) \quad H(A \rightarrow B) \rightarrow (HA \rightarrow HB)$$

We also have two versions of Necessitation, and two versions of (Dual):

$$(GDI) \quad \neg FA \leftrightarrow G \neg A$$

$$(HDI) \quad \neg PA \leftrightarrow H \neg A$$

(GNec) If A occurs in a proof, GA may be appended.

(HNec) If A occurs in a proof, HA may be appended.

In addition, we need two interaction principles, reflecting the fact that the accessibility relation for F and G is the inverse of the accessibility relation for P and H :

$$(\text{Con1}) \quad A \rightarrow G P A$$

$$(\text{Con2}) \quad A \rightarrow H F A$$

These axioms and rules, added to those of classical propositional logic, define an axiomatic calculus that is sound and complete for K_t . (Completeness is easily proved with the canonical model technique.)

Exercise 7.3

Show with the help of definition 7.2 that all instances of (Con1) and (Con2) are true at all times in all temporal models.

Exercise 7.4

Give K_t -tree proofs for the following schemas.

- (a) $A \rightarrow G P A$
- (b) $A \rightarrow H F A$
- (c) $F A \rightarrow H F F A$
- (d) $P G A \rightarrow P F A$
- (e) $H A \leftrightarrow H F H A$

For most applications, K_t is too weak. We will want to impose further restrictions on the relevant temporal models. For example, definition 7.1 allows for cases in which $t < s$ and $s < r$ without $t < r$. But if a time t is earlier than s , and s is earlier than r , then surely t must be earlier than r . For almost every application of temporal logic, we assume that the precedence relation is transitive. This corresponds to the (4)-schema for G . It also corresponds to the (4)-schema for H .

$$(4G) \quad G A \rightarrow G G A$$

$$(4H) \quad H A \rightarrow H H A$$

Exercise 7.5

Explain why, if a relation $<$ is transitive, then so is its converse. The converse $>$ of $<$ is the relation that holds between x and y iff $y < x$.

Another plausible condition is that no time is earlier than itself. Formally, $<$ should be *irreflexive*, so that no element of T is $<$ -related to itself. We know that reflexivity corresponds to the (T)-schema, whose (forward-looking) temporal analogue would be $\text{G } A \rightarrow A$. What corresponds to irreflexivity? The following observation reveals the answer: nothing.

Observation 7.1: A sentence is valid in the class of irreflexive frames iff it is valid in the class of all frames.

Proof sketch: The right-to-left direction is obvious. The left-to-right direction is implied by the answer to exercise 4.8. But we can give a more direct argument.

Suppose that some sentence A is not valid in the class of all frames. We show that A is not valid in the class of irreflexive frames. That A is not valid in the class of all frames means that there is some world w in some model $M = \langle W, R, V \rangle$ at which A is false. We will show that there is some world in some irreflexive model at which A is false.

To this end, we will construct an irreflexive model $M^i = \langle W', R', V' \rangle$ from M in which the same sentences are true at w as in M . Since A is true at w in M , it follows that A is true at w in M^i .

Initially, M^i has the same worlds, the same accessibility relation, and the same interpretation function as M . Now for any world w in M that can see itself, we add a new world w' to M^i so that

- w' verifies the same sentence letters as w : if $w \in V(P)$ then $w' \in V(P)$;
- w' can see the same worlds as w : whenever $wR'v$ then $w'R'v$; and
- w' can be seen from the same worlds as w : whenever $vR'w$ then $vR'w'$.

Finally, we make w inaccessible from itself in M^i . A simple proof by induction on complexity shows that if a sentence is true at a world w in M then it is also true at w in M^i . \square

Given transitivity, irreflexivity is closely related to asymmetry. Recall from the previous chapter that $<$ is asymmetric if whenever $t < s$ then not $s < t$. There is no modal schema that corresponds to asymmetry.


Exercise 7.6

Show that a transitive relation is irreflexive iff it is asymmetric.

Exercise 7.7

A popular idea in many cultures is that time is circular. Does this cast doubt on asymmetry? What about irreflexivity?

In the previous chapter, I mentioned that transitive and irreflexive relations are called (strict) partial orders. The name reflects the fact that such orders need not order everything. In a model of branching time, for example, we can have $t < s$ and $t < r$ but neither $s < r$ nor $r < s$; in that case, r and s are not ordered by the precedence relation.



