Comp. by: PG4118 Stage : Revises1 ChapterID: 0002108607 Date:16/5/14 Time:06:44:23
Filepath://ppdys1122/BgPr/OUP_CAP/IN/Process/0002108607.3d3

1

Introduction

1.1. Goals

This book stems from a belief that linguistic semantics is a beautiful field, that the tools used to study formal semantics have yielded a rich body of results about fascinating and subtle data, that the field continues to produce exciting new insights at an impressive rate,[1] and that there are simple and

[1] Readers wishing a taste of many of the ongoing developments in formal semantics and in the syntax/semantics interface might want to look at the journals *Linguistics and Philosophy* (Springer), *Natural Language Semantics* (Springer), *Journal of Semantics* (Oxford University Press), and *Semantics and Pragmatics* (online journal, available at <http://semprag.org/>), among many other journals. Regular conferences at which cutting-edge research is presented include the annual Semantics and Linguistic Theory (SALT) conference, Sinn und Bedeutung (also annual), the biannual Amsterdam Colloquium for Language, Logic, and Information, and Semantics of Underrepresented Languages of the Americas, as well as most of the more general regular linguistics conferences. Of course, most of the work in these venues will not be accessible to a student just learning formal semantics, but it is hoped that this book will give a large part of the necessary background for following at least some of this research. In any case, a glance at the list of papers in

elegant tools to model how the syntax and semantics of a natural language work together. We begin with a very elementary "fragment" of English and proceed to expand it further and further—adding tools as needed but aiming to keep the basic machinery relatively simple. The goal of proceeding in this way is to account for a domain of data which is sufficiently rich as to show the excitement of studying formal semantics and its interaction with syntax. We note one limitation from the outset: this book concentrates entirely on the analysis of English. The project of modeling the semantics and the syntax/semantics interaction of any single language already provides such a rich set of results that one can hopefully find this limitation justified for an introductory book like this. In fact, the results that have been gleaned from a detailed modeling of one language have in recent years allowed the field to expand so as to provide a wealth of analyses of other languages.[2] This book hopes to give the foundation to approach that literature.

1.1.1. Compositional semantics and (some of) the goals of semantic theory

One of the most striking and fundamental properties of language—any language—is that speakers have the ability to produce *and understand* an unlimited number of expressions that they have never produced or heard before (indeed many of these will have never before been uttered by anyone). This simple point is stressed in just about every introductory linguistics textbook, often phrased this way: "A speaker of a language is able to recognize as well-formed an unlimited number of expressions that s/he has never heard before." Examples that demonstrate this are easy to construct. One can, for instance, note the existence of expressions like those in (1):

any of these venues can give the reader a taste of the richness of the domain of inquiry within linguistic semantics.

[2] Much cross-linguistic semantic work can be found in the journals and conference proceedings cited in footnote 1. An early edited volume on this is Bach et al.'s *Quantification in Natural Languages.* There is now also an annual conference Semantics of Underrepresented Languages of the Americas with published conference proceedings. And many of the specialized conferences on individual languages and language families regularly include work on semantics.

(1) a. the tallest linguistics major
b. the tallest linguistics major who is graduating in December
c. the tallest linguistics major who is graduating in December who is enrolled in formal semantics
d. the tallest linguistics major who is graduating in December who is enrolled in formal semantics who took phonology last semester . . .

One can keep forming longer and longer expressions like this by adding new relative clauses (each of the phrases that begin with *who* here is what is commonly known as a *relative clause*). But while this is often put in terms of a speaker's ability to recognize that these are well-formed, that is surely only part of the story. Even more interesting (at least to a semanticist) is the fact that speakers know how to *interpret* these expressions. The rule system that speakers have unconsciously learned is hardly just a system to determine whether a given string of words is an expression of the language in question (here English); language would be quite useless if it were just a collection of meaningless strings.

And so, in modeling what a speaker of English "knows" (in an unconscious sense, of course) about her/his language we want to predict how it is that s/he can understand expressions like those in (1) no matter how many relative clauses they contain. Thus speakers obviously have as part of their knowledge a finite set of basic items—call these the *words* and call the collection of the basic items *the lexicon*. (Here and for most of this text we ignore the distinction between words and morphemes.) Since the lexicon is finite, the meanings of the basic items can be learned on a case-by-case basis. But this obviously cannot be the case for the larger expressions: there has to be some systematic set of principles that speakers have that allows them to understand their meanings on the basis of the meanings of the smaller parts (ultimately the words) that make them up. This is the system which is called the *compositional semantics*—and one of the jobs of a theory of the semantics (of any language) is to model the rules and/or principles which allow speakers to understand an unlimited number of expressions. This book is primarily about just this.

Let's look a bit more at the expressions in (1). When a speaker utters any of these expressions—perhaps as part of a fuller sentence like in (2)—the act of uttering these expressions takes place in a fuller discourse context, and we understand them relative to facts about that context:

(2) We need to make sure to order academic regalia which is long enough to fit the tallest linguistics major (who is graduating in December (who . . .))

The role of context will be discussed more formally at various points in the text, but the informal notion of a speech or discourse context is clear enough. So suppose we are using the expressions in (1) in a context in which it is obvious that we are concerned with the students at Brown University. Given this (or any other context), we can see that any speaker of English immediately knows some interesting facts about these expressions—facts which our model of the compositional semantics needs to account for. Take for instance (1a). It refers to some unique individual.[3] The hearer may well not know who exactly that is—in fact the speaker might not either (as is clear in a context like (2)). But both parties assume that there is a particular individual (and only one) referred to by each of these expressions. And there are many other inferences that can be drawn from these. For example, we immediately know that if the individual described by (1a) is Nora, then either she's also the individual described by (1b) or else she is not graduating in December. Moreover, if Nora is not the person picked out by (1b) then whoever that person is, s/he must be shorter than Nora. Similarly, with each successively longer phrase we either refer to the same person, or to one who is shorter. Suppose that Zorba is the person described by (1b). We know that he is shorter than Nora, and also know that if he is not the person described by (1c) then he is not enrolled in formal semantics. And whoever the (1c) person is—let's say Otto—Otto must be shorter than Zorba. The addition of each successive relative clause either keeps the referent constant or allows shorter and shorter people to "rise to the top." This kind of knowledge is automatic and immediate, and it is the job of a model of the compositional semantics to explicitly account for inferences like this.

