PART I

Background

CHAPTER 1

The Symmetric Paradigm

1.1 The Symmetric Paradigm

From grade school to the higher reaches of tense logic, a fundamental paradigm prevails. Human languages offer up the ability to talk about past events, present ones, and future ones. According to the paradigm, in many human languages, this is implemented by a system of tenses with three broad categories: past, present, and future. With some exceptions and complications, these tenses line up with the temporal location of the appropriate events.[1]

The insight that past and future tenses have symmetric meanings is central to research frameworks that build on the classical framework for the logic of tense (Prior, 1957, 1967, 1969). The symmetric paradigm reverberates in many theoretical decisions concerning the syntax, the semantics, and logic of discourse about the future. When Prior (1967, p. 35) provides a series of postulates for "the logic of futurity," he immediately proceeds to lay down "a series of analogous postulates … to give the logic of pastness." Fast forward fifty years, and we find that von Fintel and Heim (2011) take as a starting point an approach on which the semantic entries for future and past tense (lifted from tense logic) look like "mirror images of each other, though they are of different syntactic categories" (p. 72).

This chapter develops the core idea of the symmetric paradigm. Eventually, I will critique this idea, with an eye toward developing my own, nonsymmetric account of future discourse. But before any cards are on the table, I want to clarify that my critique of the symmetric paradigm is limited in scope. Even if it turned out that the meanings of the expressions we use to talk about the future – the *predictive expressions*, as I will say – are of an entirely different sort from the meanings of the expressions we recruit to talk about the past, there is nothing inherently objectionable

1 While some languages, such as Mandarin Chinese, lack tense markers, it is nevertheless possible to characterize their devices of temporal reference symmetrically (Bittner, 2014).

about designing a system of *logic* with symmetric tenses. This is because studying the logic of tensed statements need not be part of a model of the meaning of tense in natural language. A formal temporal logic could be part of an attempt to regiment certain philosophical arguments and to clearly display their structural features. Alternatively, its purpose could be to state and compare rules for automated temporal reasoning – the kind of rules that we might want to instill in a computer carrying out reasoning tasks involving temporally structured information.[2]

A formal framework only becomes exposed to empirical evidence, and to the sort of argumentation I will build against the symmetric paradigm, if it is embedded in a theory with the ambition of predicting linguistic phenomena. This is the case for natural language semantics. These explanatory ambitions include developing a model of temporal discourse that systematizes and explains our judgments of acceptability (or unacceptability) of various speech acts in context as well as our judgments about the acceptability (or unacceptability) of various inference patterns.

With that said, even if Prior himself might not have had natural language applications on his mind, the interpretation of the tense logic framework as a module in a theory of meaning has nonetheless been influential. For example, simple tense logic is the default model of tense that is developed in two of the most influential textbooks in natural language semantics (Dowty et al., 1981, and in the notes in von Fintel and Heim, 2011) as well as in Richard Montague's influential essay "Pragmatics" (which can now be found in Montague, 1974, pp. 95–119). Even those who have objected to this application of tense logic have often landed on symmetric analyses (we will see some examples in due course).

Our first task, then, is to develop with more precision the hypothesis that the tenses of a natural language such as English are semantically symmetric. To do so, I spell out in detail the behavior of the simplest tense operators of tense logic. Once the symmetry assumption is spotted in this context, it is easy to also identify it in a variety of alternative, more complex frameworks.

1.2 Symmetric Semantics

Consider a toy language capable of expressing temporally structured information. The core feature of the sort of language I have in mind is that basic tensed sentences are the product of composing a temporal operator

[2] See, e.g., Goldblatt (2006), §7.3 on the logic of concurrency.

with a tenseless sentence-like object. Thus the sentence *she passed* might be the result of composing a tense operator (something with the meaning of *in the past*) and a tenseless core (*she pass*). (Note that in giving examples of such tenseless cores, I will omit gender and number features.) It is a substantive empirical assumption, though one that seems initially plausible for many languages, that natural language sentences have this kind of linguistic structure. At any rate, I will accept this assumption for the sake of illustration.

