Journal Title: Compositional semantics an introduction to the syntax/semantics interface

Volume: c.1
Issue:
Month/Year: 2014
Pages: 27-41

Article Author:

Article Title: Semantic Foundations

Cited In: ScanDelver

Notes:

ODYSSEY REQUEST

Christopher D Kennedy
ck0@uchicago.edu

Sent HP 6/13

ODYSSEY REQUEST

the set of states (every state does have a capital) to the set of US cities. But it is not a one-to-one correspondence for the same reason that our original relation is not a total function; there are many cities without the honor of being a capital.

Occasionally in this text it will be useful to list out some actual functions—that is, to name every member in the domain and name what the function at issue maps that member to. There are a variety of ways one could do this. To illustrate, take a domain of four children {Zacky, Yonnie, Shelley, and Baba} (call that set C) and four men {Abe, Bert, Carl, David} (call that set M). Suppose there is a function f from C to M which maps each child to their father. Assume that Abe is the father of Zacky and Yonnie, Bert is the father of Shelley, and David is the father of Baba. Then one can write this information out in various ways. One would be to simply give the set of ordered pairs: {(Zacky, Abe), (Yonnie, Abe), (Shelley, Bert), (Baba, David)}. Usually this notation, however, is not terribly easy to read. We could also write this out in either of the ways shown in (20):

(20) a. f(Zacky) = Abe
f(Yonnie) = Abe
f(Shelley) = Bert
f(Baba) = David

b. Zacky → Abe
Yonnie → Abe
Shelley → Bert
Baba → David

Or, sometimes it is more convenient to list out the domain on the left and the co-domain on the right and connect them with arrows as in (21):

(21)

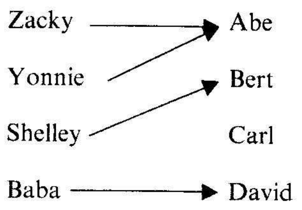

Which notation is chosen makes no difference; the choice should be dictated by clarity.

2

Semantic foundations

2.1. Model-theoretic semantics

The primary focus of this book is the syntax/semantics interface—that is, how the syntax and semantics work so that sentences (and other well-formed linguistic expressions) are paired with a meaning. But of course this task is impossible without some idea of what meaning is. Can we talk about meaning without relegating it to the realm of the mysterious, or leaving it solely to folks who work on cognition to deal with? The answer (of course) is yes—there is a rich tradition within linguistics and the philosophy of language for modeling linguistic meaning.

In some early work within linguistic theory—especially in the 1960s—meaning was taken to be just a symbolic representation (call it a Logical Form, or LF). While it remains an open question as to whether such representations play an actual role in the way the grammar pairs expressions with meanings, this book (along with much other modern work in formal semantics) assumes that meaning is not just some string of symbols, but rather some actual object out there in the world. Call this a *model-theoretic object*. (More precisely, we are taking meaning to be an object which forms part of a model which is an abstract representation of the world: hence the term *model theory*.) Of course, we need some way to name these objects, and so throughout we will use strings of symbols as ways to name them. But the point is that the grammar maps each linguistic expression into something beyond just a symbolic representation. Otherwise—as so aptly pointed out by David Lewis (1970)—we are simply mapping one language (say, English) into another (what Lewis termed "Markerese"). Yet language is used to convey facts about the world; we draw inferences about the world from what we hear and we gain information about what is true and what is not. So semantics must be a system mapping a linguistic expression to something in the world.

But what exactly is meant by model-theoretic objects? These can in fact be quite abstract. Still, they are the "stuff" that is out there in the universe—something constructed out of actual bits of the universe (or, at least, the ontology of the universe as given by language). This would include things like individuals, times, possibilities, and perhaps others; just what are the basic objects that we need is an open question and is part of what semantic theory addresses. The strategy here will be to use a fairly sparse set of primitive objects, and construct more complex objects out of these. Let us, then, begin by setting up two basic building blocks which are foundational in much of the work in linguistic formal semantics.