We won't give a serious account of any of this at this point, but can hint at one possible account. Suppose that an expression like *linguistics major* refers to some set of individuals. (Readers not familiar with basic notions of set theory should consult the Appendix to this chapter.) When this set is put together with *the tallest* (pretend that *the tallest* is a single word here), the entire expression ends up referring to the tallest member of that set. Nothing

[3] In reality there could conceivably be two individuals of exactly the same height. But use of the expressions in (1) does seem to assume that there is a unique referent for these. This is sometimes called a *presupposition*; these are rather odd expressions if the speaker knows that there are two individuals with exactly the same height (in that case the speaker might have said *the two tallest linguistics majors*).

surprising so far. But what is more interesting is what happens with the addition of further relative clauses. It seems plausible that something like *who is graduating in December* also refers to a set (obviously, the set of December graduates). The above facts will make sense if the compositional semantics first combines the two sets (the set of linguistics majors and the set of December graduates) and *intersects* them to give a new set. (The *intersection* of two sets is all things that are in both sets; again see the Appendix.) So (1b) ends up picking out the tallest member of that set. It is now possible to demonstrate that the system correctly predicts that if the referent of (1b) is not Nora, it can only be because she is not graduating in December. For if Nora is taller than anyone in the linguistics major set (call that L) then she is taller than anyone in the intersection of L with the December graduates (call that D). After all, everyone who is in that intersection of L and D is also in L. So if Nora is not the referent of (1b) it can only be that she's not in the intersection of D and L, and since she's in L (by assumption) it follows that she can't be in D. It also follows that if (1b) refers to Zorba, he must be shorter than Nora. By the definition of intersection, if Zorba is in the intersection of D and L he is in L, but we already know that Nora is taller than everyone else in L. All of this is very simple logic that we—the linguists—can work out in the form of an informal proof as above. It could also be worked more formally if one were so inclined. Pedantic though it may seem, it shows that our compositional procedure (which involves intersecting two sets) can be used to correctly model inferences that speakers of English effortlessly make.

Moreover, the appeal here is that this is perfectly general and extends no matter how many new relative clauses are added. Take (1c). The semantics set-up above extends immediately to this. The new relative clause in (1c) is *who is enrolled in formal semantics*. This picks out yet another set—and so this now intersects with the set that we already formed for (1b). The fact that the referent of (1c) can either be Zorba or someone shorter than Zorba follows by the same logic shown above; the reader can work out the details. And the procedure can be repeated over and over no matter how many relative clauses are introduced.

1.1. Because this example is just meant to illustrate the notion of a *compositional semantics*, we have made some assumptions about the order in which the semantics put things together without justifying them. Suppose that rather than the way it was set up here, the meanings of the two relative clauses (1c) first combined, and then that combined with *linguistics major.* Would that make any difference to the basic semantic compositional picture that we have set up here? Would the procedure extend correctly to (1d)?

1.1.2. Direct Compositionality—and its role in this text

This book has a rather ambitious set of goals. On the one hand, I intend this to be a stand-alone text for anyone wishing to have an introduction to formal semantics, compositional semantics, or what is commonly known as *the syntax/semantics interface*. In other words, we will be asking (as in the above example) what a compositional semantics might look like: how can we model the tools available (again, of course, unconsciously) to speakers of a language that allow them to compute meanings of larger expressions from the meanings of the smaller ones that make them up. What are the formal ways in which meanings combine? And what are the types of objects that we need in order to model that? (For example, the discussion above shows that some simple tools of set theory can be useful.)

But while most semanticists agree that (in general) the meaning of a larger expression is built in some systematic way from the meanings of the parts that make it up, just exactly how the syntactic system of a language and the compositional semantics work together is a matter of considerable controversy, and is one of the central questions addressed in this book. And so this book takes one particular point of view on this: the point of view known as *Direct Compositionality*. This view was explored perhaps most notably in Montague (1970) and was either generally accepted or at least taken as a serious desideratum in much of the work in linguistic formal semantics throughout the 1970s and 1980s (particularly work in what was then known as the Montague Grammar program). It was also taken as the foundation for semantics in syntactic theories such as Generalized Phrase Structure Grammar (Gazdar, Klein, Pullum, and Sag 1985), and is assumed in a large body within current grammatical theories that go under the rubric of Categorial Grammar, Type-Logical Grammar, and other related theories.

