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 r expresses that it is raining, we will use $\mathsf{F}\,r$ to express that is will be raining, at some point in the future. We will use $\mathsf{P}\,r$ to express that it has been raining, at some point in the past. In general:

> $\mathsf{F}\,A$ is true at a time t iff A is true at some time after t.
> $\mathsf{P}\,A$ is true at a time t iff A is true at some time before t.

The temporal operators F and P can be nested. For example, we can use $\mathsf{F}\,\mathsf{P}\,r$ to express that at some point it will have rained. $\mathsf{P}\,\mathsf{F}\,r$ means that it was once going to rain, $\mathsf{P}\,\mathsf{P}\,r$ that there was a time before which it rained, and $\mathsf{F}\,\mathsf{F}\,r$ that there will come a time after which it will rain.

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

Intuitively, $\mathsf{G}\,A$ means that A is always *g*oing to be the case. (Hence the symbol 'G'.) For example, if it is not the case that at some point in the future it will not rain ($\neg\,\mathsf{F}\,\neg r$), then it is always going to be the case that it will rain ($\mathsf{G}\,r$). Similarly, $\mathsf{H}\,A$ means that A *h*as always been the case. If it is not the case that at some point in the past it was not raining ($\neg\,\mathsf{P}\,\neg r$), then it has always been raining ($\mathsf{H}\,r$).

We can state the truth-conditions of $\mathsf{G}\,A$ and $\mathsf{H}\,A$ in parallel to the above truth-conditions for $\mathsf{F}\,A$ and $\mathsf{P}\,A$:

> $\mathsf{G}\,A$ is true at a time t iff A is true at all times after t.
> $\mathsf{H}\,A$ 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 $\mathsf{F}, \mathsf{P}, \mathsf{G}, \mathsf{H}$ is known as the **language of basic temporal logic**. I'll sometimes call it $\mathfrak{L}_t$ for brevity.

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

In the previous section, I have informally specified the truth-conditions of $\mathsf{F}\,A$, $\mathsf{P}\,A$, $\mathsf{G}\,A$, and $\mathsf{H}\,A$ by quantifying over times: $\mathsf{F}\,A$ is true at a time iff A is true at all later times, and so on. This suggests that the truth-value of every $\mathfrak{L}_T$-sentence at any time in any conceivable scenario is fixed by the truth-value of the sentence letters at each time. Once you know at which times p, q, r, etc. are true, you can figure out the truth-value of every $\mathfrak{L}_T$-sentence at any time. For example, if p is true at some time after t, then $\mathsf{F}\,p$ is true at t, $\mathsf{G}\,\neg p$ is false at t, and $\mathsf{G}\,\mathsf{P}\,\mathsf{F}\,p$ is true at t.

For the purposes of temporal logic, we may therefore represent a scenario and an interpretation of the 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.

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 $\mathfrak{L}_T$ and each member of T a truth-value (1 or 0).

I use '$M, t \models A$' as a short-hand notation to express that the sentence A is true at time t in model M. The following definition fixes the truth-value of every $\mathfrak{L}_T$-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, ρ is any sentence letter, and A, B are any $\mathfrak{L}_T$-sentences, then

(a)	$M, t \models \rho$	iff $V(\rho, t) = 1$.
(b)	$M, t \models \neg A$	iff $M, t \not\models A$.
(c)	$M, t \models A \land B$	iff $M, t \models A$ and $M, t \models B$.
(d)	$M, t \models A \lor B$	iff $M, t \models A$ or $M, t \models B$.
(e)	$M, t \models A \to B$	iff $M, t \models B$ or $M, t \not\models A$.
(f)	$M, t \models A \leftrightarrow B$	iff $M, t \models (A \to B)$ and $M, t \models (B \to A)$.
(g)	$M, t \models \mathsf{F}\,A$	iff $M, s \models A$ for some $s \in T$ such that $t < s$.
(h)	$M, t \models \mathsf{G}\,A$	iff $M, s \models A$ for all $s \in T$ such that $t < s$.
(i)	$M, t \models \mathsf{P}\,A$	iff $M, s \models A$ for some $s \in T$ such that $s < t$.
(j)	$M, t \models \mathsf{H}\,A$	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.