2.2. Truth conditions

A common adage in semantics is: "To know the meaning of a (declarative) sentence is to know what it would take to make it true." We can use this adage as a first step in constructing the building blocks for meanings: a fundamental fact about declarative sentences is that they are either true or false[1] (and since we use language to communicate information about the world, a listener will in general assume that a sentence they have just heard

[1] Henceforth we use the term "sentence" to mean a declarative sentence. There is actually no reason to consider questions to be of the same category as declarative sentences even though they also are traditionally referred to as "sentences." Questions have a different external distribution from declarative sentences (for example, *wonder* can occur only with a question as its complement, not an ordinary sentence, while the reverse is true for *believe*); they have a different kind of meaning, and they have a different internal structure. Whether imperatives and declaratives should be considered the same category is a bit less clear (they are more similar), but we will not deal with those here either.

is true, and uses that fact to enrich their knowledge of the world). Thus (1) is true and (2) is false:

(1) Barack Obama moved into the White House on Jan. 20, 2009.

(2) John McCain moved into the White House on Jan. 20, 2009.

Hence, one basic notion used for the construction of meanings is a truth value—for now assume that there are just two such values: true and false. (More on this directly.) The claim that truth values are a fundamental part of meaning is also motivated by noting that—as shown by the examples above—speakers have intuitions about truth, given certain facts about the world, just like they do about acceptability. And these judgments can be used to test the adequacy of particular theories of meaning. Following standard practice, we use 1 for true and 0 for false. Thus the set of truth values {1,0} and we will also refer to this set as t. Let us use [[α]] to mean the semantic value (i.e., the meaning) of a linguistic expression α. Then (temporarily) we can say that [[*Barack Obama moved into the White House on Jan. 20, 2009*]] = 1.

Some worries should immediately spring to mind. The most obvious is that something seems amiss in calling the *meaning* of (1) "true" even if we are willing to accept the fact that it is true. We will enrich the toolbox directly to take care of that. But there are other objections: does it really make sense to say that all declarative sentences are true or false? Clearly not—for some sentences the truth value depends on who is speaking (and on when the sentence is spoken). Take (3):

(3) I am President of the United States.

This is true if spoken by Barack Obama in 2011, but not if spoken by John McCain and not true if spoken by Barack Obama in 2006. So this has no truth value in and of itself. Nonetheless once certain parameters are fixed (time of utterance and speaker) it is either true or false. So we might want to think of the meaning of (3) as a *function* into {1,0}—it does yield a truth value but only once we fix certain parameters. But it seems inescapable that a declarative sentence is telling us something about the world, and so truth values are certainly one fundamental piece.

In fact, there are many parameters that need to be set in order to derive a truth value. Certain words like *I*, *you*, *here*, *now*, etc. quite obviously have the property that their value depends on when, where, and by whom these are spoken (these are called *indexicals*). There are also more subtle cases—such as (4) and (5)—where truth seems to depend on what is at issue in the discourse context in which these are uttered:

(4) Magic Johnson is tall.

(5) Every student got an A on their formal semantics midterm.

(4) might be true if we are comparing Magic Johnson to the general population (Magic Johnson is a former basketball player for the Los Angeles Lakers) but perhaps not true if we are comparing him to the set of basketball players. The context of utterance usually makes it clear what is the relevant comparison class. (5) may be true if we are restricting the interpretation of *students* to those students in my formal semantics class in 2011, but not if we mean every student in the world, or even every student at Brown or every student who has ever taken formal semantics from me. We return to cases like these in (among other places) sections 8.4, 10.6, and 17.4.

2.1. In what ways do the following sentences also show a dependency on context of utterance? (Some may show this in more than one way.)

(i) The dog did not get a walk today.
(ii) Samantha hasn't eaten yet.
(iii) All of the ex-presidents of the US gathered at the State of the Union address.

So far, we have two reasons that it is oversimplified to say that the meaning of a sentence is true or false: (a) even once we do determine the truth value of a sentence we surely don't want to call that "meaning," and (b) often the truth value can't be determined until we know the context of utterance. Two further worries have to do with the fact that (c) there are vague sentences where some have the intuition that the truth value is something in between 1 and 0, and (d) some sentences (even once we fix the context of utterance) seem to be neither true nor false (nor anything in between).

As to (d), these are cases which have generally gone under the rubric of *presupposition failure*. The classic example from Russell (1905) is (6); another example would be (7) said of someone who never went to an aerobics class:

(6) The present King of France is bald.

(7) He stopped going to aerobics class.

Surely (6) is not true, but do we want to say it is false?[2] For the most part, we will make the expository simplification of assuming all sentences are true or false, although we will return to the issue of presupposition from time to time.

