

1 The Basic Idea

Conditional sentences take two importantly different forms, the *indicative* conditional and the *subjunctive* conditional. Indicative conditionals say that if such and such is the case, or was the case, or will be the case, then something is the case, or was the case, or will be the case. For example, “If he is in Paris, then he is happy” is an indicative conditional. So is, “If he was in Paris, then he was happy.” Subjunctive conditionals say that if such and such were the case, or had been the case, or were going to be the case, then something would be the case, or would have been the case, or would be going to be the case. For example, “If he were in Paris, then he would be happy” is a subjunctive conditional. So is, “If he had been in Paris, then he would have been happy.”

It has often been questioned whether the terms “indicative” and “subjunctive” accurately describe the distinction between the two main kinds of conditional that we are interested in. At first sight this terminology might seem incorrect, because it is not obvious that all so-called subjunctive conditionals have the main verb of their antecedents in the subjunctive mood. In “If he were there,” “were” appears to be in the subjunctive mood. But in “If he had gone,” it is not obvious that “had” is subjunctive; it might be just the past tense of “have” (used to form the past perfect “had gone”). Subjunctive conditionals are sometimes called “counterfactual conditionals,” but that, as we will see, is a mistake, since they are not all counterfactual. Some grammarians adopt the term “open” for the class of conditionals that I am calling “indicative.” Grammarians employ a variety of other terms for the class I am calling “subjunctive,” such as “remote” (Huddleston and Pullum 2002, 738), “hypothetical” (Quirk et al. 1985, 1091; Greenbaum 1996, 340), and “unreal” (Eastwood 1994, 333). I will retain the traditional terminology of “indicative” and “subjunctive” because, as I will argue in chapter 7, it is not entirely wrong.

Not every declarative sentence formed with “if” is truly a conditional. Philosophers have taken some note of sentences like “There is beer in the refrigerator if you want some” in which the “if”-clause merely identifies the occasion for asserting the main clause and does not provide a condition under which the situation described in the main clause can be expected to hold. But there are many other sorts of nonconditional “if”-sentences that cannot be assimilated to this one, such as, “That’s a pretty nice sand castle, if I do say so myself.” It might be good to have a criterion by which to distinguish between conditional “if”-sentences and nonconditional “if”-sentences that did not presuppose a resolution of the semantic issues, so that when we debated the semantics we could be sure we were talking about the same class of sentences. I have no such criterion to offer, but there is no pressing need for one, since the question of where to draw the line has never been an important point of contention between the advocates of the various semantic theories of conditionals.¹

One way to draw the distinction apart from semantics might be to appeal to cross-linguistic comparisons. Surely every human language will exhibit a conditional construction of some sort, whereas the nonconditional functions that the word “if” is recruited to perform in English would be performed in various other ways in other languages. But this approach will not provide a very firm distinction either. There might be a kind of natural transition between true conditionals and certain sorts of nonconditional uses of the conditional construction that in any language encourages the use of the conditional construction for those purposes. For example, sentences such as “There is beer in the refrigerator if you want some” may arise as a foreshortening of conditionals such as “If you want some beer, then you should know that there is beer in the refrigerator.” Other sorts of nonconditional uses of the conditional construction might be peculiar to English, or to Germanic languages, or to Indo-European languages. In any case, I will not discuss nonconditional uses of “if” any further in this book.

In this first chapter, I will explain and illustrate the basic theory of conditionals that I wish to defend. In the course of this I will introduce a

1. For an impressive survey of the varieties of nonconditional uses of “if”, see Declerck and Reed 2001, chapter 10. I would not want to exclude all of the sorts of conditionals that Declerck and Reed call “rhetorical conditionals” from the purview of my account of conditionals, however. For example, “If you think your bike is bad, you should try mine!” (2001, 5) seems to me to be a perfectly good conditional (what Declerck and Reed call a *case-specifying-P* conditional). What is special about it is just that the point of interest to someone who asserts it may not be so much what he says but what he implies, namely, that his own bike is very bad, or that the addressee should not have such high standards.

number of important concepts such as *context* and *assertibility* and will explain what I mean by them. I hope that the basic conception of conditionals presented in this chapter will seem plausible on its own, but the heart of my defense of it is chapter 3, which is grounded in chapter 2. That defense is supplemented by critical comparisons to other people's theories in chapter 4. Chapters 5 through 7 will spell out the theory in precise, or, as they say, formal, detail. The final two chapters will extend the theory and develop in more detail its metatheory.

1 The Context-Relativity of Conditionals

As philosophers and linguists are well aware, sentences exhibit context-relativities of various kinds. Everyone acknowledges that the semantic value of a sentence like "You are taller than that" depends on the context in which it is uttered, inasmuch as it depends on who and what "you" and "that" refer to. Most philosophers will acknowledge as well that the semantic value of a sentence like "Everyone is present" depends on the context in which it is uttered inasmuch as it depends on the domain of discourse relative to which we interpret the quantifier "everyone". The semantic value of "That's a big one" depends on the standard of size that is pertinent to the context (so that a big mouse does not have to be as big as a big house). When we say "Henrietta's car", that may mean either the car that Henrietta drives and owns or the car that Henrietta merely sits in, depending on the context.

Conditionals exhibit a kind of context-relativity quite different from any of these others. Consider, for example, the following conditional:

If you turn left at the next corner, you will see a blue house at the end of the street.