We can rule out such cases by imposing the requirement of **connectedness**, also known as *completeness* or *totality*. This demands that for any points t and s , either $t < s$ or $t = s$ or $s < t$. An irreflexive, transitive, and connected relation is called a **(strict) linear order** (or a *strict total order*).

For some applications, we may want linearity in only one direction. Many philosophers have been attracted to a branching-future conception of time, where a point in time may have more than one future, but only one past. In such models, we would only require **left-linearity**: that if $s < t$ and $r < t$, then either $s < r$ or $s = r$ or $r < s$.

An axiom schema corresponding to left-linearity is (LL):


$$(LL) \quad FPA \rightarrow (FA \vee A \vee PA)$$

Right-linearity – the assumption that if $t < s$ and $t < r$, then either $s < r$ or $s = r$ or $r < s$ – corresponds to (RL):


$$(RL) \quad PFA \rightarrow (PA \vee A \vee FA)$$

The conjunction of (LL) and (RL) is valid on a frame iff the frame's precedence relation does not branch in either direction. This is not quite the same as connectedness, because it allows for frames with parallel time lines. There is no schema that corresponds to connectedness.

The tree rules for left-linearity and right-linearity directly reflect the definition of the two properties.



These rules create *three* branches. They also create “identity nodes” of the form $\nu = \nu$, stating that two world/time labels refer to the same thing. (This must be taken into account when we read off a countermodel from an open branch.) We need two further rules to deal with identity nodes. Both of these rules are called ‘Identity’.



Exercise 7.8

Use the tree method to check which of the following sentences are valid, assuming time is linear (i.e., using the Transitivity, Left-Linearity, Right-Linearity, and Identity rules).

- (a) $(F p \wedge F q) \rightarrow F(p \wedge q)$
- (b) $P G G p \rightarrow G G p$
- (c) $P F p \rightarrow (P p \vee (p \vee F p))$

- (d) $\text{P H}p \rightarrow \text{H}p$
- (e) $\text{F G}p \rightarrow \text{G F}p$
- (f) $\text{F}(\text{G}q \wedge \neg p) \rightarrow \text{G}(p \rightarrow (\text{G}p \rightarrow q))$

The precedence relation in relativistic spacetime is neither left-linear nor right-linear. But it has a weaker property: convergence. A spacetime point p_1 can precede two points p_2 and p_3 neither of which precedes the other, but these two points will always precede a common later point p_4 . Convergence corresponds to the (G)-schema. In temporal logic, we have one (G)-schema for future convergence and one for past convergence:

- (FG) $\text{F G}A \rightarrow \text{G F}A$
- (PG) $\text{P H}A \rightarrow \text{H P}A$

Exercise 7.9

Can you find schemas that correspond to the following frame properties?

- (a) There is no last time. (That is, every time precedes some time.)
- (b) There is no first time.
- (c) There is a last time.
- (d) There is a first time.

Exercise 7.10

Show that the schema $\text{F}A \rightarrow \text{F F}A$ corresponds to density. (You have to show that (a) whenever a frame is dense then $\text{F}A \rightarrow \text{F F}A$ is valid on the frame, and (b) whenever $\text{F}A \rightarrow \text{F F}A$ is valid on a frame then the frame is dense.)

Exercise 7.11

Can you find an \mathcal{E}_T -expression stating that p is true at all times? Can you do so if you make assumptions about the precedence relation?

7.4 Branching time

In section 1.5 we looked at the idea that the future is “open” while the past is “settled”, insofar as we can still influence (say) whether we will exercise tomorrow, but not whether we have exercised yesterday. Some have argued that this calls for a non-linear model of time, with multiple branches into the future. On one branch, we would exercise tomorrow, on another we would not.

This line of thought appears to conflate temporal and modal considerations. The precedence relation in models of time is normally understood as a purely temporal relation – as the earlier-later relation. The fact that we can bring about a world in which we exercise tomorrow and a world in which we don’t exercise does not entail that both kinds of tomorrow take place here in the actual world.

If we want to make explicit the connections between settledness and time, it is better to use a multi-modal language with circumstantial operators for settledness and openness in addition to the purely temporal operators F, G, P, H . We could then say things like $Pp \rightarrow \Box Pp$ to formalize the claim that if p has happened then it is settled that p has happened.