This still leaves the issue of vagueness. For example, even once we do fix the comparison class at issue in a sentence like (4) there remains some vagueness: is a 6 foot 9 inch basketball player tall for a basketball player? If so, exactly where does one draw the line between tall and not tall? There is a rich literature on the problem of modeling vagueness and some have attempted to model it using intermediate truth values (where a sentence can have any value between 1 and 0). Like many other domains, we can set this aside in this introductory text (but for a nice discussion aimed at introductory students, see Chierchia and McConnell-Ginet 1990: chapter 8). For the present purposes, take all declarative sentences (once fixed to a context) to have just one of two truth values: 1 or 0.

2.3. Possible worlds

2.3.1. *Introducing the notion*

The one important observation that we will not set aside is the first problem noted above: it hardly makes sense to think of the *meaning* of a sentence like (1) as true. There are plenty of reasons to find this absurd. First, we can know the meaning of a sentence without having the slightest idea of whether it is true or false: I know the meaning of (8) but have no idea of its truth value:

(8) The tallest man alive anywhere on January 1, 2010 had pomegranate juice for breakfast.

[2] Russell (1905) answered this question "yes"; under his analysis (6) is false. Strawson (1950) argued "no" and introduced the notion of presupposition by which a sentence like (6) is neither true nor false.

And surely we do not want all true sentences to have the same meaning (ditto for all false sentences). In fact if this were all there is to meaning we would be able to express only two things: "1" or "0." Obviously a rather strange notion of what we do with language! But let us return to the adage: "To know the meaning of a sentence is to know what it would take to make the sentence true." Thus enters the notion of a set of *possibilities*, otherwise known as a set of *possible worlds*.[3] This gives us just the tool we need: to know the meaning of a sentence is to know what the world would have to be like to make it true: i.e. the meaning of a declarative sentence, then, is a function from ways the world could be (possible worlds) to {1,0}. (Readers not familiar with the notion of a function should consult the appendix to Chapter 1.)

This makes good sense. We can know the meaning of a sentence without knowing its truth value not because of our lack of linguistic knowledge but precisely because we don't know what is the actual world in which we live. But if we did—if we were omniscient—then for every sentence we would of course know whether it is true or false. Moreover, think of what it means for one's information to grow as we communicate. As we hear a sentence (and if we also assume that the speaker is speaking truthfully) we narrow down the space of possible worlds that we could be in; all worlds mapped to 0 by the relevant sentence are eliminated. This also provides a nice way to capture the notion of an *entailment*. (9), for example, entails (10) (Mitka is the name of the author's former dog):

(9) Mitka killed the bird that had been trapped on the porch.

(10) The bird that had been trapped on the porch died.

Hence, we can say that a sentence S_1 *entails* S_2 if and only if every world mapped to true by S_1 is mapped to true by S_2. Put differently (but equivalently), the set of worlds in which (9) is true is a subset of worlds in which

[3] The terminology of "possible worlds" often causes people to cringe—it brings with it some unfortunate terminological baggage. While there may be deep philosophical questions that ultimately need to be addressed here, one can go quite far without worrying too much about the *real* status of "alternate possible worlds." Surely we do talk about other ways the world might be (see the discussion below). And surely there are many facts about the actual world that we do not know, hence think of a possible world as a way the world might be compatible with our knowledge state.

(10) is true. (The reverse of course does not hold; (10) can be true because a different dog killed the unfortunate bird, or because it died of a broken heart.) Similarly, we can call two sentences *synonymous* if and only if they entail each other—i.e., they have exactly the same truth conditions.[4]

As soon as one begins to talk about possible worlds and truth conditions, an interesting question arises. Consider a sentence that is necessarily true—that is, one which seems to be true no matter how the world happens to be and hence true in all possible worlds. If the meaning of a sentence is a function from worlds to values, then all necessarily true sentences have the same meaning. But do we want to say that all necessarily true sentences really have the same meaning? (Similarly for all sentences that are necessarily false.) These are thorny issues. But some of the discomfort posed by modeling the meaning of a sentence as a function from possible worlds to the set t can perhaps be allayed by taking a rather charitable view of what it means to be a possible world. Consider the (often cited) case of mathematical facts. Such facts are what they are—they are not really contingent facts about the world (the world could not change in such a way as to make them false). Yet there is a strong intuition that (11a) and (11b) do not have the same meaning:

(11) a. The largest prime number less than 1000 is 997.
b. The largest prime number less than 2000 is 1999.