The basic building blocks of the toy language are *sentence radicals*. Radicals are tenseless descriptions of events (*Bea run*) or states (*Al be happy*). I will not explore the inner structure of tense radicals, assuming instead that they are directly interpretable for truth and that if they are not in the scope of a tense operator, they default to a present interpretation. So *Al be happy* describes a state of happiness by Al that happens at the time at which it is uttered. (As the semantic development approaches something closer to English, in Chapter 7, I will require that radicals always combine with tense before being even interpretable for truth or falsity.) In addition to radicals, our toy language contains Boolean connectives (*and, or, not*, etc.) and the temporal operators *was* and *will*.

In this language, one can say things with obvious English analogues such as

will(Bea run) and not was (al be happy)

The temporal operators of the toy language have meanings that roughly translate as *at some point in the past* and *at some point in the future*.

We can already read off these informal glosses the expectation that the meanings of *was* and *will* should be symmetric. This expectation is borne out upon development of standard techniques of model-theoretic semantics.[3]

The first thing we need toward that development are, of course, models. Think of models as abstract objects that depict the features of reality that are needed to capture the semantic properties of the expressions of the language. To use Etchemendy's (1990, chapter 2) terminology, this means that models are understood *representationally*. That is to say, models are simplified representations of temporally structured worlds. Under this conception, two different models represent two different ways a temporally

[3] A word about notation: When ignoring what happens at a sub-sentential level, I generalize over sentences and sentence radicals of this language by means of variables such as A, B, C, etc. In these cases, I will only be concerned with connectives and temporal operators as means of composition of sentence radicals.

structured world could be – fixing the meanings of our words.[4] Needless to say, much as our toy language does not aim to represent all the complexities of a natural language, these models incorporate substantial idealizations. They do not aim to represent all the complexity of our world and they do not aim to be completely accurate representations of those features that they do represent. With these clarifications in mind, define:

Definition 1.1 A *model* $\mathcal{M}$ is a triple $\langle \mathcal{T}, <, v \rangle$ with $\mathcal{T}$ a set of times; $<$ a linear order on $\mathcal{T}$; and v a valuation function.

A linear order on a set S is a transitive, antisymmetric, and total relation on S. The valuation function v maps each sentence radical of the target language $\mathcal{L}$ to a truth-value relative to a time. So if Bea does run at t_0, we have $v(\textit{Bea run}, t_0) = 1$. The possible truth-values in our interpretation schema are 1, for true, and 0, for false.

Suppose that our toy language contains exactly four radicals

Bea run, Al run, Bea be happy, Al be happy

Figure 1.1 diagrams the structure of a simple model and possible assignments to sentence radicals by the valuation function. In these simplified settings, we can diagram the valuation function at each point by a sequence of four 0's and 1's, where a 1 in the nth place of the sequence means that the nth radical (in the preceding list) is true at the given time, and a 0 means that it is false. In the model diagrammed by Figure 1.1, *Bea* and *Al* are both running and happy at t_2; they both stop running at t_3, but only *Al* is happy then, and he stays happy for one more tick of the clock. After that, the world ends.

Models have two jobs: to contribute to a theory of the truth-conditions of sentences of the language and to contribute to an account of logical entailment. However, models (as just characterized) are not enough to state the truth-conditions for the sentences that are of interest. The question whether *Bea be happy* is true in the model of Figure 1.1 is ill posed: Bea is happy at some but not all times. To determine whether the radical *Bea*

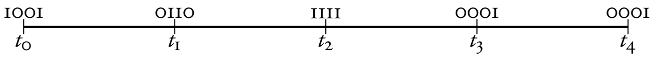

Figure 1.1 Valuation function for a model with four radicals

4 Representational models contrast with "interpretational models," which represent different ways in which we might interpret the nonlogical fragment of the language fixing what the world is like.

be happy is true, one must take the perspective of a particular point in the temporal development of the model.[5] Speaking formally, the compositional semantics for this language centers not on a definition of truth in a model $\mathcal{M}$ but on a definition of truth at a *point of evaluation*. In the semantics for the toy language, a point of evaluation consists of a model together with a time that is drawn from the stock of times of that model.