Fix the time and place of utterance. Fix the people, places, and times referred to. Let it be decided whether or not the addressee will walk to the next corner and turn left and whether or not the addressee will see a blue house at the end of the street. Still, the semantic value of this sentence depends on something else. My claim in this book will be that the semantic value of this sentence depends on the *context* pertinent to the conversation in which it is, or might be, uttered, and that the way in which the context makes a difference is by providing a range of *relevant prospects*. If in each of these prospects in which the addressee turns left at the next corner, the addressee sees a blue house at the end of the street, then the conditional has a positive semantic value; if in some of them the addressee

turns left at the next corner but does not see a blue house at the end of the street, then it does not.

A relevant prospect has to be distinguished from a mere possibility. The above conditional may have a positive semantic value even if it is *possible* for the antecedent to be true while the conclusion is false. If the addressee turns left at the next corner, he might fail to see a blue house because a large moving van blocks his view, or because the house has been painted pink or has burned down, or because a crow swoops down and pecks out his eyes. Despite all these possibilities, this conditional may, in some contexts, have a positive semantic value, for in some situations we may be entitled to ignore some of these possibilities. Whether we are entitled to ignore them is not independent of how the world actually is apart from what we happen to believe. At the very least, if in fact there is a moving van parked in front of the house that blocks the view of it from the corner, then this conditional is just wrong (for whatever kind of wrongness is the object of our semantics). But it is not the case that we are entitled to ignore all those possibilities that will not in fact be realized. What we are entitled to ignore depends on the context.

Subjunctive conditionals differ from indicative conditionals only in that each subjunctive conditional is evaluated with respect to a broader range of prospects than the corresponding indicative conditional would be evaluated with respect to in the same context. This broader range too is a feature of the context. The relation between the range of prospects relative to which the indicative conditional is evaluated and the range relative to which the subjunctive conditional is evaluated may be of various kinds. For instance, there might be a situation in which the opportunity to turn left at the pertinent corner has passed. In that case, the context might not immediately provide any prospects in which the addressee goes to that corner. Nonetheless, the context might also provide a range of less immediately relevant prospects in some of which the addressee turned left at the pertinent corner. If in each of those less immediately relevant prospects in which the addressee turns left, the addressee sees a blue house at the end of the street, then the following subjunctive conditional will be assertible:

If you had turned left at the next corner, then you would have seen a blue house at the end of the street.

In this case, the difference between the range of immediately relevant prospects, relative to which we will evaluate indicative conditionals, and the broader range of more remote prospects, relative to which we will

evaluate subjunctive conditionals, may be that the latter includes prospects that might have been realized if something different had been done in the past but which, in consequence of what the agent actually did, can no longer be realized.

The term “context” could be used in at least two different ways. First, it could be used to describe certain aspects of the actual environment in which a conversation takes place (or might take place): roughly, the arrangement of objects in the vicinity and the mental states of the interlocutors. Second, it could be used to refer to the values of a certain variable in a semantic theory. For example, in evaluating an utterance (or potential utterance) of the sentence “Everyone is present”, we will have reason to consider the people who are relevant to the discourse and whether they are present. Using the term “context” in the first way, we could describe those people and their status as present or absent as part of, or aspects of, the context. But in formulating a semantic theory, we will also have reason to speak of a “domain of discourse”—a set of objects, or perhaps names of objects—as the value of a variable relative to which, according to our semantic theory, quantified sentences are to be evaluated. Using the term “context” in the second way, we could describe this domain of discourse as part of the context. Likewise, given that the evaluation of a conditional depends on a set of relevant prospects, we could use the term “context” to refer to those aspects of the situation that determine which prospects are relevant, or we could use it to refer to the set of relevant prospects itself.

In this book I will use the term “context” in only the second way, to refer to the values of a certain variable posited by a semantic theory. A context will be a certain sort of structure, and a semantic theory for a language will specify, for each sentence of the language, the conditions that a structure of that kind must meet in order for the sentence to have a positive semantic value relative to that structure. For purposes of evaluating conditionals, for example, the content of the context that matters will be the range of prospects that the context includes (or *is*). The semantic value of a sentence relative to a context will have a bearing on conversations because, for each conversation, there will be a particular context that *pertains* to it. So we may identify what is semantically good or bad in a given conversation in terms of the semantic values that sentences have relative to the context that pertains to it.

Certainly there will also be something about the character of the environment in which a conversation takes place that determines which context—which value of a certain variable in our semantic theory—is

the context that pertains to a given conversation. So the other thing we might have called “the context” is those features of the world that determine, for a given conversation, which range of prospects is the range of prospects relative to which we ought to evaluate any conditional that may be uttered in the course of that conversation. But in fact I will not call those things “the context”, but will prefer the terms “situation” and “circumstance” to refer to such context-determining aspects of the environment.