To elucidate, a fairly uncontroversial claim is that the grammar of any natural language is a system of rules (or principles, if one prefers) that define the set of well-formed expressions of the language (i.e., the syntax) and a set of rules (or principles) pairing these with meanings (i.e., the semantics). The hypothesis of Direct Compositionality is a simple one: the two systems work in tandem. Each expression that is proven well-formed in the syntax is assigned a meaning by the semantics, and the syntactic rules or principles which prove an expression as well-formed are paired with the semantics which assign the expression a meaning. (An interesting consequence of this view is that every well-formed syntactic expression does have a meaning.[4]) It is not only the case that every well-formed *sentence* has a meaning, but also each local expression ("constituent") within the sentence that the syntax defines as well-formed has a meaning. Of course putting it this way is arguably not much more than just a slogan: the empirical content of this depends in part on just how the syntax works and what one takes to be a *meaning*. This will be filled in as we proceed. It might also seem at first glance that the hypothesis of Direct Compositionality is a fairly trivial one. But in fact it is not always immediately obvious how to give a Direct Compositional analysis. Even the example in 1.1.1 is a case in point. If the syntax and semantics work together, then the analysis given above leads to the conclusion that *in the syntax* a relative clause like *who is graduating in December* combines with *linguistics major* rather than with *the tallest linguistics major.* But this very question regarding the syntax of relative clauses has been debated in the literature since the 1960s, and many researchers have claimed that the syntactic constituent structure of *the tallest linguistics major who is graduating next year* is not the structure that was used above for the semantic analysis. We will actually revisit this particular question in later chapters (see, e.g., Chapter 13.2). So one of the goals of this book will be to see what it takes to give Direct Compositional analyses of a variety of constructions.

While the material in this book is generally exposited from the Direct Compositional point of view (along with discussion of the challenges to this hypothesis), the book is also intended to be a perfectly reasonable

[4] Of course one of the earliest arguments in Generative Grammar *for* divorcing the syntax from the semantics (and thus a putative argument *against* Direct Compositionality) is based on the claim that there are well-formed expressions that don't have any meaning (Chomsky 1957). This is addressed in Chapter 7.1.

Comp. by: PG4118 Stage : Revises1 ChapterID: 0002108607 Date:16/5/14 Time:06:44:23
Filepath://ppdys1122/BgPr/OUP_CAP/IN/Process/0002108607.3d10

OUP UNCORRECTED PROOF – REVISES, 16/5/2014, SPi

stand-alone textbook for any formal semantics course. Thus it is suitable for any linguistics student or linguist wanting a ground-up introduction to formal semantics, and for a philosophy or logic student wanting a background in formal semantics within linguistics. In the service of being a stand-alone text in modern formal semantic theory, the book will, where relevant, also develop the mechanics of at least one fairly standard non-Direct Compositional theory of the syntax/semantics interface. This is done especially in Part III (Chapters 13–15) where some phenomena are discussed from both direct and non-Direct Compositional points of view. There are several reasons for expositing parallel Direct and non-Direct Compositional accounts of some domains. One is to enable readers to approach the wide range of literature written from either point of view. Second, this allows for a serious comparison of two different approaches. Third, learning more than one set of details for the analysis of any construction allows for a deeper understanding of the basic generalizations and results—generalizations which often transcend the particulars of one theoretical implementation. Finally, a student who has already learned formal semantics from a non-Direct Compositional point of view can hopefully also profit from this book by seeing an interesting fragment of English explicitly analyzed from the Direct Compositional point of view.

1.2. A brief note on the history of semantics within modern linguistic theory

The subfield of semantics as a core field in modern linguistic theory is relatively recent and is one of the fastest growing subfields.[5] Early work within the general enterprise of generative grammar had little to say about semantics. To be sure, by the end of about the 1960s and the early 1970s there was considerable discussion as to how the syntax and the semantics interacted; such discussion was mostly framed in terms of a debate between *Generative Semantics* (see, e.g., McCawley 1971; Lakoff 1971) and *Interpretive Semantics* (see Chomsky 1970; Jackendoff 1972). We will not discuss the content of that debate here, but much of the work framed within these

[5] A much more extensive and authoritative history of the development of formal semantics within modern linguistic theory can be found in Partee (forthcoming).

two competing points of view did not incorporate a systematic view of the semantics itself. Of course, the linguistics literature during that period contained many seminal observations about semantic notions such as scope, negation, and "binding," but these were generally not embedded within a full-blown theory of semantics designed to capture semantic notions like entailment and truth conditions (see Chapter 2), although they easily could have been embedded into such a theory. The fact that semantics was not taken a subfield in and of itself during this period comes—at least in part—from Noam Chomsky's emphasis on syntax during the early development of generative grammar. Chomsky (1957) explicitly rejects the notion that semantics is relevant in the construction of grammars, and this notion persisted for quite some time.

It is probably fair to say that modern formal semantics as a subfield within linguistic theory began in the early to mid-1970s with the cross-fertilization of linguistic theory and philosophy of language (including semantics) sparked by Barbara Partee, Richmond Thomason, David Lewis, and others. Partee's work was a particularly influential bridge between linguistics and philosophy as she had originally been a student of Chomsky's at MIT and always had a strong interest in the connections between language and logic, and hence in topics like quantifiers, negation, etc. As an assistant professor at UCLA, she became acquainted with the seminal work of Richard Montague, a philosopher and logician who (among his many other contributions within philosophy and logic) had a major interest in modeling the semantics (and the syntax/semantics interaction) of natural language (although Montague himself dealt only with English). In fact, the program of Direct Compositionality is advocated in his work (see especially Montague 1970). We will have more to say about his specific contributions as this book proceeds; for now, we note that one of the appeals of his work from the point of view of a linguist was his notion that the semantic composition of natural language reflects and respects its syntax. Partee saw the relevance of Montague's work to linguistic theory and wrote a series of papers aimed at synthesizing some of the insights from Montague's work with results within Transformational Grammar (see, for example, Partee 1973). At the same time, the appearance of Lewis (1970), Stalnaker and Thomason (1973), and other work in the philosophy of language also helped launch modern formal semantics and cement its connection to linguistic theory. Such work within philosophy as well as Partee's early group of students (both at UCLA and later at the University of Massachusetts) continued the tradition, broadening the domain of inquiry and results. From there was born the enterprise known as Montague grammar[6] which eventually gave rise to the more general subfield of formal semantics. Montague himself died in 1971,[7] and the field of formal semantics evolved in many ways quite different from the original work in Montague grammar. Nonetheless, many of the basic tools of linguistic formal semantics as it is developed to this day stem from some of this early work cited above. Since the late 1970s the field has blossomed, and is now within linguistics generally considered as one of the core areas along with at least phonology and syntax.