[4] Many find this an unsatisfying definition of "synonymy" since two sentences can have the same truth conditions but a variety of subtle differences such as differences in when they are appropriately used, what they draw attention to, etc. Consider active–passive pairs:

(i) a. Mitka killed the bird that had been trapped on the porch.
b. The bird that had been trapped on the porch was killed by Mitka.

One can imagine a variety of situations in which it would be appropriate to utter (ia). For example, if I come home at the end of the day and ask my husband what Mitka did during the day, (ia) would be a reasonable answer. It is also a reasonable answer to the question of why the porch appears to be full of feathers. But while (ib) is reasonable in this second scenario, it is an odd answer to the question of what Mitka did during the day. These kinds of differences often go under the rubric of *information structure*. Ultimately, then, one might well want a much finer-grained notion of synonymy that includes not only the truth conditions of a sentence but other factors as well. But this is simply a matter of terminology: here we elect to use the term synonymy in this rather coarse-grained way (referring to having the same truth conditions). One could substitute in "mutual entailment" if one prefers.

Arguably, though, this is because the set of possible worlds includes more than just the mathematically possible worlds; rather, it includes all worlds which we can conceive of. Until writing out (11a) and consulting a number theorist I had no idea whether it was true or not: surely the space of epistemically possible worlds (worlds compatible with my knowledge state) includes worlds in which (11a) is true and worlds in which it is false. Similarly for (11b). Indeed, the truths of mathematical statements are subject to discovery (and ditto for truths that hold via physical laws of the universe) and so there's no real reason to restrict the notion of possible worlds to those in which those laws hold. (If one prefers, we should perhaps rename these as *imaginable worlds*—but we stick here to the term *possible worlds* out of convention.) Thus at least some of these concerns can be allayed in this way.

This still leaves the case of sentences known as *tautologies* (true by logic and/or by definition) such as the pair in (12). No amount of imagination would make these false (someone who thinks they are false presumably just does not know the meaning of *or*) and so the theory says they have the same meaning:

(12) a. Henry had breakfast yesterday or Henry didn't have breakfast yesterday.
b. Sally climbed Mt Washington or Sally didn't climb Mt Washington.

We will simply agree to live with this problem (if it is really is one). Perhaps literally these two do have the same meaning. As will be discussed at points throughout this text, there is a distinction between the literal meaning of sentences and other things that they might be used to convey (the latter being the domain of *pragmatics*). It is not unreasonable to think that the pair in (12) have the same meaning but could be used for different purposes. A full-blown theory of pragmatics should be able to predict that while both sentences are always true, they can indirectly convey different information and hence be used in different circumstances. So the notion that the meaning of a sentence is a function from possible worlds (broadly construed) to truth values does not seem fundamentally threatened by these cases.

Incidentally given the meaning of a sentence as a function from worlds to the set {1,0}, the notion of a *partial function* can be used to model presupposition failure. Hence a common account of the fact that *The present King of France is bald* seems somehow just nonsensical in our world is to suppose

that it assigns neither true nor false to this world. The meaning of this sentence is thus a function from worlds (and times) to truth values, but a partial one—undefined for this world. Whether this is the right way to think about presupposition is a topic that we leave open in this work, although we will be confronting the issue(s) from time to time.

2.3.2. Characteristic function of a set

Take any function f from some domain D into the set t (i.e., the set {1,0}). Then consider that subset of D consisting of all and only the members of D that are mapped to 1 by f; call it D'. Then we can say that D' is the set *characterized by* f. (Note that D' can be Ø.) Similarly, if we have a set D' which is a subset of D, then we can construct the function f that maps every member of D' to 1 and every member of D not in D' to 0. f is called the *characteristic function* of the set D'. Note that for any function f into {1,0} we can recover the set D'. Note moreover that for any D' a subset of D and any D, we can uniquely recover the total function f that maps all members of D' to 1 and everything else to 0. In this way we can go back and forth between talking about certain functions (those whose domain is {1,0}) and the sets that they characterize.