Given a model $\mathcal{M}$ and a time t that belongs to the model's stock of times, truth-values (again, 1 for true and 0 for false) can be assigned to all sentences of the language. The double brackets $[\![\cdot]\!]^{\cdot}$ denote the interpretation function. This function inputs an expression and a point of evaluation and outputs the semantic value at a point of evaluation. Since points of evaluation are pairs consisting of a model and a time, I will write $[\![\cdot]\!]^{\mathcal{M},t}$. As our points of evaluation change, so will the parameters in the double-bracket notation. Unless it is important to remind the reader that everything is relativized to a model, I omit the $\mathcal{M}$ superscript.

The assignment starts with sentence radicals (this language's equivalent of "atomic sentences") that are interpreted directly by the valuation function. Suppose we want to evaluate the semantic value of *Bea run* at t_1 in model $\mathcal{M}$. First, extract the model's valuation function $v_{\mathcal{M}}$, and then check what the valuation function assigns to *Bea run* at t_1. In the model of Figure 1.1, this is 0, so we write $[\![\textit{Bea run}]\!]^{t_1} = 0$.

The pattern of propagation is determined by the *lexical entries* of the various expressions that are used to generate complex sentences. For sentential connectives, adopt standard Boolean entries.

$$[\![\mathsf{A}\ \textit{and}\ \mathsf{B}]\!]^t = \textit{min}([\![\mathsf{A}]\!]^t, [\![\mathsf{B}]\!]^t)$$
$$[\![\mathsf{A}\ \textit{or}\ \mathsf{B}]\!]^t = \textit{max}([\![\mathsf{A}]\!]^t, [\![\mathsf{B}]\!]^t)$$
$$[\![\textit{not}\ \mathsf{A}]\!]^t = 1 - [\![\mathsf{A}]\!]^t$$

A conjunction is true just in case both its conjuncts are true; a disjunction is true if even one disjunct is; and negations flip truth-values around.

Now for the queen of the entries. The goal is to assign semantic values to *was* and *will* that capture the symmetric approach to truth-conditions of past and future claims. To this end, say that *was*(*Bea run*) is true at t if there is a prior time at which Bea runs; on the future side, say *will*(*Bea run*) is true at t if there is a future time at which Bea runs. Figure 1.2 illustrates the idea.

[5] This parallels one of the standard moves to interpret models for modal logic: In that context, the points at which we evaluate atomic sentences are possible worlds. If the worlds in a model disagree on the truth-value of some sentence A, and for whatever reason it matters to fix that truth-value, we can designate a world in the model as actual.

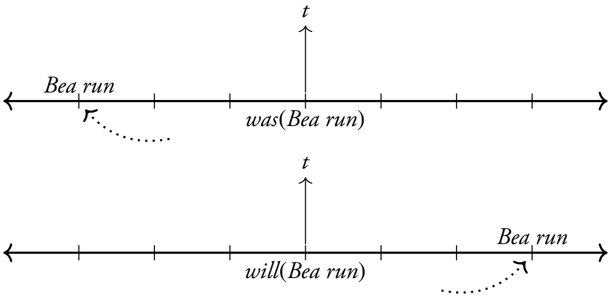

Figure 1.2 Temporal shift for symmetric semantics

More formally, and more generally, we can state the following pair of lexical entries:

LINEAR SYMMETRIC SEMANTICS

a. $[\![\mathit{was}(\mathsf{A})]\!]^t = 1$ iff $\exists u < t, [\![\mathsf{A}]\!]^u = 1$
b. $[\![\mathit{will}(\mathsf{A})]\!]^t = 1$ iff $\exists u > t, [\![\mathsf{A}]\!]^u = 1$

These entries complete the definition of truth at a model/time pair for our toy language.