We'll say that an $\mathfrak{L}_T$-sentence is **valid** iff it is true at every time in every (suitable) temporal model; an $\mathfrak{L}_T$-sentence B is a **logical consequence** of sentences $A_1, \ldots, A_n$ iff B is true at every time in every (suitable) model at which $A_1, \ldots, A_n$ are all true. I say "suitable" because we will want to put some constraints on the precedence relation $<$. More on that in a moment.

All this should remind you of our Kripke semantics for $\mathfrak{L}_M$ in chapter 3. In fact, temporal models *are* Kripke models, as defined on page 48. I have merely relabelled the set 'W' as 'T', and the relation 'R' as '$<$'.

Definition 7.2 resembles definition 3.2 from page 48, 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) which are not definable in terms of each other. However, 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 choose 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 $1 < 25$, for example. (Note that 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 the truth-value 1 (True) to p at a time t iff t is an even number.

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

- $M, 0 \models p$ (because $V(p, 0) = 1$);
- $M, 0 \models \mathsf{F}\, p$ (because $V(p, 2) = 1$ and $0 < 2$);
- $M, 0 \models \mathsf{G}\,\mathsf{F}\, p$ (because for every number there is a greater number that is even);
- $M, 0 \models \neg\,\mathsf{F}\,\mathsf{G}\, p$ (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 $\{0, 1, 2, \ldots\}$, and $t < s$ iff t is smaller than s. This time, $V(p, t) = 1$ iff $p < 10$. Which of the following statements are true?

(a) $M, 0 \models \mathsf{F}\, p \land \mathsf{F}\,\neg p$
(b) $M, 0 \models \mathsf{G}\,\neg p$
(c) $M, 0 \models \mathsf{F}\,\mathsf{G}\,\neg p$
(d) $M, 0 \models \mathsf{G}\,\mathsf{F}\, p$
(e) $M, 0 \models \mathsf{G}(\mathsf{F}\, p \to \mathsf{F}\,\mathsf{F}\, p)$
(f) $M, 0 \models \mathsf{F}\,\mathsf{H}\, p$
(g) $M, 0 \models \neg\,\mathsf{P}(p \lor \neg p)$
(h) $M, 0 \models \mathsf{H}\, p$

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 particular 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 relevant process can lead from state t to state s. If the process 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 another application, 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 would perhaps have the right formal structure.

For 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. For example, there is no 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; $(x_1, x_2, x_3, t) < (y_1, y_2, y_3, s)$ holds iff the second point can be reached from the first without travelling faster than the speed of light.

7.3 Logics of time

Since temporal models are just Kripke models, the standard proof systems for the minimal modal logic K also apply to temporal logic. They apply twice over, once for the forward-looking operators F and G, and once for the backward-looking P and H.

For example, the tree rules for minimal temporal logic are just the K-rules, but with the accessibility relation reversed for H and P.

$$
\begin{array}{cc}
\mathsf{G}\,A & (\omega)\\
\omega < \nu & \\
\vdots & \\
A & (\nu)
\end{array}
\qquad
\begin{array}{cc}
\mathsf{F}\,A & (\omega)\\
\vdots & \\
\omega < \nu & \\
A & (\nu)\\
 & \uparrow\\
 & \text{new}
\end{array}
\qquad
\begin{array}{cc}
\mathsf{H}\,A & (\omega)\\
\nu < \omega & \\
\vdots & \\
A & (\nu)
\end{array}
\qquad
\begin{array}{cc}
\mathsf{P}\,A & (\omega)\\
\vdots & \\
\nu < \omega & \\
A & (\nu)\\
 & \uparrow\\
 & \text{new}
\end{array}
$$

$$
\begin{array}{cc}
\neg\mathsf{G}\,A & (\omega)\\
\vdots & \\
\omega < \nu & \\
\neg A & (\nu)\\
 & \uparrow\\
 & \text{new}
\end{array}
\qquad
\begin{array}{cc}
\neg\mathsf{F}\,A & (\omega)\\
\omega < \nu & \\
\vdots & \\
\neg A & (\nu)
\end{array}
\qquad
\begin{array}{cc}
\neg\mathsf{H}\,A & (\omega)\\
\vdots & \\
\nu < \omega & \\
\neg A & (\nu)\\
 & \uparrow\\
 & \text{new}
\end{array}
\qquad
\begin{array}{cc}
\neg\mathsf{P}\,A & (\omega)\\
\nu < \omega & \\
\vdots & \\
\neg A & (\nu)
\end{array}
$$

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:

(**FK**) $\mathsf{G}(A \to B) \to (\mathsf{G}\,A \to \mathsf{G}\,B)$

(**BK**) $\mathsf{H}(A \to B) \to (\mathsf{H}\,A \to \mathsf{H}\,B)$

Similarly, we have two versions of the Necessitation rule:

(**FNec**) If A occurs in a proof, $\mathsf{G}\,A$ may be appended.