Some people, having vaguely perceived the context-relativity of conditionals, might conclude that conditionals are neither true nor false. The assumption would be that if a conditional is either true or false, then it must be true or false absolutely. More precisely, if the context-relativity cannot be attributed to one of the familiar brands of context-relativity, then conditionals must be neither true nor false. Rather, a conditional might be thought to be merely an expression of the speaker’s state of mind (for example, an expression of the subjective conditional probability that the speaker assigns to the consequent given the antecedent). The conclusion I draw instead is that we ought to acknowledge that conditionals are context-relative and then write that context-relativity tightly into the sinews of our semantics so that the semantic value of every sort of sentence whatsoever is relative to context in the same way.²

According to the prevailing point of view in semantics, a speaker’s aim in making an assertion is to express a proposition, that is, to reveal to the audience that he or she stands in a certain cognitive relation to a certain proposition. A speaker can succeed in revealing his or her relation to a proposition because sentences express propositions, or express propositions in context. Normally, the speaker intends the audience to recognize that the speaker stands in the pertinent cognitive relation to the very same proposition that his or her words express in context. The speaker can expect this intention to be realized insofar as the audience can be counted on to know which proposition the speaker’s sentence expresses in the con-

2. One author who I think is misled by his failure to countenance the context-relativity of conditionals is Jonathan Bennett. Confronting indicative conditionals that, if they had truth values at all, would have different ones in different contexts, he concludes that indicative conditionals do not have truth values at all (2003, 83–94). (I should add, though, that Bennett’s book is a very clear and thorough review of most of the issues that have been treated as central in the literature on conditionals. Reading it while in the midst of writing this book, I found it very useful as a reminder of some of the issues I needed to address.) Similarly, Stephen Schiffer (2003, 184) notes that our evaluation of indicative conditionals with false antecedents varies with “contextually determined interests” and concludes that such conditionals are “indeterminate.”

text in which it is uttered. As for propositions, we can think of a proposition as a classification of the world. So the proposition expressed by “The dodo is extinct” classifies the world as a world in which the dodo is extinct. Thus it is useful to think of a proposition as a set of possible worlds. To make an assertion is to express a proposition, which is true or false depending on whether the actual world belongs to the set of possible worlds with which we identify the proposition. So too, conditional sentences are supposed to express propositions in context and thereby classify the world. According to this view, the main problem in explicating the meaning of conditionals is to identify the sets of possible worlds that conditionals express.

The point of view on semantics that will be adopted in this book is quite different. We will conceive of an assertion as a contribution to a conversation, but we will not suppose that the way in which a speaker makes a contribution is by expressing a proposition. We will suppose that conversations typically serve the achievement of goals that the interlocutors share. Even if two people are just shooting the breeze, they may feign some goal in order to give their conversation some direction. What has to be said in a given conversation depends on the goal of the conversation and the situation in which it takes place. What I am calling the *context* that pertains to a conversation is a product of the goals of the conversation and the environmental circumstances under which the conversation takes place. The context-relative semantic values that we will attribute to sentences relative to a context will be what I call *assertibility* and *deniability* relative to a context. For each sentence of the language, and each context, the sentence will be either assertible, deniable, or neither in the context. A speaker’s objective in participating in a conversation will be to utter only sentences that are actually assertible in the context pertinent to the conversation. By doing so, the speaker helps to identify the content of the context pertinent to the conversation and in that way contributes to the achievement of the goals of the conversation. It is approximately correct to say that an assertible sentence is one that *may* be asserted (but see section 9.3 below). But it is certainly not the case that an assertible sentence *must* be asserted; sometimes a sentence may be assertible though there is no need to assert it, because it already goes without saying. The problem of explicating the semantics of conditionals is to spell out the conditions under which a conditional is assertible in a context.

As part of our writing context-relativity into the sinews of our semantics, even logical validity will be defined in terms of assertibility in a

context. According to the standard conception of logical validity, an argument is valid, roughly, if and only if for every possible world in which the premises are true, the conclusion is true. (Chapter 2 will include a more careful exposition of the standard conception.) According to the conception of logical validity to be advanced in this book, in contrast, an argument is valid if and only if for every context in which the premises are assertible, the conclusion is assertible as well. This conception of logical validity will be integral to the present theory of conditionals for at least two reasons. First, it will be this conception of logical validity that we will employ in deciding which arguments involving conditionals are logically valid. A frequent claim will be that other theorists have misjudged the validity of arguments because they have allowed an illegitimate shift in context between their evaluation of the premises and their evaluation of the conclusion. Second, this theory of logical validity provides a perspective on the linguistic function of conditionals. In light of it, an indicative conditional may be conceived as a kind of rule of inference that is valid only with respect to a restricted range of contexts—a context-relative rule of inference. I will postpone further exposition of this conception of logical validity to the next chapter.

My term for the context-relative semantic property in terms of which I explicate conditionals and define logical validity might have been simply *truth in a context* rather than *assertibility in a context*, but there is perhaps one good reason to prefer the latter term over the former. Very roughly, contexts will comprise what is, in a sense, *objectively relevant* to a conversation. For most of us, our training leads us to balk at the idea that what is true depends on what is relevant, although perhaps we ought to get used to that. Just because the term *assertible* is less deeply entrenched in tradition, we can more readily accept that assertibility is something that depends on what is relevant. Unfortunately, the term does already have a use in philosophy to mean, roughly, the property that a sentence has when a speaker is epistemically justified in using it to make an assertion. I ask the reader to dissociate the term from that use and to allow me to explicate the concept of assertibility in my own way. Even on the present theory, there will be cases in which the content of the context pertinent to a conversation reflects constraints on what the interlocutors can reasonably be held responsible for knowing; however, the determinants of the content of the context will often *not* be epistemic in this way at all. (Let me point out that I spell “assertible” with an “i”. I do not spell it “assertable”, with an “a”, although many other authors do so. For a reason to spell it with an “i”, see Fowler 1965, pp. 2–4.)

2 Conversations in the Abstract

Conversations, I will assume, have goals. In the paradigmatic case, the goals are practical goals, such as obtaining food. The goals may also be, so to speak, intellectual. For instance, the goal may be to resolve some question of principle before moving on to more practical goals on which these matters of principle may have some bearing. Not all actual conversations have real goals of course. A lot of the time, we merely pretend to have some goals in order to have something to say and in that way entertain ourselves. Sometimes conversations are exercises in routine, and while there may be some goal, such as introducing a newcomer, it is the routine and not the goal that dictates the course of the conversation. But in the cases of interest—the ones we must understand first of all in a study of language—conversations have goals.