Exercise 7.12

Suppose we endorse all instances of the schema (S1) $PA \rightarrow \Box PA$. Suppose we also endorse all instances of (S2) $\neg PA \rightarrow \Box \neg PA$, on the grounds that if something has failed to happen then there is nothing we can do that would make it have happened. Let’s further assume that the present time is not the first, and that the box is closed under logical consequence, meaning that if $\Box A_1, \dots, \Box A_n$ are true at a time, and B is entailed by A_1, \dots, A_n , then $\Box B$ is true (at the time) as well. Show that we can then derive the fatalist conclusion that anything that never actually happen is settled to never happen: all instances of $(\neg PA \wedge \neg A \wedge \neg FA) \rightarrow \Box \neg FA$ are true. (Hint: use instances of (S1) and (S2) in which A is a statement about the future.)

There are nonetheless good reasons to consider branching models of time. I already mentioned that such models are widely used in computer science, where the “times” represent states of a computational process and the precedence relation has a semi-modal interpretation, holding between two states iff the first can lead to the second. I also mentioned that the precedence relation in relativistic spacetime al-

allows for branching, although diverging spacetime branches ultimately reconverge. A more classical form of branching (without reconvergence) has been argued to follow from a certain interpretation of quantum physics. On this interpretation, what are normally understood to be chance events are really branching events in which all possible outcomes actually take place.

Another way to motivate a branching conception of time arises from a metaphysical view called *presentism*. According to presentism, only the present is real; all truths that seem to concern other times are reducible to more fundamental truths about the present. If, for example, it is true that there was a sea battle yesterday, then according to presentism this must ultimately be explained by what is true *now*; there must be facts about the present state of the world that entail (and explain) yesterday's sea battle. Different forms of presentism disagree over what the relevant facts about the present might be. On one view, they are particular facts about the distribution of physical particles and fields etc. together with the general laws of nature. If the laws of nature are deterministic, then the complete truth about the present distribution of particles and fields etc. together with the laws fixes all truths about the past and about the future. But suppose the laws are indeterministic towards the future: they merely settle that if the present physical state of the world is so-and-so, then the future is *either like this or like that*. In that case, the presentist will regard both of these futures as equally actual.


Let's assume, then, that we want to reason about branching time. This is less straightforward than it might at first appear.

Two pieces of terminology will be useful. First, let's define a **history** in a model $\langle T, <, V \rangle$ as a maximal linearly ordered subset of T . That is, a history is a collection of times H such that

- (i) for all t and s in H , either $t < s$ or $t = s$ or $s < t$, and
- (ii) no further member of T could be added to H without making (i) false.

The model (or rather, frame) depicted on the right contains two histories: $\{t_1, t_2, t_3\}$ and $\{t_1, t_2, t_4\}$.

For the second piece of terminology, let t be any time in any model. Any maximal linearly ordered set of times *later than* t will be called a **future of** t . In the model on the right, t_1 has two futures: $\{t_2, t_3\}$ and $\{t_2, t_4\}$.



If you look back at definition 7.2, you can see that in the standard semantics for temporal logic, Gp is true at t iff p is true at all times *in all futures of t* ; Fp , on the other hand, is true at t iff p is true at some time *in at least one future of t* . This ensures that G and F are duals, but it is often thought to be problematic if we want Fp to translate ‘it will be the case that p ’.

To illustrate, suppose I’m about to toss a coin. In one future (let’s assume), the coin will land heads, in another it will land tails. By definition 7.2, both Fh and Ft are true. But should we say that the coin will land heads and also that it will land tails?

We could adopt an alternative semantics for F according to which Fp is true at t iff p is true at some time in *all* futures of t :

$$M, t \models FA \text{ iff every future of } t \text{ contains some } s \text{ such that } M, s \models A.$$

This is known as the **Peircean interpretation** of F (after Charles S. Peirce; the name is due to Arthur Prior).

On the Peircean account, Fp is false whenever p only takes place in one of several futures. If we keep the classical interpretation of G , both Fp and $G\neg p$ can be false; the two operators are no longer duals. The dual of F is a strange operator that applies to a sentence A iff there is *some* future in which A is always true.

Exercise 7.13

Explain why the Peircean interpretation renders $p \rightarrow HFp$, an instance of (Con2), invalid.

A rather different approach is taken by (what Prior called) the **Ockhamist** approach. According to Ockhamism, if there are several futures then it doesn’t make sense to say – without qualification – that p will be the case, or that p won’t be case. To talk about what will or won’t be the case we must specify which future we have in mind.

Formally, in Ockhamist semantics, the truth-value of every sentence is evaluated at a pair consisting at a time and a history. Histories are linear by definition, so the problems raised by multiple futures disappear. To say that p is the case in *some* history, or in *all* histories, Ockhamists add new operators \Diamond and \Box that quantify over histories. The Peircean F operator is equivalent to $\Box F$ in Ockhamism. $\Box Fp$ says

that every future contains a time at which p is true; $\Diamond Fp$, by contrast, would say that some future contains a time which p is true.