(**BNec**) If A occurs in a proof, $\mathsf{H}\,A$ 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.

(**Con1**) $A \to \mathsf{G}\,\mathsf{P}\,A$

(**Con2**) $A \to \mathsf{H}\,\mathsf{F}\,A$

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

Exercise 7.3

Show with the help of definition 7.2 that **Con1** and **Con2** are valid in the class of all temporal models.

Exercise 7.4

Show that **Con1** and **Con2** can be derived by the tree rules for minimal temporal logic.

The set of sentences that are valid at all times in all temporal models is sometimes called system K_t, because it is the temporal analogue of system K.

For most applications, this system is too weak. Not every model in the sense of definition 7.1 is an adequate model of time.

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 will therefore want to assume that the precedence relation is transitive. This makes the **4**-schema for G valid.

(**4G**) $\mathsf{G}\,A \to \mathsf{G}\,\mathsf{G}\,A$

Exercise 7.5

Transitivity of $<$ not only gives us the **4**-schema for G, but also the **4**-schema for H. Explain why.

Another plausible condition is that no time is earlier than itself; so $<$ should be *irreflexive*, meaning that we never have $t < t$.

We know that reflexivity corresponds to principle **T**, whose (forward-looking) temporal analogue would be $\mathsf{G}\,A \to A$. 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. For the left-to-right direction, suppose that some sentence A is not valid in the class of all frames. We need to 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 have to 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: for all ρ, $V'(\rho, w') = V'(\rho, w)$;
- 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. □

Observation 7.1 tells us that there is no modal principle that is valid in all and only the irreflexive frames. So the logic of irreflexive frames is the same as logic of all frames (system K). The proof carries over to many other classes of frames. For example, the logic of irreflexive and transitive frames is the same as the logic of transitive frames (namely, K4).

Given transitivity, irreflexivity is closely related to asymmetry. Recall from the previous chapter that $<$ is asymmetric if $t < s$ entails $s \not< t$. Asymmetry is also plausible if we interpret $<$ as the precedence relation between times. But as with irreflexivity, 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?

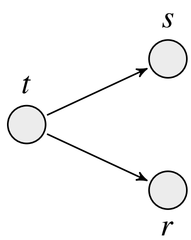

If $<$ is transitive and irreflexive (or equivalently: transitive and asymmetric), then it is a partial order. Partial orders are called "partial" because they don't necessarily order everything. For example, in a model of branching time we can have $t < s$ and $t < r$ but neither $s < r$ nor $r < s$; so r and s are not ordered by the precedence relation.

For many applications, we may want to rule out such cases, by imposing another requirement of **(strong) connectedness**. This demands that for any points $t, s \in T$, either $t < s$ or $t = s$ or $s < t$. (Note that this allows for cases where $t < s$ *and* $s < t$.) A relation that is irreflexive, transitive, and (strongly) connected is called a **linear order**.

For other applications, we may want linearity in only one direction. Many philosophers have been attracted to a branching-future conception of time, on which a given 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 $r < s$.

The axiom schema corresponding to left-linearity is **BL** (for "backwards-looking linearity”).

(**BL**) $\mathsf{F}\,\mathsf{P}\,A \rightarrow (\mathsf{F}\,A \lor A \lor \mathsf{P}\,A)$

Right-linearity similarly corresponds to **FL**.

(**FL**) $\mathsf{P}\,\mathsf{F}\,A \rightarrow (\mathsf{P}\,A \lor A \lor \mathsf{F}\,A)$

Right-linearity means that if $t < s$ and $t < r$, then either $s < r$ or $s = r$ or $r < s$. Together, left-linearity and right-linearity are equivalent to connectedness. So the conjunction of **BL** and **FL** corresponds to connectedness.

The tree rules for left-linearity and right-linearity are as you might expect from the definition of these two properties.