This is useful in thinking about the meaning of sentences. We have taken the meaning of a sentence to be a function from worlds to {1,0} but one can equivalently take it to be a *set of possible worlds* (namely, all worlds mapped to 1 by the relevant function). It is useful to be able to go back and forth between these two conceptions. Note that the function is uniquely recoverable from the set only under the assumption that every sentence is indeed assigned 1 or 0 in every world. Thus if we treat presupposition failure as best modeled by allowing partial functions—a sentence which is true in some worlds, false in others, and assigned no value in yet other worlds—then there is no way to uniquely recover the relevant function from the set of worlds mapped to true. (Those not in the set might have been mapped to 0 or might have been given no value.) However, as long as the function is total (thus all worlds are assigned 1 or 0 by the relevant function) it can be uniquely recovered. The reverse holds in either case: one can always construct the set characterized by a function f whose co-domain is {1,0}: that set is the set assigned 1.

2.2. Consider the function f whose domain is the set of New England states and co-domain is {1,0} which is as follows: f(Maine) = 0, f(New Hampshire) = 1, f(Vermont) = 1, f(Massachusetts) = 1, f(Rhode Island) = 0, f(Connecticut) =1. Show the set it characterizes.

2.3. Now consider the case of partial functions from the set of New England states to {1,0}. That is, we are considering functions which map each state into either 1, 0, or where the function can be undefined for some states (some states will be mapped into neither value). Consider the set {Maine, Vermont, Rhode Island}. List out all of the possible functions (partial or total) which characterize this set. (For readers unfamiliar with American geography, the remaining New England states are as given above—they are New Hampshire, Massachusetts, and Connecticut.) See Appendix to Chapter 1 for ways to write out functions.

2.3.3. *Notation and terminology*

Recall that t is the set {1,0} (i.e., the set of truth values). We will use w to mean the set of worlds. We will use the notation <a,b> to mean the set of all functions whose domain is the set a and co-domain is the set b.[5] Thus <w,t> is the set of functions from worlds to truth values. We will speak of the *semantic type* of a syntactic category such as a "sentence" (S) as being <w,t>; this means that every expression of this category has as its meaning some member of <w,t>. (Thus <a,b> itself means the set of all functions from a to b; the terminology "of type <a,b>" means that an expression has as its meaning a member of <a,b>.)

Take any expression whose meaning is a function from possible worlds to something else. (So far we have seen this just for the case of sentences.) We call that function the *intension* of the expression. Thus the intension of an expression is its actual meaning, not its value in some particular world.

[5] This is very slightly different from the way this notation is used in some other works in semantics (in particular this differs slightly from its use in Montague 1973 which is how it was introduced into linguistic semantics). However, the differences are inconsequential for the purposes here.

When talking about the value of some expression α at some particular world, we will use the term *extension of* α *at a world w*. Most often we are concerned with its value in the actual world, and in such cases we will simply refer to the *extension* of the expression.

2.3.4. *Talking about worlds*

Possible worlds have been added to the basic toolbox as a way to capture the fact that the meaning of a sentence is not simply 1 or 0, but rather can be thought of as a function from worlds to truth values. But there is a second reason for incorporating possible worlds into the theory, which is that natural language contains all sorts of expressions and constructions which actually make reference to worlds other than the actual one. Put differently, there are words whose meanings crucially combine with the *intensions* of sentences (or other expressions).

A rather striking example of how language allows us to talk about possible worlds is the case of *counterfactual conditionals*. These are sentences of the form *If* S_1 *then* S_2, where there are special forms of the verbs (or auxiliaries) in these sentences (in particular, the *then*- clause often contains the auxiliary *would*). Consider (13):

(13) If the US Supreme Court had not stopped the Florida recount in 2000, nuclear power plants would no longer be in use in the US.

Some relevant background: the outcome of the 2000 US Presidential election (between Al Gore and George W. Bush) hinged on a few hundred votes in Florida. Although Bush was initially declared the winner of Florida (and hence of the US election), Gore asked for a recount of the votes in key precincts. There was a complex series of court decisions, leading up to the Florida Supreme Court ruling in favor of Gore's request. Bush then took this to the US Supreme Court. While the recount was in process (and hence no one knows for sure what the outcome would have been), the US Supreme Court halted the recount, and subsequently ruled that it should not proceed. Bush therefore remained the victor in Florida and hence won the US presidency. An additional bit of relevant background is that Al Gore is very opposed to the use of nuclear power.