From the model-theoretic perspective, a near-classical way to define entailment is as preservation of truth at a point of evaluation (a pair $\langle \mathcal{M}, t\rangle$). An argument with premises $\mathsf{A}_1, \ldots, \mathsf{A}_n$ and conclusion C is valid (written $\mathsf{A}_1, \ldots, \mathsf{A}_n \models \mathsf{C}$) just in case there is no model $\mathcal{M}$ and time $t \in \mathcal{T}_{\mathcal{M}}$ such that all the A's are true at $\langle \mathcal{M}, t\rangle$ but C is not. It is invalid otherwise (in which case, we write $\mathsf{A}_1, \ldots, \mathsf{A}_n \not\models \mathsf{C}$).

The system of logic resulting from this interpretation of temporal language has been studied in A. N. Prior's seminal work (Prior, 1957, 1967, 1969) and, later, in the context of the explosion of model theoretic techniques for modal logic (Rescher and Urquhart, 1971; Burgess, 1979). The appendix to this chapter reviews a sound and complete axiomatization of (a slight refinement on) this consequence relation.

At a less formal level, it is important to understand the proper interpretation of verdicts about validity and invalidity in this system. Here are four sample patterns, a valid one and three invalid ones.

(1) a. $\mathit{will}(\mathit{will}(\mathsf{A})) \models \mathit{will}(\mathsf{A})$
b. $\mathsf{A} \not\models \mathit{will}(\mathsf{A})$
c. $\mathsf{A} \not\models \mathit{will}(\mathit{was}(\mathsf{A}))$
d. $\mathsf{A} \not\models \mathit{was}(\mathit{will}(\mathsf{A}))$

It is tempting to assess these verdicts by considering parallel inferences in English and evaluating whether they sound valid. It is remarkable, however, that the inferences in (1) do not have immediate, natural-sounding translations. Consider (1-d). The tempting thought is to view this as a theoretical prediction about the invalidity of an inference that in English we might put as follows:

I am happy

It was the case at some point in the past that I was going to be happy

The problem is that the English inference features additional material that is not exactly semantically inert: *was going to* is not precisely the same as *will*. Another problem is that, as I will show shortly, even holding fixed the semantics, we can change the validities by imposing constraints on the class of models. For this reason, we need to be careful in making judgments about the acceptability of this inference as evidence for or against the semantics.

It is possible to justify the verdicts in (1) by appealing to intuitions of a different sort. We might have intuitions about whether these inferences should come out valid *given the informal glosses for the tense operators*. To illustrate with (1-a), we might reason that if at some future point, there is a further future point at which A is true, then it has to be true that at the origin point there is a future point at which A is true. As for (1-b), it is clearly invalid under the informal glosses we have been operating with: The present truth of A at some point in time is not enough to establish the truth of A at a later point. While these judgments cannot be used to address the *empirical adequacy* of the symmetric semantics, they can help us investigate whether the formalism is a correct implementation of its informal design specifications.

In addition, the semantics can illuminate correspondences between logical validities and the temporal reality they represent. Consider again (1-c) and (1-d). As we think through the meanings of the symmetric tense operators, we might intuit (1-c) and (1-d) as valid. Surprisingly, the current semantics invalidates them both. Consider evaluating *Al be happy* at the last time t^* of a finite model. Because this is the last time in the model, every sentence of the form *will* A is false at t^*. After all, if there is no time after t^*, there is no time after t^* at which A is true. As a result, both *will(Al be happy)* and *will not(Al be happy)* are false. To enforce the validity of (1-c) and (1-d), we could constrain time to have no endpoints toward the future, which validates (1-c), or toward the past, which validates (1-d). The familiar point here is that, given the standard tense logic entries, there is a correspondence between the metaphysical assumption that timelines are infinite and the validity of (1-c) and (1-d).

Alternatively, we could adopt a more pragmatic perspective: In any context in which speakers can convince themselves that they are not evaluating from the temporal edge of the world, these schemas will preserve truth. So in practice, we might be warranted in treating them as valid inferences even if they are not valid in full generality. This perspective might be bolstered by noting that the following are valid entailments in this system.