Left-Linearity

$$\begin{array}{ccc} & \nu < \omega & \\ & \upsilon < \omega & \\ & \vdots & \\ \nu < \upsilon & \nu = \upsilon & \upsilon < \nu \end{array}$$

Right-Linearity

$$\begin{array}{ccc} & \omega < \nu & \\ & \omega < \upsilon & \\ & \vdots & \\ \nu < \upsilon & \nu = \upsilon & \upsilon < \nu \end{array}$$

These rules create *three* branches. Also, they create nodes of the form $\nu = \upsilon$, stating that two world/time labels refer to the same thing. We need two further rules to deal with this kind of node. Both of these rules are called ‘Identity’.

$$\begin{array}{ll} A & (\omega) \\ \omega = \nu & \\ \vdots & \\ A & (\nu) \end{array} \qquad\qquad \begin{array}{ll} A & (\omega) \\ \nu = \omega & \\ \vdots & \\ A & (\nu) \end{array}$$

Exercise 7.8

Give tree proofs for the following schemas, assuming time is linear and transitive (i.e., using the Transitivity, Left-Linearity, Right-Linearity, and Identity rules).

(a) $\mathsf{P}\,\mathsf{G}\,A \to \mathsf{P}\,\mathsf{F}\,A$

(b) $\mathsf{P}\,\mathsf{G}\,\mathsf{G}\,A \to \mathsf{G}\,\mathsf{G}\,A$

(c) $\mathsf{P}\,\mathsf{F}\,A \to (\mathsf{P}\,A \lor (A \lor \mathsf{F}\,A))$

(d) $(\mathsf{F}\,A \land \mathsf{F}\,B) \to (\mathsf{F}(A \land B) \lor (\mathsf{F}(A \land \mathsf{F}\,B) \lor \mathsf{F}(\mathsf{F}\,A \land B)))$

(e) $\mathsf{F}(\mathsf{G}\,B \land \neg A) \to \mathsf{G}(A \to (\mathsf{G}\,A \to B))$

Exercise 7.9

Use the tree method to find countermodels for the following sentences, assuming time is linear, transitive, and serial.

(a) $(\mathsf{F}\,p \land \mathsf{F}\,q) \to \mathsf{F}(p \land q)$

(b) $\mathsf{F}\,p \to \mathsf{F}\,\mathsf{F}\,p$

(c) $\mathsf{P}\,\mathsf{H}\,p \to \mathsf{H}\,p$

(d) $\mathsf{F}\,\mathsf{G}\,A \to \mathsf{G}\,\mathsf{F}\,A$

The precedence relation in relativistic spacetime models is neither left-linear nor right-linear. But it has a weaker property that we already know: 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**) $\quad \mathsf{F}\,\mathsf{G}\,A \to \mathsf{G}\,\mathsf{F}\,A$

(**BG**) $\quad \mathsf{P}\,\mathsf{H}\,A \to \mathsf{H}\,\mathsf{P}\,A$

Exercise 7.10

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 7.11

Show that the schema $\mathsf{F}\,A \to \mathsf{F}\,\mathsf{F}\,A$ corresponds to density.

Exercise 7.12

Can you define *Always* and *Sometimes* in terms of F, G, P, and H? Can you do so if you make assumptions about the flow of time $<$?

7.4 Branching time

It is natural to think of the future as "open" in a way that the past is "closed". I might decide to go for a walk this afternoon, or I might decide to stay at home. So there appear to be multiple futures: in some I go for a walk, in others I stay at home. One might conclude that time is left-linear, but not right-linear.

But this argument is too quick. If I haven't decided what to do in the afternoon, then arguably there are several *possible* futures – several ways the worlds *might* evolve. But it is far from clear that there are several *actual* futures. If there were, it would make little sense to wonder what I will do, or to contemplate whether I should go for a walk or stay at home: I would end up doing both anyway, albeit on different temporal branches.

To formalize the above intuition we could use a multi-modal language with both temporal and circumstantial operators. The openness of the future could then be expressed by a statement like $\Diamond\,\mathsf{G}\,p \land \Diamond\,\mathsf{G}\,\neg p$. To prevent a corresponding openness of the past, the accessibility relation for the circumstantial diamond would have to hold fixed the past, so that a world v is accessible from a world w only if the past of v coincides with the past of w.