The goals of a conversation, as I define them, are *collective* goals. They are shared by all participants in a conversation. Even in cases of conflict, there must be a shared goal that keeps the interlocutors talking to one another. If one of the interlocutors is lying, then their “conversation” is not what I am here calling a conversation, for it is not the sort of case to which we must refer in formulating the basic norms of discourse; it is rather a case in which a person exploits those norms for his or her own purposes. Even in genuine conversations it may happen that not all of the participants equally participate in *setting* the goals of a conversation. In some cases, goals may be established by consensus. In other cases, a single interlocutor may have the status of leader and be privileged to set the goals for the whole group. In reality, interlocutors, or subgroups of interlocutors, may have their own goals that conflict with the goals of the group. These may be the source of some of the more subtle features of a conversation, which I will have to pass over in this book.

For each conversation, there will be a *context* that, as I will say, *pertains* to it, or *governs* it. What the interlocutors ought to do in order best to achieve the goal of their conversation will depend on the content of the context that pertains to their conversation. More precisely, what the interlocutors ought to do in order to reliably achieve the goal of their conversation with a minimum of wasted effort depends on the content of the context. What sort of thing contexts intrinsically are will be dealt with later in this chapter. For the moment, I simply assume that the content of the context pertinent to a conversation is determined by two things: the goal of the conversation and the way the world really is. The content of the context pertinent to a conversation depends on these two things

because what the interlocutors ought to do in order best to achieve the goals of their conversation depends on these two things—what the goal is and the way the world is when the conversation takes place.

The context pertinent to a conversation is objective in the sense that it depends on things that may fall outside of the awareness of the several interlocutors. That is so since the context pertinent to their conversation depends on the way the world really is and the interlocutors may be unaware of the pertinent features of the world. And if the goal is set by a leader, some of the interlocutors may even be unclear about the goal. So in contrast to some other popular conceptions of context, contexts are here not to be conceived in terms of interlocutors' beliefs. In particular, the context pertinent to a conversation may not consist of what the interlocutors believe in common or what an interlocutor supposes the interlocutors believe in common. (I will return to the objectivity of contexts in section 9.1 below.)

Since contexts are objective, we have to distinguish between the context that objectively does pertain to a conversation and each interlocutor's *take* on the context pertinent to their conversation. An interlocutor's take on the context is a mental representation of the context pertinent to the conversation in which he or she is engaged, which may or may not match the context that objectively does pertain to the conversation. Here I make no claims about the intrinsic nature of this representation in an interlocutor's mind or brain. (For some discussion of what *not* to say about it, see my 2003b, chapter 2.) Since both contexts and takes on contexts may be modelled, as we will see, as structures built up from sentence-like entities, I do not expect there to be any problem in defining the sense in which an interlocutor's take on the context *matches* the objective context. In particular, we will not need to appeal to a relation of *reference* to explain this.

An interlocutor's take on the context may consist in a representation of a single context that the interlocutor takes to be the context pertinent to the conversation, or it may consist in a representation of a whole range of alternative contexts among which the interlocutor supposes the pertinent context to lie. At the start of a conversation, interlocutors may have no very clear take on the context governing their conversation. They may conceive of a whole range of contexts among which they expect the context pertinent to their conversation to be found. So there may be a whole range of contexts such that an interlocutor considers each context in that range to be potentially the context that really does pertain to the conversation in which he or she is engaged. An interlocutor's take on

the context pertinent to his or her conversation comprises that whole range.

While what interlocutors *ought* to do in order to achieve the goals of their conversation depends on which context objectively does pertain to their conversation, what they will do *in fact* in order to achieve the goals of their conversation depends on what they *take* the context pertinent to their conversation to be. So, in order that what they *do* be what they *ought to do* in order to achieve their goals, it is important that what each of the interlocutors takes to be the context pertinent to their conversation match the context that actually does pertain to their conversation. An assertion can help to bring this about inasmuch as it serves to restrict each interlocutor's take on the context to a narrower range. An assertion can have this narrowing-down effect because for each sentence and each possible context, that sentence will be either *assertible* in that context or not. When an interlocutor uses a sentence to make an assertion, each interlocutor who *accepts* that assertion restricts his or her take on the context to a range of contexts in each of which the sentence used in the assertion is assertible. So if the sentence used in making the assertion really is assertible in the context that objectively pertains to the conversation, then, insofar as the assertion is accepted, each interlocutor's take on the context will better approximate to the context that objectively pertains to the conversation. So speakers will make it their aim to make assertions that will be accepted and to assert only what is assertible in the context that really is pertinent to their conversation and in that way to restrict every other interlocutor's take on the context to a narrower range that still contains the context that objectively pertains to their conversation.

Sentences may fail to be assertible in a context in either of two ways. Sentences may be positively deniable in the context, or they may be simply unassertible without being deniable. The deniable ones are those the assertion of which would be positively misleading. If the interlocutors take any sentences to be assertible that in reality are deniable in the context that pertains to their conversation, then their pursuit of the goals of the conversation will be positively hindered. Other sentences may be both unassertible and undeniable. The pursuit of the goal will be neither aided nor hindered if all of the interlocutors confine their search to contexts in which such sentences are assertible or to contexts in which such sentences are deniable.