Given this, it is easy to understand (13), and to understand the reasoning leading a speaker to say this. What is interesting about such sentences is that we can argue about whether they are true or false but we can never really know, as we are arguing about a chain of events in another possible world. A reasonable approximation of the truth *conditions* of a counterfactual conditional sentence of the form *If* S_1 *then* S_2 is that it is true if in the closest world to ours where S_1 is true S_2 is also true. "Closest" means the world just like ours where we make minimal changes. In other words, we look at a world in which the *if*-clause (known as the *antecedent*) is true and where nothing else changes except those facts that follow from the truth of the antecedent. (Usually these are uttered in a situation where the antecedent clause is not true in the real world and hence the closest world is not our world, but this is not necessary; the closest world in question could be the actual one.[6]) Of course in those cases where the relevant other world is not ours, we are not omniscient; we can't just jump over to another possible world. And this is exactly why we argue about the truth of counterfactual conditionals: we don't really know what the facts are in our relevant alternate reality. In fact, (13) is interesting because we don't even know for sure whether Gore would have won the election had the Supreme Court not stopped the recount. People argue about this (different newspapers tried to do counts, and came up with different results), so we actually don't even know the truth value of a simpler sentence like *If the Supreme Court had not stopped the Florida recount, Gore would have become President*. But we do know what these sentences all mean: we know what it would take to make them true.

Notice, incidentally, that we have restricted the semantics to looking only at the *closest* world in which the antecedent is true (the world most like ours where we change as little as possible). It would be incorrect to posit, for example, that the consequent must be true in *all* worlds in which the

[6] An illustration of the fact that the construction does not require the antecedent to be false in the actual world is the following "detective" scenario (modified from discussion in Anderson 1951): A great detective is analyzing the scene of a murder in which there are bloody footprints of a size 12 shoe, an empty vial of an extremely rare poison available only in Saskatchewan, and a note clearly scribbled with a magenta felt-tip marker. Our detective reports his thinking as follows. "Let's see now. If Mr Magoo had committed this murder he would have left bloody footprints from a size 12 shoe. If he had committed the murder he would have purchased his poison in Saskatchewan. And if Mr Magoo had committed the murder he would have left a note with his favorite magenta felt-tip marker. I therefore conclude that Mr Magoo is indeed our murderer."

antecedent is true. For surely there is a possible world in which the recount proceeded, Gore did get more votes, was assassinated his second day in office, and where his successor Lieberman had no qualms about nuclear power. Or for that matter, one cannot doubt the existence of a *possible* world (quite remote) in which the majority of votes in Florida had actually been cast for Patrick Buchanan or even Mickey Mouse.

The interesting point here, then, is that while we cannot really determine the truth *value* of (13), it does indeed have truth *conditions* (otherwise there would be nothing to argue about), and we can rather precisely state the truth conditions—if not the truth value—using the notion of possible worlds. Our ignorance, as always, comes from the fact that we are not omniscient: we do not have access to the full set of facts in our world nor to the full structure of other possible worlds. Thus in order to give a semantics for *if... then*, the theory needs access to possible worlds.

In fact, one does not even need to consider constructions with such exotic syntax to see that we reason and talk about alternative possibilities in everyday language. Even the semantics of the verb form called the *progressive* (as in *I am writing a book*) arguably requires access to possible worlds. (The discussion below—although not the particular example—is based on the analysis in Dowty 1979.)

Consider the following (silly but wonderfully illustrative) two scenarios.

Scenario 1: Mr Peepers lives alone and spends a lot of time watching the activity in his street out his window. A turtle who he has named Myrtle lives in a nearby pond, but every morning Ms Goodheart who lives across the street from Mr Peepers places food on the curb (across the street from Mr Peepers) for Myrtle. And so every morning at the same time, Myrtle slowly ambles across the street, gets to the curb on the other side, eats her food, and slowly ambles back. One fateful day, Myrtle gets exactly a third of the way across the street, and a large truck comes and hits her and Myrtle dies. Mr Peepers is reporting this, and can quite naturally say:

(14) Poor Myrtle the Turtle was crossing the street when she got hit by a truck and died.