On the other hand, there are also views that assume a genuinely branching flow of time. Some of these are motivated by discoveries in 20th century physics. I already mentioned that the precedence relation in relativistic spacetime allows for branching, although all these branches ultimately reconverge. A more classical form of branching (without reconvergence) has been argued to follow from the so-called "Everett 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.

I already mentioned that branching time models are widely used in computer science, where the "times" represent states of a computational process and the precedence relation holds between two states iff the first can lead to the second.

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 talk about other times are reducible to more fundamental truths about the present. For example, if 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 which entail that (and explain why) there was a sea battle yesterday. On one form of presentism, the relevant facts about the present state of the world are (a) particular facts about the distribution of physical particles and fields etc., and (b) the general laws of nature. Many laws of nature are dynamic, specifying how closed physical systems evolve over time. So the laws might entail that if the present physical state of the world is so-and-so, then there was a sea battle yesterday. This is how facts about the present might entail that there was a sea battle in the past.

But now suppose the laws of nature 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 deny that exactly one of these futures is actual.

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

The models we are interested in are not right-linear. I will, however, assume that they satisfy the following weaker property:

$$\text{if } t < s \text{ then for any } r, \text{ either } t < r \text{ or } r < s.$$

This is sometimes called "weak connectedness" or "transitivity of nonprecedence", because it is equivalent to the assumption that if $t \not< s$ and $s \not< r$ then $t \not< r$. It slightly simplifies our models, for example by ruling out multiple parallel time lines which never connect to each other.

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

(a) for all t and s in H, either $t < s$ or $t = s$ or $s < t$, and

(b) no further member of T could be added to H without making (a) false.

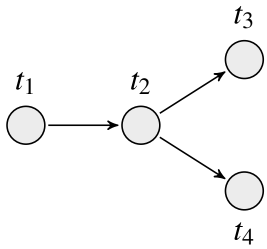

The model (or rather, frame) depicted on the right, for example, contains two histories: $\{t_1, t_2, t_3\}$ and $\{t_1, t_2, t_4\}$.

Second, if t is any time in any model, then any maximal linearly ordered set of times *later than t* will be called a **future of** t. In the example 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, $\mathsf{G}\,p$ is true at t iff p is true at all times *in all futures of t*; $\mathsf{F}\,p$, 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 may be regarded as problematic if we want $\mathsf{F}\,p$ to translate 'it will be the case that p'.

To illustrate, suppose I'm about to toss a coin. In one future, the coin will land heads, in another it will land tails. By definition 7.2, both $\mathsf{F}\,h$ and $\mathsf{F}\,t$ are true. But it is not clear if we should say that the coin will land heads and also that it will land tails.

One might therefore suggest an alternative semantics for F according to which $\mathsf{F}\,p$ is true at t iff p is true at some time in *all* futures of t:

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

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

On the Peircean account, $\mathsf{F}\,p$ is false if p only takes place in one of several futures. If we keep the classical interpretation of G, both $\mathsf{F}\,p$ and $\mathsf{G}\,\neg p$ can be false; so F and G 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 Peirceanism renders **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 $\Box$ which quantify over histories. So the Peircean F operator is equivalent to $\Box\,\mathsf{F}$ in Ockhamism. $\Box\,\mathsf{F}\,p$ is true if p will eventually be true at some time in every future; $\Diamond F p$, by contrast, would express that p is the case in some future.

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, ρ is any sentence letter, and A, B are any sentences in the Ockhamist language, then

(a)	$M, H, t \models \rho$	iff $V(\rho, t) = 1$.
(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 $M, H, t \models B$ or $M, H, t \not\models A$.
(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)	$M, H, t \models \mathsf{F}\,A$	iff $M, H, s \models A$ for some $s \in H$ such that $t < s$.
(h)	$M, H, t \models \mathsf{G}\,A$	iff $M, H, s \models A$ for all $s \in H$ such that $t < s$.
(i)	$M, H, t \models \mathsf{P}\,A$	iff $M, H, s \models A$ for some $s \in H$ such that $s < t$.
(j)	$M, H, t \models \mathsf{H}\,A$	iff $M, H, s \models A$ for all $s \in H$ such that $s < t$.
(k)	$M, H, t \models \Box A$	iff $M, H', t \models A$ for all histories H' that contain t.
(l)	$M, H, t \models \Diamond A$	iff $M, H', t \models A$ for some history H' that contains t.