To say that a sentence is assertible is not to say that it positively should be asserted. Some sentences that are assertible in the context pertinent to a conversation may *go without saying* in the sense that their assertion

would do nothing to narrow the range of contexts that the interlocutors regard as potentially the context pertinent to their conversation. That is, an assertible sentence will go without saying if it is assertible in every context that belongs to any interlocutor's take on the context for the conversation. Also, to say that a sentence is assertible is not to say that there can be no harm at all in asserting it. As I will explain in section 9.3 below, it can happen, in view of the accidental features of a discursive practice, that an assertible sentence is nonetheless, in a way, *misleading*. So we can say that a speaker's fundamental objective in contributing to a conversation ought to be to assert whatever is assertible in the context pertinent to the conversation, provided it does not go without saying and is not misleading.

On this conception of the nature of conversation, a semantics for a language will do basically two things: First, it will explain the formal structure of contexts for that language. Second, for each sentence of the language, it will identify the conditions under which that sentence is assertible in a context and the conditions under which that sentence is deniable in a context. In particular, a theory of conditionals will take the form of an account of the conditions under which conditionals are assertible in contexts of the sort appropriate for a language containing such conditionals.

3 Primitive Contexts

A definition of contexts will be relative to a language. That is, we will always suppose that what we are defining are the contexts for a given language. For purposes of defining the formal structure of contexts for a given language, we will suppose that the language in question may be augmented by any countable number of *individual terms*. We can think of these individual terms as like names, in that each one will correspond to an individual that is relevant to the conversation, but they will be like demonstrative expressions, such as "this" and "that", in that we can recycle them and use them to refer to different individual objects in different situations.³ By an *atomic sentence* for the language, I mean a sentence

3. What I have said here is a bit of simplification. What if more than denumerably many objects are relevant to the conversation, as might be the case in a conversation about the real numbers? In that case, some individual terms will hold the place of a multitude of objects. That can happen even when only countably many objects are relevant to the conversation. For further discussion, see chapter 2, section 5.

composed of a single simple predicate of the language and an appropriate number of such individual terms. *Literals* are sentences that are either atomic sentences or negations of atomic sentences.

The simplest kind of context for a language will be what I call a *primitive context*. A primitive context is, basically, a set of literals. However, I will assume that a primitive context is formally *consistent* in the sense that it never contains both a sentence and the negation of that same sentence. So, *formally*, a primitive context is a formally consistent set of literals.

In addition to this formal account of primitive contexts, we need a *substantive* account that explains the conditions under which a given formal structure of the kind described is the one that pertains to a given conversation. In the previous section, I have already indicated the kind of answer I will give to this. In general, the context pertinent to a conversation is a function of the goals of the conversation and the way the world is when the conversation takes place. Now that we have seen that the context is itself a linguistic entity, we can add that the conventions of the particular language spoken will have a role to play here as well. In light of the present account of the structure of a primitive context, we can be more specific about the content of the context. The primitive context that actually pertains to the conversation can be approximately defined as the *smallest* formally consistent set of literals such that the interlocutors can reliably be expected to achieve the goal of the conversation if what each of them takes to be the primitive context pertinent to their conversation is that set of literals.⁴

Let us say that a sentence that does not contain conditional connectives (either as the dominant connective or in its components) is a *conditional-free* sentence. (I ignore altogether modal connectives such as “possibly”.) In the previous section, I said that an assertion functions to get the interlocutors to confine their search for the context pertinent to their conversation to contexts relative to which the sentence asserted is assertible. We may now add that the *primary* function of an assertion of a conditional-free sentence is to narrow down the range of *primitive* contexts (i.e., structures satisfying the formal definition) that might be *the* primitive context pertinent to the conversation in which the assertion takes place. That is,

4. Elsewhere, e.g., in my 2003b, chapter 3, I have defined the primitive context pertinent to a conversation as the smallest, consistent set of literals such that every action in *accordance* with that set of literals is a good way of achieving the goal. That definition comes to the same as the present definition on the assumption that each interlocutor definitely will act in accordance with that set of literals that he or she *takes* to be the primitive context pertinent to the conversation.

the primary function of an assertion using a conditional-free sentence is, for each interlocutor, to narrow down the range of primitive contexts that each interlocutor conceives of as containing the primitive context pertinent to the conversation. (I will identify a secondary function in the next section.) Insofar as each interlocutor confines himself or herself to asserting only what is assertible in the context that really does pertain to the conversation, interlocutors who narrow their search to those primitive contexts in which each interlocutor's assertions are assertible will not thereby exclude the context that really is pertinent to their conversation.

Thus, for each primitive context and each conditional-free sentence, that sentence will be either *assertible*, *deniable*, or *neither* in that context. The condition that a context has to meet in order for a given sentence to be assertible in it will depend on the logical form of the sentence. An atomic sentence will be assertible in a primitive context if it is actually a member of the context; it will be deniable if its negation is a member. A negation (of any kind of sentence) will be assertible in a primitive context if the sentence that it is the negation of is deniable in that context; a negation will be deniable if the sentence that it is the negation of is assertible in that primitive context. A disjunction will be assertible in a primitive context if one or the other of the disjuncts is assertible in it; it will be deniable if both disjuncts are deniable. A conjunction will be assertible in a primitive context if both conjuncts are assertible in it; it will be deniable if one or the other of the conjuncts is deniable in it. In order to be assertible in a context, a sentence must satisfy one or the other of the conditions on assertibility; otherwise, it is not assertible in the context; likewise for deniability.