1.3. The notion of a "fragment" and its use in this text

Inspired by the work of Montague in papers such as Montague (1973), much work in formal semantics within the 1970s and 1980s took it as axiomatic that a goal was to formulate fully explicit grammars (in both syntactic and semantic detail) of the fragment of the language one is concerned with (English in most such work). The term "fragment" got extended to mean not only the portion of the language being modeled, but also the portion of the grammar being proposed as an explicit account of the facts. The strategy of writing fragments (of grammars) has the advantage of giving an explicit theory which makes testable predictions, and of making theory and/or proposal comparison easier.

Unfortunately, the goal of formulating fully explicit fragments went out of style during the last two decades or so. This is in part due to the fact that linguistic theories often promised that many of the particular details did not need to be stated as they would fall out from very general principles. It is certainly reasonable to hope that this is ultimately true, but the relevant principles often go unstated or are stated only rather vaguely, making it extremely difficult to really compare proposals and/or evaluate theories and theoretical claims. Having rules and principles be as general as possible is, of course, highly desirable. But this does not mean that they should not be

[6] An excellent introduction to the general program of Montague semantics and an explication especially of Montague (1973) can be found in Dowty, Wall, and Peters (1981).

[7] Montague was murdered on March 7, 1971. No arrest was ever made in conjunction with the murder.

formulated explicitly—only that more mileage will be gotten out of explicit formulations.

The present text is therefore committed to trying to revive the notion of explicit fragment construction. We cannot promise to give every detail of the domain of English syntax and semantics we are trying to model. Some parts will be left tentative, some stated informally, and some simply omitted. Nonetheless, the goal is to give a reasonable amount of an explicit fragment. We will therefore periodically take stock by summarizing the fragment constructed so far, and a full summary is provided at the end of Part III.

*1.4. An intriguing puzzle

This introductory chapter concludes with an illustration of a puzzle, a solution to which is proposed in Chapter 15.6. However, the goal here is not to champion any one particular solution, and readers may safely skip this section and return to the data only in Chapter 15.6. But we include this in the introductory remarks for the reader who wants a preview of just what kinds of complex and subtle data a theory of syntax and semantics ultimately hopes to account for. To fully appreciate the particular puzzle here, one should keep the following in mind. The contrasts are quite real; the judgments have been checked with many speakers over the years by myself and many others. Yet—like other subtle facts in syntax, phonology, and semantics—these are not generalizations which we have ever been consciously taught nor even generalizations that most of us are even aware of until we see them in a linguistics course (or book). What, then, is there about our unconscious knowledge of the grammatical system that predicts these judgments? This is the sort of puzzle that theories of semantics and its interaction with syntax ultimately seek to solve.

So, consider what we will call the *A-B-C party scenario*. I go to a small party consisting of only myself and three married couples: Alice and Abe, Betty and Bert, and Cathy and Carl. I learn that Alice and Abe met each other only a few years ago, and similarly for Cathy and Carl. But interestingly, I also find out that Betty and Bert have been sweethearts since childhood. I like Betty a lot, and spend a good part of the evening talking to her. The next day, you ask me how I enjoyed the party and if there was anyone that I especially enjoyed meeting. I certainly can answer with (3):

(3) Oh yes, I especially enjoyed talking to Betty.

But now suppose that I can't remember Betty's name, although I do remember that her husband's name is Bert. I can answer with either (4a) or (4b):

(4) a. Oh yes, I especially enjoyed talking to—oh, I can't remember her name—you know, the woman who is married to Bert.
b. Oh yes, I especially enjoyed talking to—oh, I can't remember her name—you know, the wife of Bert.

((4b) would sound more natural if we substituted *Bert's wife* for *the wife of Bert*; this will not impact on the ultimate point and the exposition is simplified using (4b).)

Now, let us tweak the scenario slightly and assume that I am one of those people who just doesn't remember names very well. As a result, I remember neither Betty's name nor Bert's name, although I do remember the interesting fact that they are the only couple at the party who have been sweethearts since childhood. As an answer to your question, (5) would be quite natural:

(5) Oh yes, I especially enjoyed talking to—oh, I can't remember her name—you know, the woman who is married to her childhood sweetheart.

But what is striking is that I can't answer with (6):

(6) *Oh yes, I especially enjoyed talking to—oh I can't remember her name—you know, the wife of her childhood sweetheart.

(We are taking liberties with the * notation here. This is generally used in works in syntax to indicate something that is ill-formed. (6) in fact is fine, just not on the intended reading; and we will continue to notate a sentence with an asterisk in front of it when we mean "bad on a particular reading" provided that it is clear what the intended reading is.) It should be noted that some speakers find the contrast rather subtle but there is general agreement that (6) is stranger than (5).

All of these examples contain various extra material (the parentheticals, etc.) which are there to make them sound natural and conversational. But as we proceed it will be convenient to strip away the parts that are irrelevant to figuring out the semantics, so we can recast (6) as (7)—also impossible as an answer to the question and in this context:

(7) *I especially enjoyed talking to the wife of her childhood sweetheart.

Of course, (7) is a perfectly good sentence, but it cannot be used in our party scenario as a way to identify Betty.