Scenario 2: Exactly the same, except that in this case Ms Goodheart always puts the food right in the middle of the street. Of course this means that every day, Myrtle goes just halfway across the street, has her breakfast, and then returns. But the rest is the same, and so on the fateful day, Myrtle gets exactly a third of the way across the street, and a large truck comes and hits her and Myrtle dies.

Can Mr Peepers report Scenario 2 with (14)? The judgment here is quite robust: no. Yet—in terms of the event that would be described by (14)—the two scenarios contain *exactly identical events!* What is the difference? It is clear that in Scenario 1 we expect that Myrtle would have crossed the street had she not been hit by the truck. (Notice the use of the counterfactual conditional in reporting this.) In other words, the semantics of the progressive is appealing to another possible world in which the truck had not killed Myrtle. (Dowty calls the relevant other possible world the "inertial" continuation of our world: the world that would have materialized had nothing extraordinary intervened.) So even ordinary everyday constructions like the progressive require reference to possible worlds to fully flesh out their truth conditions.[7]

[7] One might dispute the claim above that exactly the same events happened in the two scenarios: in Scenario 1 Myrtle had the intention to go to the other side of the street and in Scenario 2 she did not. So one might argue that the difference does not come from the different possible worlds which are the inertial continuations, but rather is located in the intentions of the subject of a progressive verb.

However, one can repeat the scenario (albeit not as colorfully) with subjects that denote non-sentient beings. *Scenario 1:* A bowler named Jack always throws his bowling ball exactly the same way, and it always lands in the gutter. On the day in question, I watch him play a game—he throws his ball in the usual way, it goes on the usual path. But just as it is five inches from the gutter, a very heavy overhead light falls from the ceiling of the bowling alley and smashes the ball. I can report this as (i):

(i) Jack's ball was rolling into the gutter when an overhead light fell on it and destroyed it.

Now we modify the scenario much as we did with Myrtle. *Scenario 2:* Again, Jack always throws the ball in exactly the same way, and he throws it in exactly the same way as in the first scenario. However, the bowling alley is also home to a small genie who magically appears every time that the ball is three inches from the gutter, corrects its course, and Jack gets a strike every time. (I have resorted to the genie here to make sure that the ball itself is thrown in exactly the same way in the two scenarios. There is no hidden torque on it in Scenario 2; the change in outcome is due neither to the ball nor to Jack, and until the genie appears the events in the two scenarios are exactly identical.) Once again the light falls on the ball when it is five inches from the gutter. Here again (i) is not a good report of Scenario 2.

2.4. Times

It is not enough to think of the meaning of sentences as functions from ways the world might be to truth values. Their truth is also time-dependent. As noted earlier, (3) is true when uttered by Barack Obama in 2010 but not in 2006. Similarly, consider an ordinary past-tense sentence:

(15) George W. Bush was president of the US for 8 years.

This is true now, but when uttered in 1998 it was not true. So the value of a sentence is not only a function from worlds but also times to truth values.

The relation between worlds and times is a tricky one. For the present purposes it will work well to think of the set of times and the set of worlds as completely separate sets. Notice that we can easily talk about "now" in other possible worlds:

(16) If I weren't here now, I'd be outside enjoying the beautiful weather.

It seems that here we are importing the "now" time to another possible world—which is most easily made sense of if we treat the worlds and the times as independent sets.

Thus the full meaning or *intension* of a sentence is a function from worlds and times to truth values, and since the two are independent we can think of it as a function from world–time pairs to truth values. (And its *extension* is the value of this function taken at some particular world and some particular time.) We will use i to represent the set of times. Whenever it is convenient to ignore the worlds but not the times, we will speak as if the value of a sentence is something in the set <i,t>. When it is convenient to ignore the times but not the worlds, we will speak of it as having a meaning of type <w,t>. In reality, though, it is a function from both: and the standard terminology is to use s to name the set of world–time pairs; thus the semantic type of a sentence is <s,t>. As is the case with worlds, there are many linguistic expressions whose meanings crucially combine with time-dependent objects (the past-tense morpheme is one). We explore this briefly in the final chapter. The strategy in this book will be, for the most part, to give an *extensional* semantics (ignoring both worlds and times) as one can get quite far in considering the syntax/semantics interface by looking only at constructions in which intensions don't interact with the syntax. However, the final chapter is concerned with the fact that some expressions are crucially sensitive to intensions, and considers the implications of this to further issues concerning the syntax/semantics interface.