Here is an example on which I will build when I introduce the theory of conditionals. Suppose two people, Ailard and Arno, have it as their goal to erect a wooden fence. They have purchased the posts and rails at the lumber yard and have them in the back of their pickup truck. The post hole auger has to be borrowed from Mr. Frug, who is the only person in the neighborhood who owns one. Mr. Tim used to own the post hole auger, but Mr. Frug traded it to Mr. Tim for a weed wacker. The task immediately at hand is only to collect all the necessary ingredients together in one place in order to erect the first few posts. The primitive context for their conversation might be something like this:

{*a* is lumber. *b* is our truck. *a* is in the back of *b*. *c* is a post hole auger.
Mr. Frug has *c*. Mr. Tim does not have *c*.}

We may think of this context as comprising the essential basic facts pertinent to Ailard and Arno's proximal goal of collecting in one place the objects that they will need in order to begin building the fence. That is, it is the smallest consistent set of literals such that the interlocutors, Ailard and Arno, can reliably be expected to achieve the goal of the conversation, if what each of them takes to be the primitive context pertinent to their conversation is that set of literals. For example, if we remove "*b* is our truck", then if both Ailard and Arno took the resulting set to be the primitive context for their conversation, they might not look in the right place for the lumber. Or if we removed from the set "Mr. Tim does not have *c*", then there might be some temptation on the part of one or both of them to go looking for a post hole auger from Mr. Tim (despite the fact that "Mr. Frug has *c*" would still belong to the set). But if we added "*d* is a weed wacker", then the set would be too big. Though the addition of this sentence might not mislead them, Ailard and Arno could reliably achieve their goals even if they did not take the context to include this sentence.

Each of the sentences in the primitive context pertinent to Ailard and Arno's conversation can be said to be *assertible* relative to the set. For example, "*b* is our truck" is assertible relative to the set. And each of the sentences whose negations belong to the set is *deniable* relative to the set. For example, "Mr. Tim has *c*" is *deniable* relative to the set. "Mr. Frug has *c* or Mr. Tim has *c*" is an example of a disjunction that is assertible in the primitive context, because one of the disjuncts, "Mr. Frug has *c*", is assertible in it. Other sentences, such as "*d* is a weed wacker", are neither assertible nor deniable relative to this primitive context.

Actually, though, we have to draw a distinction between what is "assertible" in the technical sense and what the interlocutors really need to assert. The lower case letters that I have used as individual terms in representing the context may not actually be words of the interlocutors' common language. There may be words in the common language, such as "this" and "that", that we could use instead in specifying the content of a context, but there may not be enough of them, and they may not always be understood when actually spoken. So the language by means of which we theorists specify the content of a context governing a conversation may be supplemented by additional individual terms not part of the actual spoken language, as I have done above. And to say that a sentence formed of one of these additional individual terms is assertible relative to a context is not to say that the speaker should actually assert that very sentence. Rather, what a speaker should assert is only what is, in view of

the individual terms that will be understandable to his or her interlocutors, a *suitable rewrite* of a sentence or sentences that are assertible.

For example, one of the atomic sentences assertible relative to the above context is, I said, “*b* is our truck”. But the letter “*b*”, used in this way, may not actually be a word of the language that Ailard and Arno speak. Nonetheless, a suitably uttered token of the demonstrative “that” might draw Arno’s attention to the truck. For that reason “*That* is our truck” might be a suitable rewrite of “*b* is our truck”. Or instead of asserting “*a* is in the back of *b*”, Ailard might exploit the assertibility of “*a* is lumber” and “*b* is our truck” to produce the following suitable rewrite: “The lumber is in the back of our truck”. This topic, the ways in which the literals in a context determine what can actually be asserted, deserves a careful examination and theoretical codification, but I will not attempt to develop a proper theory of it here. Where we seem to require such a theory, we can avoid it by simply pretending that the letters that we use to specify the context are actually words in the common language.

In another way too, what interlocutors actually need to assert may differ from what, in our technical sense, is “assertible” relative to a context. As I noted in the previous section, some things that are, technically speaking, assertible may *go without saying*. If Ailard and Arno both know where the lumber is, it may not be necessary for Ailard to say anything to Arno about it; that might go without saying. “The lumber is in the back of our truck” is assertible then, but it goes without saying. But if only Ailard knows that it is Mr. Frug who has the auger, then Ailard might find it necessary to assert to Arno, “Mr. Frug has the post hole auger”. Or Ailard may not know exactly who has the auger, and in that case he might assert only, “Either Mr. Frug has the post hole auger, or Mr. Tim has it”.

4 Multicontexts and Conditionals

The assertibility and deniability conditions of indicative conditionals will be formulated as conditions on assertibility and deniability with respect to more complex structures that I will (when I need to distinguish them from primitive contexts) call *multicontexts*. Conditionals will be assertible or deniable in primitive contexts only as a limiting case. Formally, the simplest kind of multicontext is just a set of primitive contexts. Shortly, I will introduce the full range of multicontexts, but for the moment I will concentrate on those that are sets of primitive contexts. When I speak of the

“prospects” that a context presents us with, as I have already been doing, that is my pretheoretic way of referring to those things that, in the precise theory, turn out to be the members of a multicontext.