Here is the full Ockhamist semantics.

Definition 7.3: Ockhamist Semantics

If $M = \langle T, <, V \rangle$ is a temporal model, H is a history in M , t is a member of H , P is any sentence letter, and A, B are any sentences in the Ockhamist language, then

- (a) $M, H, t \models P$ iff t is in $V(P)$.
- (b) $M, H, t \models \neg A$ iff $M, H, t \not\models A$.
- (c) $M, H, t \models A \wedge B$ iff $M, H, t \models A$ and $M, H, t \models B$.
- (d) $M, H, t \models A \vee B$ iff $M, H, t \models A$ or $M, H, t \models B$.
- (e) $M, H, t \models A \rightarrow B$ iff $M, H, t \not\models A$ or $M, H, t \models B$.
- (f) $M, H, t \models A \leftrightarrow B$ iff $M, H, t \models (A \rightarrow B)$ and $M, H, t \models (B \rightarrow A)$.
- (g) $M, H, t \models FA$ iff $M, H, s \models A$ for some s in H such that $t < s$.
- (h) $M, H, t \models GA$ iff $M, H, s \models A$ for all s in H such that $t < s$.
- (i) $M, H, t \models PA$ iff $M, H, s \models A$ for some s in H such that $s < t$.
- (j) $M, H, t \models HA$ iff $M, H, s \models A$ for all s in H such that $s < t$.
- (k) $M, H, t \models \Box A$ iff $M, J, t \models A$ for all histories J that contain t .
- (l) $M, H, t \models \Diamond A$ iff $M, J, t \models A$ for some history J that contains t .

A sentence is *valid* in Ockhamist semantics if it is true at all times t on all histories H (containing t) in all models. As always, we can get stronger conceptions of validity – stronger logics – by adding further constraints on the precedence relation.

Exercise 7.14

Which of the following schemas are valid in Ockhamist semantics?

- (a) $\Box A \rightarrow A$
- (b) $\Box A \rightarrow \Box \Box A$
- (c) $\Diamond A \rightarrow \Box \Diamond A$
- (d) $\Box FA \rightarrow F \Box A$
- (e) $PA \rightarrow \Box P \Diamond A$

There is something odd about the Ockhamist approach. Consider a scenario in which there are multiple futures; one future holds a sea battle, another holds no sea battle. Let p translate ‘there is a sea battle’. Is Fp true in this scenario (under the given interpretation of p)? What about $F(p \vee \neg p)$? Or $Gp \rightarrow GGp$?

Ockhamism refuses to give an answer. In Ockhamism, sentences are only true or false relative to a model and a time *and a history*. A branching-time scenario, however, does not fix a particular history. We’d like to know which sentences are true today if there are multiple futures. Ockhamism only tells us which sentences are true relative to each of the different futures. Relative to a history that contains a sea battle, Fp is true. Relative to other histories, Fp is false.

If we insist that logical validity should formalize the idea of truth in all scenarios under all interpretations of non-logical vocabulary then we can’t accept the official definition of validity in Ockhamist semantics. We have to extend the Ockhamist semantics to specify under what conditions a sentence is true *in a model at a time*, without fixing a history. Then we can say that a sentence is valid iff it is true at all times in all models.

A simple way to do this is to stipulate that a sentence is true at time in an (Ockhamist) model iff it is true relative to *all* histories that contain the time:

$$M, t \models A \text{ iff } M, H, t \models A \text{ for all histories } H \text{ that contain } t.$$

This is known as a **supervaluationist** semantics.

Supervaluationism is often used when a formal semantics defines truth relative to an “extra” parameter that doesn’t correspond to any feature of a conceivable scenario. In Ockhamist semantics, that parameter is H . For a different application, consider vagueness. If p translates ‘it is warm’, and the temperature is borderline warm, it is not clear what we should say about the truth-value of p , and about various complex sentences containing p . One popular approach to vagueness is to first define truth relative to a *sharpening* of vague expressions. Relative to a sharpening on which temperatures above 15.0 degrees Celsius are warm, p has a clear truth-value in any conceivable scenario, as do complex sentences containing p . Since an actual scenario does not fix a particular sharpening, this semantics contains an extra parameter. We can define a notion of truth without that parameter by saying that a sentence is true in a scenario iff it is true in that scenario relative to every eligible sharpening.