A sentence is *valid* in Ockhamist semantics if it is true at all times on all histories in all models (ignoring histories that don't contain the relevant time). To my knowledge, no-one has yet found an axiomatic or tree method that can prove all and only the Ockhamistically valid sentences.

Exercise 7.14

Which of the following schemas are valid in Ockhamist semantics?

(a) $\Box A \to A$
(b) $\Box A \to \Box\Box A$
(c) $\Diamond A \to \Box\Diamond A$
(d) $\Box \mathsf{F} A \to \mathsf{F} \Box A$
(e) $\mathsf{P} A \to \Box \mathsf{P} \Diamond A$

Something is strange about the Ockhamist approach. Informally, a sentence is logically true if it is true in every conceivable scenario under every interpretation of the non-logical vocabulary. Now 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 $\mathsf{F}\,p$ is true in this scenario (under the given interpretation of p)? What about $\mathsf{F}(p \vee \neg p)$? Or $Gp \to GGp$?

Ockhamism doesn't say. In Ockhamism, sentences are only true or false relative to a model and a time *and a history*. Intuitively, however, a branching-time scenario does not include 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, $\mathsf{F}\,p$ is true, relative to other histories $\mathsf{F}\,p$ is false.

If we insist that logical validity should formalize the idea of truth in all scenarios under all interpretations of the non-logical vocabulary, then we can't accept the official definition of validity in Ockhamist semantics. We have to extend the Ockhamist semantics by specifying 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.

One simple strategy is to stipulate that a sentence is true at time in a 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 of various complex sentences containing p. In response, one option is to first define truth relative to a *sharpening* of vague expressions. Relative to a sharpening on which temperatures above 15 degrees Celsius are warm, p has a clear truth-value in any conceivable scenario, and so do complex sentences containing p. But actual scenarios arguably don’t fix a sharpening. Supervaluationists about vagueness therefore suggest that a sentence is true in a scenario iff it is true in that scenario relative to *all* possible sharpenings of the language.

Exercise 7.15

Explain why the very same sentences are valid on the original Ockhamist definition of validity and on the revised supervaluationist definition of validity.

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.

Supervaluationism is not the only possible way to define truth at a time in a model, without fixing a history. One alternative is to say that a sentence is true at a time in a model iff it is true relative to *some* history containing the time. Another interesting approach is to say that a sentence is true if it is true relative to all histories, false if it is false relative to all histories, and neither true nor false if it is true relative to some but not all histories. This leads to a **three-valued logic**: 'there will be a sea battle' is neither true nor false in a scenario in which one future holds a sea battle and other futures do not.

7.5 Extending the language

The expressive resources of standard modal 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 '$\bigcirc$'). Informally, $\mathsf{X}\,A$ means that A is true at the next point in time. Formally:

$$M, t \models \mathsf{X}\,A \text{ iff } M, s \models A \text{ for some } s \text{ such that (a) } t < s \text{ and (b) } s < r \text{ for all } r \text{ such that } t \neq s \text{ and } t < r.$$

With the help of X, we can also say that A is true in two units of time: $\mathsf{X}\,\mathsf{X}\,A$, or that A is true in three units of time: $\mathsf{X}\,\mathsf{X}\,\mathsf{X}\,A$, and so on. The corresponding operator for talking about the *previous* point in time is usually written Y.

A much more powerful extension of $\mathfrak{L}_T$ 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, $\mathsf{S}(A, B)$ is true iff A was true at some time in the past and B has always been true since then; $\mathsf{U}(A, B)$ is true iff A will be true at some time in the future and B will always be true until then. Formally:

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

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

If we have S and U, we don't actually need $\mathsf{F}, \mathsf{G}, \mathsf{P}$, and H, because these can be defined in terms of S and U: $\mathsf{P}\,A$ is equivalent to $\mathsf{S}(A, p \vee \neg p)$, and $\mathsf{F}\,A$ to $U(A, p \vee \neg p)$. Conversely, however, $\mathsf{S}(A, B)$ and $\mathsf{U}(A, B)$ are not expressible in $\mathfrak{L}_t$.

Exercise 7.17

Define $\mathsf{X}\,A$ 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) Bob already knew yesterday that there would be a test today.