In addition to characterizing the formal structure of a multicontext, we need to say something about what might make it the case that a conversation is governed by a structure of that kind. There may be situations in which a whole range of primitive contexts is relevant to our conversation. In that case, a conversation might be characterized by a multicontext consisting of a set of primitive contexts. Typically, when it is a multicontext, rather than a primitive context, that pertains to a conversation, the goal of the conversation will be in some way “intellectual” rather than immediately practical. For instance, at an earlier stage in their planning, Ailard and Arno were at the lumber yard and had it as their goal to decide what kind of lumber to buy. They could build their fence out of cedar or pine, and they could choose between two distances between posts. A cedar fence would last longer, but it would cost more. A shorter distance between posts would provide better resistance to storms and strong animals. Under these circumstances, the multicontext for Ailard and Arno’s decision might comprise the following primitive contexts.

{We buy lumber. Lumber is cedar. We buy thirteen posts. We buy nine-foot rails.}

{We buy lumber. Lumber is pine. We buy nineteen posts. We buy six-foot rails.}

{We buy lumber. Lumber is pine. We buy thirteen posts. We buy nine-foot rails.}

The prospect of buying nineteen posts and six-foot rails in cedar is out of the question, we might suppose, since that would make the job too expensive. Since “We buy lumber” belongs to every one of these primitive contexts, we will say that it is assertible in the multicontext as a whole. Since in every one of these primitive contexts either “Lumber is cedar” or “Lumber is pine” is assertible, we will say that the disjunction, “Lumber is cedar or lumber is pine”, is assertible in the multicontext as a whole.

Conditionals do their work in just this sort of situation. In each of the primitive contexts in this multicontext in which “We buy thirteen posts” is assertible, “We buy nine-foot rails” is assertible as well. “We buy nine-foot rails” is not assertible in the collection as a whole, but neither is “We buy thirteen posts”. So for each context in the multicontext, and for the

multicontext itself, if “We buy thirteen posts” is assertible there, then so is “We buy nine-foot rails”. A conditional can be used to express this fact. The conditional “If we buy thirteen posts, then we buy nine-foot rails” is assertible relative to the multicontext pertinent to Ailard and Arno’s conversation for that very reason. Similarly, “If lumber is cedar, then we buy thirteen posts and nine-foot rails” is assertible in this multicontext since in every context in or identical to this multicontext in which “Lumber is cedar” is assertible, so is “We buy thirteen posts and nine-foot rails”. But “If we buy thirteen posts, then lumber is pine” is not assertible in this multicontext, because buying thirteen posts in cedar is also a prospect represented in this set of primitive contexts.

In general, we may say that an indicative conditional “If \mathbf{p} then \mathbf{q} ” is assertible in a context Γ (primitive or multi-) if and only if for every context Δ in or identical to Γ , if \mathbf{p} is assertible in Δ , then \mathbf{q} is assertible in Δ as well. And we may say that an indicative conditional “If \mathbf{p} then \mathbf{q} ” is deniable in Γ if and only if for at least one context Δ in or identical to Γ , \mathbf{p} is assertible in Δ , and \mathbf{q} is deniable in Δ . And if it is simply not the case that the consequent is assertible in every context in or identical to Γ in which the antecedent is assertible, though the consequent is not deniable in any context in or identical to Γ in which the antecedent is assertible, the conditional will fail to be either assertible or deniable in Γ . A special case that we will have occasion to take note of later is that in which a conditional is assertible in a context Γ just because the antecedent is not assertible in any context in or identical to Γ . In such a case, we will say that the conditional is *vacuously* assertible in Γ . (Notice that these statements afford assertibility and deniability conditions to indicative conditionals both in primitive contexts and in multicontexts. In the case of a primitive context Δ , the only context in or identical to Δ will be Δ itself. However, the cases of interest will be those in which a conditional is assertible or not in a multicontext.)

To round out this account of assertibility in a multicontext, I need to state the conditions under which other kinds of sentences may be assertible or deniable in a multicontext. As for literals, and, more generally, sentences containing no conditional connectives, we may say that they are assertible in a multicontext if they are assertible in every member of the multicontext, and that they are deniable in a multicontext if they are deniable in every member of a multicontext. A negated sentence is assertible in a multicontext if the sentence it negates is deniable; and it is deniable if the sentence it negates is assertible. A disjunction is assertible in a multi-

context if either disjunct is assertible in it; it is deniable if both disjuncts are. A conjunction is assertible in a multicontext if both conjuncts are assertible in it; it is deniable if one of the conjuncts is deniable. No sentence is assertible or deniable in a multicontext unless it qualifies as such by one of the sufficient conditions on assertibility and deniability. (Other conditions, for other sorts of sentences, will be laid down later on.) Notice that these conditions offer two routes to assertibility or deniability for sentences that contain no conditional connectives. For example, a disjunction will be assertible if either disjunct is assertible; but it may also be assertible in a multicontext though neither disjunct is assertible in the multicontext if neither disjunct contains conditional connectives and the disjunction is assertible in every member of the multicontext. (A succinct formulation of these assertibility and deniability conditions is given in chapter 5, section 1.)

In the example of Ailard and Arno, the “cause” of the fact that a plurality of primitive contexts pertained to their conversation was their indecision. In other cases, there are other causes. Another possible cause is unpredictability. Suppose that elections to the U.S. House of Representatives will take place in one month. There are a number of close races that could determine whether the majority will be Democratic or Republican. We are in the process of drafting environmental legislation and could write provisions that are in some ways stronger or could write provisions that are in some ways weaker. Moreover, a research team is presently working on a study that is expected to bolster the argument for the stronger provisions, but we are uncertain whether the study will be completed in time for us to use it. Then the primitive contexts that constitute the multicontext for our conversation might be as follows. (Here, just for the sake of ease of illustration, I loosen up a little on the requirement that the members of the primitive context be literals.)