Supervaluationist accounts tend to have some non-classical features. Suppose we live in a branching world in which one future contains a sea battle and another doesn't. Let p express that a sea battle takes place. According to supervaluationist Ockhamism, neither Fp nor $\neg Fp$ is true in that scenario. Both are true relative to some but not relative to all histories. So neither is simply true. Assuming that a sentence is *false* if its negation is true, Fp is neither true nor false!

Logics in which a sentence can have a third status besides (mere) truth and (mere) falsity are called **three-valued**. Three-valued approaches to branching time are sometimes defended by the intuition that if a sea battle occurs on some but not all branches of the future, then one can't truly assert that a battle *will* occur nor that it *won't* occur.

The Polish logician Jan Łukasiewicz argued that statements about the future are either true, false, or "indeterminate". To accommodate this third truth-value, he proposed three-valued truth-tables specifying how the truth-value of complex sentences are determined by the truth-value of their parts. For example, he suggested that if two sentences A and B are indeterminate, then their conjunction $A \wedge B$, disjunction $A \vee B$, and negations $\neg A$, $\neg B$ are also indeterminate.

In the sea battle scenario, Łukasiewicz's account renders $Fs \vee \neg Fs$ indeterminate, assuming Fs is indeterminate. This is often regarded as problematic: even if we shouldn't assert that there will be a sea battle, it is argued that we are justified to assert that there either will or there won't be a sea battle. The supervaluationist form of Ockhamism, while also three-valued, avoids this problem. On the supervaluationist interpretation, Fs and $\neg Fs$ are neither true nor false in the sea battle scenario, but $Fs \vee \neg Fs$ is true.

Exercise 7.15

Let's say that a sentence is *super-valid* if it is true at all times in all models, where truth at a time in a model is understood in accordance with supervaluationist Ockhamism. Explain why the super-valid sentences are precisely the sentences that are valid by the original Ockhamist definition of validity (just below definition 7.3).

Exercise 7.16

Things are more complicated for entailment. Let's say that A *Ockham-entails* B iff there is no time on any history in any temporal model at which A is true and B false. Let's say that A *super-entails* B iff there is no time in any temporal model at which A is true and B false, where truth at a time in a model is defined in accordance with supervaluationism. Is Ockham-entailment equivalent to super-entailment? Explain.

7.5 Extending the language

The expressive resources of standard modal and temporal logic are weak. There are many things we might want to say about the unfolding of events in time that can't be said with F , G , P , and H . The Ockhamist history quantifiers are one way of adding expressive power to the basic language of temporal logic. In this section, we will look at some others.

A useful operator for logics of discrete and linear time is the “next” operator X (also written ‘ \circ ’). Informally, XA means that A is true at the next point in time. Formally:

$$M, t \models XA \text{ iff } M, s \models A \text{ for some } s \text{ such that (i) } t < s \text{ and (ii) } s < r \text{ for all } r \text{ such that } r \neq s \text{ and } t < r.$$

With the help of X , we can also say that A is true in two units of time (XXA), in three units of time ($XXXA$), and so on. The corresponding operator for talking about the *previous* point in time is usually written Y .

A more powerful extension of \mathcal{L}_T adds binary operators for “since” and “until”, which can be used to translate sentences like (1) and (2).

- (1) Ever since we left the house it has been raining.
- (2) It will be raining until we go back inside.

Informally, $A S B$ is true iff B was true at some time in the past and A has always been true since then; $A U B$ is true iff B will be true at some time in the future and A will always be true until then. Formally:

$M, t \models A \text{ S } B$ iff there is some s with $s < t$ for which $M, s \models B$, and for all r with $s < r < t$, we have $M, r \models A$.

$M, t \models A \text{ U } B$ iff there is some s with $t < s$ for which $M, s \models B$, and for all r with $t < r < s$, we have $M, r \models A$.

The operators F, G, P, and H can all be defined in terms of S and U. For example, PA is equivalent to $(p \vee \neg p) \text{ S } A$. And FA is equivalent to $(p \vee \neg p) \text{ U } A$.

Exercise 7.17

Define XA in terms of U.

Another noteworthy addition to temporal logic is the “Now” operator N. To see the point of this operator, consider the following multi-modal statement.

(3) We already knew yesterday that there would be a test today.

Using Y for ‘yesterday’, we might try to translate (3) as YKp , where p translates ‘there is a test’. But that’s wrong. By the semantics for Y, YKp is true today iff Kp is true yesterday (using days as temporal units). Since Kp entails p , it follows that YKp is true today only if p is true *yesterday*. But the test takes place today, not yesterday.