Since some readers do find the contrast subtle, two points are worth noting. First, one should resist the temptation to recast (7) in one's mind as *I especially enjoyed talking to the one who's the wife of her childhood sweetheart* or *I especially enjoyed talking to the woman who's the wife of her childhood sweetheart*. That would be cheating; the point is not to find a closely related way to say the same thing but to notice that the actual way in (7) contrasts with *I especially enjoyed talking to the woman who is married to her childhood sweetheart* (and contrasts with the above variants too). As to why these variants are good, we return to that shortly. Moreover, while the contrasts above may be subtle for some speakers, there is a related mystery where the facts dramatically pop out. Thus take (8) in the same scenario, where the only people at issue are Alice, Betty, and Cathy:

(8) Betty is the only woman who is married to her childhood sweetheart.

This can be making two different claims. The obvious one in this scenario is that Cathy is not married to Cathy's childhood sweetheart, and Alice is not married to Alice's childhood sweetheart. The other is the "non-polygamous" reading: it asserts that Bert (or whoever Betty's husband might be) has only one wife. Since we (generally) assume that people have just one wife, this reading (given standard assumptions) is not the first one that someone would think of since it is less likely to be conveying any interesting information. But despite the fact that the non-polygamous reading is the less obvious one for (8), it is the only reading (or at least the one that pops out first) for (9):

(9) Betty is the only wife of her childhood sweetheart.

Why should that be? We'll put (8) and (9) aside for the moment, and return to the simpler case of (7).

So the mystery is why (7) is bad as a way to identify Betty. This is especially puzzling in that both (4b) and (5) are perfectly good—or, to give their stripped-down versions, (10) and (11) are both fine. Each one differs minimally from our bad case, yet neither of these two has any problem.

(10) I especially enjoyed talking to the woman who is married to her childhood sweetheart.

(11) I especially enjoyed talking to the wife of Bert.

So surely there is nothing incoherent or wrong with the meaning that (7) is trying convey, for (10) is just a slightly different form and conveys exactly this meaning. Hence the puzzle has something to do with the mapping between syntax and semantics: why one is a good way to package the relevant information while the other is not.

We can informally recast the puzzle in the following way. Compare the two expressions *the woman who is married to Bert* and *the wife of Bert*. (Following a long tradition within linguistics, we will refer to these as NPs, which comes from "noun phrases." They are also in much modern literature referred to as DPs, for "determiner phrases," but we stick to the more traditional terminology in this text.)[8] Both of these can correspond to meanings that we can (roughly and informally) represent as (12):

(12) the x: x is a woman and x is married to Bert

But while the object NP in (10) can be represented as in (11), the object NP in (7) cannot:

(13) the x: x is a woman and x is married to x's childhood sweetheart

The basic phenomenon here was discussed in, among others, Jacobson (1977) (where it was called *Langendoen's constraint*), Chomsky (1981 under the rubric of *i-within-i condition*), and many since. As there seems to be nothing wrong with the meaning, we can assume that the phenomenon in question has something to do with the way the syntax and semantics interact.

Notice that we have given a kind of formula (and one that uses a "variable" *x*) to represent the meanings in question, but for now we should think of these simply as placeholders to bring out the intuition. After all, recasting *her* in the above examples as *x* doesn't really immediately give us the tools for computing the meanings of the expression: we have traded a

[8] In theories which use the term DP, the NP is used instead to refer to material after the Determiner; e.g., *mother of Romeo* in an expression like *the mother of Romeo.* Here we will be calling this simply N (i.e., a "noun") and allowing terms like N to refer both to simple material consisting of just one word and to complex material. This is discussed further in Chapter 6. We are aware that this will initially cause some confusion to a reader who is used to using "NP" to mean a noun and its complement, but it is well worth becoming fluent in both sets of terminologies. The terminology here is the traditional one found in large amounts of literature not only in syntax and semantics but in neighboring fields like psycholinguistics, philosophy, cognitive science, etc.

pronoun *her* for a variable *x*. But this accomplishes little until we have a way to think about what a variable like *x* means. (Indeed this is explored in detail in Chapters 9 and 15, including developing an alternative view that does not make use of variables in the semantics.) We thus caution that formulas like (12) and (13) are best seen simply as informal and helpful ways to bring out the intended meanings. Similarly, one often sees indices used in the literature as a way to bring this out; one will find discussions using the notation in (14) and (15) to make the point, where the indexation in (14) indicates a good possible reading for the NP while (15) cannot be understood in the intended way:

(14) the woman_i who_i is married to her_i childhood sweetheart

(15) *the wife_i of her_i childhood sweetheart

Much work in grammatical theory actually assumes that NPs and pronouns come with (obviously silent) indices in the syntax; here we will be using indices from time to time simply as a way to notate intended readings without any commitment to their being actual pieces of grammatical machinery.

Before leaving this (for now), there's one other interesting point to notice. However we ultimately state the principle, the claim is that an NP like *the wife of her childhood sweetheart* cannot correspond to the meaning shown earlier in (13):

(13) the x: x is a woman and x is married to x's childhood sweetheart

But one might think that this is not really correct, since it is in fact just fine to use (16) as a way to identify Betty:

(16) I especially enjoyed meeting the woman who is the wife of her childhood sweetheart.

This point was made earlier; many speakers on reading (7) tend to recast it in their minds as (16). Similarly, (17) is impeccable on the understanding where *her* is Betty:

(17) Betty is the wife of her childhood sweetheart.

But a close reflection reveals that this does not threaten the generalization. Again, using indices or variables simply as a convenient way to elucidate the point, it is easy to see in (17) that *her* need not be "coindexed with" (or "correspond to the same variable as") *wife* but rather it just refers directly to *Betty*. That is, we can represent it as in (18a) using indices, or as in (18b) using the more spelled-out formula (though still quite informal).