Using Y for 'yesterday', we might try to translate (3) as $\mathsf{Y}\,\mathsf{K}_b\,p$, where p translates 'there is a test'. But that's wrong. By the semantics for Y, $\mathsf{Y}\,\mathsf{K}_b\,p$ is true today iff $\mathsf{K}_b\,p$ is true yesterday (using days as temporal units). And since knowledge is factive, if $\mathsf{K}_b\,p$ is true at some time, then p is true at that time. So $\mathsf{Y}\,\mathsf{K}_b\,p$ wrongly entails that there was a test *yesterday*.

Intuitively, the problem is that 'today' in (3) refers to the present day, even though it occurs in the scope of 'yesterday'. 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, (3) can be translated as $\mathsf{Y}\,\mathsf{K}_b\,\mathsf{N}\,p$, and (4) as $\mathsf{F}\,\forall x(\mathsf{N}\,Rx \to Px)$. (We will have a closer look at quantified modal logic in later chapters.)

The semantics of N raises a problem. Consider a simpler example, $\mathsf{P}\,\mathsf{N}\,p$. Since 'now' always picks out the present time, even when embedded under other temporal operators, $\mathsf{P}\,\mathsf{N}\,p$ should be equivalent to p (provided that the present time is not the first point in time). By the semantics of P,

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

Now we want $M, s \models \mathsf{N}\, p$ 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 standard way to achieve this is to define truth relative to *pairs* of times. One of the times is shifted by the modal 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, ρ is any sentence letter, and A, B are any $\mathfrak{L}_T$-sentences, then

(a)	$M, t_0, t \models \rho$	iff $V(\rho, t) = 1$.
(b)	$M, t_0, t \models \neg A$	iff $M, t_0, t \not\models A$.
(c)	$M, t_0, t \models A \land B$	iff $M, t_0, t \models A$ and $M, t_0, t \models B$.
(d)	$M, t_0, t \models A \lor B$	iff $M, t_0, t \models A$ or $M, t_0, t \models B$.
(e)	$M, t_0, t \models A \to B$	iff $M, t_0, t \models B$ or $M, t_0, t \not\models A$.
(f)	$M, t_0, t \models A \leftrightarrow B$	iff $M, t_0, t \models (A \to B)$ and $M, t_0, t \models (B \to A)$.
(g)	$M, t_0, t \models \mathsf{F}\, A$	iff $M, t_0, s \models A$ for some $s \in T$ such that $t < s$.
(h)	$M, t_0, t \models \mathsf{G}\, A$	iff $M, t_0, s \models A$ for all $s \in T$ such that $t < s$.
(i)	$M, t_0, t \models \mathsf{P}\, A$	iff $M, t_0, s \models A$ for some $s \in T$ such that $s < t$.
(j)	$M, t_0, t \models \mathsf{H}\, A$	iff $M, t_0, s \models A$ for all $s \in T$ such that $s < t$.
(k)	$M, t_0, t \models \mathsf{N}\, A$	iff $M, t_0, t_0 \models A$.

Like in Ockhamism, we also need to specify under what conditions a sentence is true (in a model) *at a time*, not at a pair of two times. Here, the standard approach is not supervaluation but "diagonalization":

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

Now we can show that $\mathsf{P}\,\mathsf{N}\, p$ is equivalent to p, provided there are earlier times. I will go through the left-to-right direction.

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

4. Then $M, t, t \models p$ for some $s < t$, by clause (k) of definition 7.4.
5. Then $M, t \models p$ for some $s < t$, 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, $\mathsf{N}\, p \rightarrow p$ is valid, in the sense that it is true at all times in all models, but $\mathsf{G}(\mathsf{N}\, p \rightarrow p)$ is invalid: if p is true at t and false at some time after t, then $\mathsf{G}(\mathsf{N}\, p \rightarrow p)$ is false at t. So we must give up the forward and backward Necessitation rules. The fact that something is logically valid no longer entails that it will always be the case.

Exercise 7.18

Suppose we add to the language of basic temporal logic with N another operator $\Box$ that applies to a sentence iff the sentence is valid. Explain why at least one of the following assumptions must then be false. Which of them do you think we should reject?

(a) Whenever $\Box A$, then $\mathsf{G}\,\Box A$.
(b) For any A, $\mathsf{G}(\Box A \rightarrow A)$.