Intuitively, the problem is that ‘today’ in (3) refers to the present day, even though it occurs in the scope of the ‘yesterday’ operator. The same thing happens in the quantified statement (4).

(4) One day everyone who is now rich will be poor.

Here, ‘now’ refers to the present time, even though it is in the scope of the F operator ‘one day’.

With the “Now” operator N, we can translate (3) as $YKNp$, and (4) as $F\forall x(NRx \rightarrow Px)$. (We will have a closer look at quantified modal logic in later chapters.)

Intuitively, the N operator allows us to look outside the scope of an embedding operator. PNp , for example, is true if there is some time in the past such that p is true not at that time, but at the present. How does this work formally?

By the semantics of P,

$$M, t \models P N p \text{ iff } M, s \models N p \text{ for some time } s < t.$$

Now we want $M, s \models N p$ to be true iff p is true at the original time t . So we need to keep track of the original time at which we evaluate a sentence, even if a temporal operator shifts the time at which a subsentence is evaluated.

The simplest way to achieve this is to define truth relative to pairs of times. One of the times is shifted by the temporal operators, the other is held fixed.

Definition 7.4: Two-Dimensional Temporal Semantics

If $M = \langle T, <, V \rangle$ is a temporal model, t, t_0 are members of T , P is any sentence letter, and A, B are any \mathcal{L}_T -sentences, then

- (a) $M, t_0, t \models P$ iff t is in $V(P)$.
- (b) $M, t_0, t \models \neg A$ iff $M, t_0, t \not\models A$.
- (c) $M, t_0, t \models A \wedge B$ iff $M, t_0, t \models A$ and $M, t_0, t \models B$.
- (d) $M, t_0, t \models A \vee B$ iff $M, t_0, t \models A$ or $M, t_0, t \models B$.
- (e) $M, t_0, t \models A \rightarrow B$ iff $M, t_0, t \not\models A$ or $M, t_0, t \models B$.
- (f) $M, t_0, t \models A \leftrightarrow B$ iff $M, t_0, t \models (A \rightarrow B)$ and $M, t_0, t \models (B \rightarrow A)$.
- (g) $M, t_0, t \models F A$ iff $M, t_0, s \models A$ for some s in T such that $t < s$.
- (h) $M, t_0, t \models G A$ iff $M, t_0, s \models A$ for all s in T such that $t < s$.
- (i) $M, t_0, t \models P A$ iff $M, t_0, s \models A$ for some s in T such that $s < t$.
- (j) $M, t_0, t \models H A$ iff $M, t_0, s \models A$ for all s in T such that $s < t$.
- (k) $M, t_0, t \models N A$ iff $M, t_0, t_0 \models A$.

Like the Ockhamist semantics from the previous section, this semantics has an extra parameter. An ordinary scenario is represented by a single time in a model, not by a pair of times. So we need to specify under what conditions a sentence is true at a (single) time. Here, the standard approach is not supervaluation but “diagonalization”:

$$M, t \models A \text{ iff } M, t, t \models A.$$

This “two-dimensional” semantics correctly predicts that $P N p$ entails p .

1. Assume $M, t \models P N p$.
2. Then $M, t, t \models P N p$, by the definition of truth at a time in a model.
3. Then $M, t, s \models N p$ for some $s < t$, by clause (i) of definition 7.4.
4. Then $M, t, t \models p$, by clause (k) of definition 7.4.
5. Then $M, t \models p$, by the definition of truth at a time in a model.

The presence of a “Now” operator has far-reaching consequences for the logic of time. For example, $Np \rightarrow p$ is valid, in the sense that it is true at all times in all models. But $G(Np \rightarrow p)$ is invalid. If p is true at t and false at some time after t , then $G(Np \rightarrow p)$ is false at t . So we must give up the forward and backward Necessitation rules. The fact that something is logically true does not entail that it will always be true!

Exercise 7.18

‘It might have been that everyone who is actually rich is poor.’ This says that there is a world w such that everyone who is rich *at the actual world* is poor *at w* . To formalize statements like these, we need a modal operator analogous to N that takes us back to the actual world, even in the scope of other modal operators. This operator is called the *actually* operator. Let’s write it as A and add it to \mathcal{L}_M . Can you find a sentence B in this language that is logically true but not necessarily true, in the sense that B is true at all worlds in all models but $\Box B$ is not?