(18) a. Betty$_j$ is [$_{NP}$ the wife$_i$ of her$_j$ childhood sweetheart].
b. Betty, y [y = the x: x is a woman and x is married to y's childhood sweetheart]

Since we are asserting identity between Betty and the person married to Betty's childhood sweetheart, it of course follows that Betty is married to Betty's childhood sweetheart and so the full sentence (16) will end up with the relevant meaning.[9] But the claim that the object NP itself (*the wife of her childhood sweetheart*) does have the meaning represented in (12) is not threatened. The same point holds for (17), whose meaning can be represented as (19a) or (19b).

(19) a. the woman$_j$ who$_j$ is the wife$_i$ of her$_j$ childhood sweetheart
b. the y: y is a woman and y = the x: x is a woman and x is married to y's childhood sweetheart

Is there a way to confirm that this is the right sort of explanation for these apparent counterexamples? Indeed there is, and it centers on the contrast between (8) and (9) which was discussed earlier. We leave it to the interested reader in the exercise to play with this and get a sense of why (8) is ambiguous and (9) is not. Having completed that, one should be able to see how it is that this gives support for the explanation offered above as to why (17) does not threaten the claim that *the wife of her childhood sweetheart* cannot correspond to the meaning shown informally in (12).

***1.2.** Work out—using the informal representations either with indices or the representations with variables—why it is that (7) is ambiguous and (8) is not. Of course you will need to think a bit about how to treat *only*, but nothing very complex is required. You can be perfectly informal in your treatment of *only*, but you should be able to get a feel for why these two differ.

[9] This general observation—although for a slightly different case—was made in Postal (1970) who distinguished between "presuppposed" coreference and "asserted" coreference. Here the fact that *Betty* and *the wife of her childhood sweetheart* end up "referring" to the same individual is exactly what the sentence is asserting.

As noted at the outset of this section, the goal here is just to provide a mystery to whet the reader's appetite; the tools needed to provide a hypothesis as to the explanation of the mystery will be developed later.

1.5. Appendix: Sets and Functions

1.5.1. Sets, members, and subsets

Since the notions of sets and of functions are crucial throughout this book, some formal definitions and discussion are provided here for readers not entirely familiar with these notions. We begin with the notion of a *set*. A set is simply any collection of objects (it can have a finite number of objects, an infinite number, or none at all). For example, we can talk about the set of positive integers less than 10; sets can be notated by listing the members and enclosing the list in curly brackets: {1,2,3,4,5,6,7,8,9}. The order in which they are listed makes no difference; a set is just a collection of things without any order. So if we were to write {2,5,3,4,9,7,8,1,6}, this names the same set. Each integer in this set is called a *member* or an *element* of the set. If we were to name this set A, then the notation $4 \in A$ means that 4 is a member (or element) of A. Something either is or is not in a set; it makes no sense to say it occurs twice (or more) in the set. Note also that a set can have a single member; this is called a *singleton set*. Thus {4} is the set with only one member; this set is distinct from 4 itself. (4 is a member of {4}.)

A set can have an infinite number of members; the set of positive integers for example is infinite. Obviously this can't be named by listing the members. One can in this case specify the set by a recursive procedure. Call the set I, then one can specify I by two statements: (a) (what is known as the base step): $1 \in I$, and (b) (the recursion step) if $n \in I$ then $n+1 \in I$. (It is understood when one lists things this way that nothing else is in I.) One will also often see a notation which describes rather than lists the members. For example, we can write the following set, call it B: {x|x is a New England state}. This names a finite set, and so we could also give B in list form as follows: {Maine, New Hampshire, Vermont, Massachusetts, Rhode Island, Connecticut}. These two are just different notations for naming the same set. This can also be used, of course, for infinite sets. Take, for example, the set {x|x is an integer and x > 9}. This names the set of integers greater than 9. And, a set can have no members. There is only one such set; its name is *the null set* or *the empty set*, and is generally written as Ø. Of course, there are other ways one can describe the null set. For example, the set of integers each of which is greater than 9 and less than 10 is the empty set. The *cardinality* of some set refers to the number of elements in that set; the notation |B| means the cardinality of B. Hence, given our set B above, |B| is six.

Take some set A. Then a *subset* of A is any set all of whose members are also in A. Suppose, for example, we begin with a set C which is {1,2,3}. Then {1,2} is a subset of C, as is {1,3} and so forth. The notation for the subset relation is $\subseteq$. The full definition of subset is as follows: B $\subseteq$ A if and only if every member of B is a member of A. From this it follows that every set is a subset of itself (so for the set C above, one of its subsets is the set {1,2,3}). It is, however, sometimes convenient to refer to those subsets distinct from the original set; in that case we can talk about a *proper subset* of some set. The symbol for this is $\subset$, so B $\subset A$ if and only if B $\subseteq$ A and B $\neq$ A. Since the definition of subset says that B is a subset of A if and only if everything that is in B is also in A, it follows that if nothing is in B then B is a subset of A. thus the null set is a subset of every other set. Sets themselves can have sets as members, and so one can talk about the set of all subsets of a set A. This is called the *power set* of A, written as $\mathscr{P}$(A). For example, given the set C above, $\mathscr{P}$(A) = {Ø, {1}, {2}, {3}, {1,2}, {1,3}, {2,3}, {1,2,3}}.

***1.3.** If a set A has n members, then the number of subsets of A is 2^n. Try to see why this is true. *Hint:* for every member x of some set A, then for each subset B of A, x is either in B or is not in B.

1.4. How many members does the following set have: {Ø}?

***1.5.** What is $\mathscr{P}$(Ø)?

We will also have occasion to talk about the reverse of the subset relation—i.e., the *superset* relation. A is a superset of B if and only if B is a subset of A. The notation for this is A $\supseteq$ B. Once again this is defined in such a way that every set is a superset of itself; a superset of B which is not identical to B is called a *proper superset*, and the notation for this is $\supset$.