(2) a. $\mathsf{A}, \mathit{will}(\top) \models \mathit{will}(\mathit{was}(\mathsf{A}))$
b. $\mathsf{A}, \mathit{was}(\top) \models \mathit{was}(\mathit{will}(\mathsf{A}))$

(Here $\top$ stands in for an arbitrary tautology.) We might take this to mean that the inferences (1-c) and (1-d) should be acceptable (again, given an informal gloss on the temporal operators) whenever it is presupposed that there is a future (by implicitly accepting *will*($\top$)) or that there is a past (by implicitly accepting *was*($\top$)).

1.3 The Symmetric Paradigm Contextualized

The semantic analysis I provided in Section 1.2 exemplifies but does not exhaust the symmetric paradigm. Much pioneering work in tense logic was not explicitly addressed to the semantic analysis of natural language. And it is misleading to take Prior's work as directly providing a semantic analysis of English tenses (Ogihara, 2007, pp. 393–397).[6] Moreover, the empirical hypothesis that natural language tenses might work like tense logic operators quickly found a rival approach – namely, the view that instead they work as object-language referential devices (Partee, 1973; Dowty, 1982; see King, 2003, chapter 6 for an extensive overview).

The essence of the symmetric paradigm is not a specific semantic assumption, however. Instead, it is the general idea that the semantics of

[6] Ogihara discusses the tricky case of theorists such as Montague. Unlike Prior, Montague is clearly concerned with providing a model of the semantic functioning of the natural language tense system. However, Montague's methodology involves translating from English to an intermediate language like the toy language of Section 1.2. It seems possible to argue that it is not a problem if the intermediate translation language is *too* expressive. A more significant problem is that, even when natural language tenses do stack, the truth-conditions are generally not what is predicted by the tense operator account.

future and past tenses ought to be mirror images of each other. Even as the semantics of tense has become more nuanced and expressive, the symmetric paradigm has maintained some of its shine and theoretical grip. In this section, I consider some simple ways of refining our understanding of the semantics of tense while sticking to the central tenets of the symmetric paradigm.

One way in which the operators of tense logic are at best coarse approximations of natural language tenses is that not every future moment can matter to our evaluation of future claims and not every past moment can matter to our evaluation of past claims. Ordinarily, as speakers, we restrict attention to certain specific points in time. Sometimes we do that by explicitly restricting the relevant temporal range. If I say *I cooked dinner*, I arguably do not speak truly if I cooked dinner once in 1995. In practice, I must be talking about some restricted interval of time that is easily identifiable by my conversational partners. This suggests that temporal talk has a context-sensitive dimension that is absent from my initial formulation of the symmetric clauses. What is more, these salient intervals are not controlled only by the linguistic context. They can be explicitly restricted or modified.

(3) a. Last May, Isabella visited Canada
    b. In 1984, Los Angeles hosted the Olympics

We cannot correctly analyze these phrases if English past tense just meant *at some point in the past* and if it did not compositionally interact with these restricting expressions.

There is a natural way to enrich the symmetric semantics to provide it with this kind of flexibility. Suppose that instead of a time, the points of evaluation also keep track of a time interval.[7]

Definition 1.2 An *interval* $\mathcal{I}$ is a set of times satisfying the property that for any two times t and v both in $\mathcal{I}$, if $t \leq u \leq v$, then $u \in \mathcal{I}$.

More succinctly, intervals are *convex* sets of times (with respect to the temporal precedence ordering).[8]

This modification allows a distinction (as in Figure 1.3) between those cases in which A occurs within the designated range I and those in which

[7] The idea of using intervals in the semantics comes from Bennett and Partee (1972), and has been widely applied in many subsequent frameworks (Dowty, 1982; Condoravdi, 2002).

[8] Proponents of intervals in temporal semantics commonly stipulate that the set of times has the cardinality of the continuum (this is the approach of Bennett and Partee, 1972, p. 69). Definition 1.2 is neutral on matters of cardinality.