1.5.2. Union, intersection, and complement

Take any two sets A and B. Then there is a set C which consists of everything that is in A and everything that is in B. This is called the *union* of A and B, and is written A ∪ B. For example, if A is {1,2,3} and B is {2,4,6} then A ∪ B is {1,2,3,4,6}. Moreover, for any two sets A and B the *intersection* of A and B is the set of all things that are in both A and B. This is written A ∩ B. So, for example, in the case directly above, the intersection of A and B is {2}. Or, if we were to intersect the set of integers which can be evenly divided by 2 (the set of even integers) with the set of integers which can be evenly divided by 3, we end up with the set of integers that can be evenly divided by 6.

1.6. a. For any two sets A and B such that A ⊆ B, what set is A ∪ B?
b. For any two sets A and B such that A ⊆ B, what set is A ∩ B?

One final useful notion here is the *complement* of a set. The complement of some set A is the set of all things which are not in A (this is sometimes notated as A′). Usually one talks about this notion with respect to some larger domain. Strictly speaking, the complement of {1,2,3} would include not only all integers greater than 3 but also all sorts of other numbers (like 1/3), the sun, my dog Kiana, and the kitchen sink. Rarely are we interested in that sort of set; so in practice when one talks about "the complement of some set A" this is generally with respect to some larger set B of which A is a subset. Then the complement of A refers to all things in B that are not in A (this is notated as B-A). For example, when restricting the discussion to the set of positive integers, the complement of {1,2,3} is the set of all integers greater than 3.

1.5.3. Ordered pairs, relations, equivalence relations, and partitions

Sets are unordered collections of objects. But it is quite useful (as will become very apparent as this book proceeds) to be able to talk about pairs of objects that are ordered in some way. An *ordered pair* is just that:

it is two objects with some ordering between them. If the two objects are a and b, then (a,b) is an ordered pair; (b,a) is a different ordered pair. An ordered pair need not contain distinct items: (a,a) is an ordered pair. In applying this to actual natural relations that exist in the world we are generally interested in *sets of ordered pairs*. (One can generalize this notion to ordered triples and so forth; an *ordered n-tuple* means an ordered list of n items.)

This notion is easiest to grasp with some concrete examples. Take again the set {1,2,3}, and take the relation "is greater than." Then this can be seen as a set of ordered pairs; if we are restricting this to items from our little 1-2-3 set, this would be the set {(2,1), (3,1), (3,2)}. Now suppose we instead take the following set of ordered pairs: {(2,1), (3,1), (3,2), (1,1), (2,2), (3,3)}. Then (restricting this again to our 1-2-3 set) we have now actually listed the relation "is greater than or equal to." Or, take the set {(1,1), (2,2), (3,3)}. That is the relation "is equal to" (defined for the set of integers {1,2,3}).

In other words, what we are calling a *relation* is just some set of ordered pairs. In the example above, both the first and second member of each ordered pair was drawn from the same set (the set {1,2,3}). But this is not necessary; we can have a set of ordered pairs each of whose first member is drawn from some set A and the second member from some set B where A and B are different (they can, but need not, have some of the same members). For example, the relation "is the capital of" is a relation between cities and states; it can be expressed as a set of ordered pairs of the general form {(Providence, Rhode Island), (Boston, Massachusetts), (Springfield, Illinois), (Pierre, South Dakota), . . . } (the . . . here is a shorthand for the remaining 46 pairs).

Take two sets A and B (they could be the same set or different). Then A x B refers to the set of all ordered pairs whose first member is in A and whose second member is in B. (This is also called the *Cartesian product* of A and B.) As in the case above, it is helpful to give the intuition of this by coming up with some concrete example. Suppose we take as our set A some group of professors—say, Professor Magoo, Professor Carberry, and Professor Glazie. Call that set P (for shorthand, let's call its members m, c, and g, so P is the set {m,c,g}). Now suppose we have a set S which consists of three students who we will just indicate as x, y, and z (so S = {x,y,z}). Then P x S = {(m,x}, (m,y), (m,z), (c,x), (c,y), (c,z), (g,x), (g,y), (g,z)}. Suppose that Magoo wrote a letter of recommendation for all three students, Carberry wrote one for only y, and Glazie wrote one for y and z. Then the relation "wrote a recommendation for" is a subset of P x S, and is the set of ordered pairs {(m,x}, (m,y), (m,z), (c,y), (g,y), (g,z)}.

More generally, we define a *relation* (between members of A and B) as any subset of A x B. There are some special and interesting properties that can be defined when the two sets are the same. That is, we are now looking at subsets of A x A. Consider a relation R (some subset of A x A) which is such that for all x in A, (x,x) is in R. Such a relation is called a *reflexive* relation. (These need not be the only kinds of pairs to be in R for R to be reflexive; other pairs can be in there too.) For example, if talking about the set of integers again, the relation "is greater than or equal to" is reflexive; for all numbers n, (n,n) is in the set of ordered pairs described by that relation. A relation R is called *irreflexive* if for all x in A, (x,x) is *not* in R. Further, consider any two members x and y (both members of A). Then if it's the case that for all x and y if (x,y) is in R then (y,x) is also in R, the relation is called *symmetric.* Imagine, for example, a lovely world with no unrequited love. Then *is in love with* is symmetric in that world. If our set were {m, c, g, and p}, then if the pair (m,c) were in our relation R (i.e., "is in love with") the fact that R is symmetric means that (c,m) is also in R. (Notice that our definition neither requires (c,c) to be in R nor excludes that; either is possible.) Or, to look at a relation which is symmetric by definition: consider the relation *is a sibling of*. (While *is a sibling of* is symmetric, *is a sister of* is not. Why not?) One final useful definition is a *transitive relation*. A transitive relation R is one for which for every x, y, and z, if (x,y) is in R and (y,z) is in R, then (x,z) is in R. (The relation "is greater than" is transitive, as is the relation "is greater than or equal to").

Any relation R which is reflexive, transitive, and symmetric is called an *equivalence relation*. As an example of such a relation, consider the set of students (call it S) at an elementary school that services grades 1 through 6. Then "is in the same grade as" is an equivalence relation in S x S. (While it is unusual to use the phrase "in the same grade as" when referring to the same person it seems false to say *Johnny is not in the same grade as himself* so we can see that this relation is reflexive.) It is also obvious that it is symmetric and transitive. Note that this—and any other equivalence relation—divides up the original set (here, S) into a group of non-overlapping subsets. The set of these subsets is called a *partition*. Thus, a *partition* of any set S is a set of subsets of S such that for each distinct subset A and B, $A \cap B = \emptyset$, and the union of all the subsets is S. To show that any equivalence relation induces such a partition, take any x in S and define S_x as {y|(y,x) is in R}. Since R is reflexive, we know that x is in S_x (and hence we know that S_x is guaranteed not to be empty). Moreover, the fact that R is reflexive means that each member of S is guaranteed to be in at least one such subset, so we know that the union of all of these is S. We can further show that for any two such subsets S_a and S_b, they either have no members in common (i.e., they have a null intersection) or they are the same. Thus, take any c which is in both S_a and S_b. By definition, this means that (c,a) is in R and (c,b) is in R. By the fact that R is transitive and symmetric, it follows that (a,b) and (b,a) are in R (the reader can work through the necessary steps). But then, for all x such that (x,a) is in R, (x,b) is also in R. To show this note again that R is transitive. If (x,a) is in R and (a,b) is in R then (x,b) is also in R. Hence given the initial premise that there is a non-empty intersection between S_a and S_b, it follows that everything in S_a is in S_b. That everything in S_b is also in S_a follows in the same way, and so the two are the same set. Each subset in a partition is called a *cell* in that partition.

In the example above, the cells correspond to the different grades. (There don't have to be six cells—it could be that one of the grades has no student in it. But there can be no more than six; recall that by definition a cell can't be empty.) Just as any equivalence relation induces a partition, given any partition one can give an equivalence relation that corresponds to any partition; this is the relation that holds between any two a and b in S such that a and b are in the same cell in the partition.

1.5.4. Functions

A function takes every member of some set A and assigns it a value from a set B (B could be the same set as A, but need not be). This can also be formalized using the notion of a set of ordered pairs. Thus, consider two sets A and B (which again could be the same but need not be). Then, a (total) function from A to B is any set of ordered pairs (i.e., any subset of A x B) such that for each a in A, there is one and only one ordered pair with a as first member. Thus if we think of the function f as assigning to each a in A some values in B, note that the criterion above ensures that each member of A is indeed assigned a value, and is assigned a unique value. A is referred to as the *domain* of the function, and B is referred to as the *co-domain*. For any function f and any a in the domain of f, we write f(a) to indicate the value that f assigns to a. (To use other common terminology, f(a) means the result that one gets by applying the function f to a.) There is no restriction that each member of B must appear as second member of some ordered pair; the term *range* of the function f is the set of all b in B such that there is some a such that f(a) = b. Note that these definitions are such that the range of a function is a subset of the co-domain. In practice (at least in works within linguistics) the terms "range" and "co-domain" are often not distinguished.

As noted above, there is no restriction that each member of B appear as second member of an ordered pair. Nor is there a restriction that it appear only once. If each member of B is used as a value only once (that is, for each b in B, there is a unique a such that f(a) = b) then B obviously can be no smaller than a. It can have more members, or it can be the same size. If the latter is the case, then it also follows that for every b in B, there is some a such that f(a) = b. When both conditions above hold (i.e., for each b in B, there is one and only one a such that f(a) = b, we say that there is a *one-to-one correspondence* between A and B. Note that for any function f which is a one-to-one correspondence, there is a corresponding function f^{-1} which is just the reverse: it is a function mapping each member of B to a member of A such that for all a in A and b in B, if f(a) = b then $f^{-1}(b) = a$.[10]

We will have some occasion to talk about the notion of a *partial function*. A partial function is one where not every member of A is actually assigned a value by f; f is *undefined* for some subset of A. (Of course any partial function f is also a total function with a smaller domain.) We can illustrate this by returning to our earlier example of ordered pairs of US cities and states, where the first member of each ordered pair is the capital of the second. This is a *partial function* from the set of US cities to states (not every US city is a capital). We can reverse it, and have each state as the first member of the ordered pair and the second as its capital (this function could be expressed in prose as *has as its capital*). This is now a total function from

[10] Incidentally, the notion of the availability of a one-to-one correspondence can be used to define what it means for two sets to have the same cardinality. Obviously for two finite sets it is clear what it means to have the same cardinality, since we can count the members. But consider the case of infinite sets. Take the following two sets: A = the set of positive integers {1,2,3, . . . } and B = the set of positive even integers {2,4,6, . . . }. Both are infinite. Surprisingly (when one first hears this) they are also of the same cardinality, because one can establish a one-to-one correspondence between them (each member of A is paired with a member of B by multiplying by 2: we will never run out of members in B).

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)
$$\begin{tikzcd}
\text{Zacky} \arrow[r] & \text{Abe} \\
\text{Yonnie} \arrow[ur] & \text{Bert} \\
\text{Shelley} \arrow[ur] & \text{Carl} \\
\text{Baba} \arrow[r] & \text{David}
\end{tikzcd}$$

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