{The Democrats will win a majority. The study will be complete. We propose the strong provisions.}

{The Democrats will win a majority. The study will not be complete. We propose the strong provisions.}

{The Republicans will win a majority. The study will be complete. We propose the strong provisions.}

{The Republicans will win a majority. The study will not be complete. We propose the weak provisions.}

Here the conditionals that are assertible will include: “If the Democrats will win a majority, then we will propose the strong provisions”; “If the study will be complete, then we will propose the strong provisions”; “If we propose the weak provisions, then the Republicans will have won a majority.” A conditional that we can flatly deny in this context is: “If the Democrats win a majority, then the study will be complete.” We can deny this because the second primitive context in our list contains both the antecedent of this conditional and the negation of the consequent of this conditional. (This is the sort of case I alluded to at the end of section 1, when I said, “there will be cases in which the content of the context reflects constraints on what the interlocutors can reasonably be held responsible for knowing.”)

Another cause for the presence of a plurality of primitive contexts in our context may be that our purposes are didactic. Suppose a parent is teaching a child how to care for a potted plant in the child’s bedroom. They are standing over the plant, and the parent presses the child’s finger onto the soil. The soil is moist, and the plant plainly does not need to be watered. Nonetheless, the multicontext for their conversation may comprise the following three primitive contexts:

{The soil is moist. One week has passed since the last watering. The plant does not need to be watered.}

{The soil is moist. One week has not passed since the last watering. The plant does not need to be watered.}

{The soil is dry. One week has passed since the last watering. The plant needs to be watered.}

{The soil is dry. One week has not passed since the last watering. The plant does not need to be watered.}

In this case, the conditionals that are assertible include: “If the soil is moist, then the plant does not need to be watered,” and “If the soil is dry and one week has passed since the last watering, then the plant needs to be watered.” But the following conditional is deniable: “If the soil is dry, then the plant needs to be watered.”

Yet another cause for a plurality of primitive contexts is fairness in discourse. Suppose that an employee for a company has been accused of graft and has hired an attorney to prove his innocence. In fact, the employee is innocent, and his attorney is fully persuaded of this. In a conversation between members of the attorney’s law firm, the sentence, “My

client is innocent” might be assertible. However, in the multicontext pertinent to the court proceedings at one stage, this sentence might fail to be assertible. It might not be assertible just because it remains to be shown that the client is innocent. At one stage in the attorney’s argument in court, the question might be what became of the missing funds. The primitive contexts in the multicontext pertinent to her discourse at that stage might be:

{My client is innocent. My client’s bank account did not grow. My client’s investments did not grow.}

{My client is guilty. My client’s bank account grew. My client’s investments did not grow.}

{My client is guilty. My client’s bank account did not grow. My client’s investments grew.}

{My client is guilty. My client’s bank account grew. My client’s investments grew.}

In that case, the attorney may assert the following conditional: “If my client is guilty, then either his bank account or his investments grew.”

So far we have considered only multicontexts whose members are themselves all primitive contexts. Call such multicontexts *first-level multicontexts*. In addition, we may reckon among the multicontexts contexts whose members may be first-level multicontexts as well as primitive contexts. Call these *second-level multicontexts*. Beyond that, we may countenance contexts whose members may be either primitive contexts, first-level multicontexts, or second-level multicontexts. And so on. In general, a multicontext is any member of any of the levels built up in this way (and in addition, the union of all of these levels is itself a multicontext). Here, for example, is a case in which we might think of a conversation as governed by a second-level multicontext. Suppose three men share a house and have slightly complicated arrangements for sharing household responsibilities. Let us employ the following abbreviations:

A = Al will mop the floor

B = Bill will mow the lawn

C = Chuck will make dinner

In view of the recent history in the household, only certain combinations of the events that these sentences describe are relevant prospects, but each

resident has his own ideas about which combinations those are. In that case, in fairness to the points of view of the three residents, the context for their conversation, Γ , may be a second-level multicontext containing three first-level multicontexts, as follows:

$$\Gamma = \{\Delta_1, \Delta_2, \Delta_3\}.$$

$$\Delta_1 = \{\{A, B, C\}, \{\text{not-}A, B, \text{not-}C\}\}.$$

$$\Delta_2 = \{\{A, B, C\}, \{A, \text{not-}B, C\}\}.$$

$$\Delta_3 = \{\{A, B, C\}, \{A, \text{not-}B, \text{not-}C\}\}.$$

“If **A**, then if **B** then **C**” is assertible in the context comprising these three contexts. To see this, observe that **A** is a member of both of the primitive contexts in Δ_2 ; so **A** is assertible in Δ_2 . Likewise, **A** is assertible in Δ_3 . **A** is not assertible in Δ_1 or in the multicontext comprising Δ_1 , Δ_2 , and Δ_3 . So to test whether “If **A**, then if **B** then **C**” is assertible in this multicontext, we may confine our attention to Δ_2 and Δ_3 and ask whether “If **B** then **C**” is assertible in each of them. Indeed, in each of the contexts in Δ_2 in which **B** is assertible, **C** is assertible as well. So “If **B** then **C**” is assertible in Δ_2 . Similarly, “If **B** then **C**” is assertible in Δ_3 . So “If **A**, then if **B** then **C**” is assertible in Γ . The significance of this example is that it illustrates the conditions under which we may usefully assert a conditional having a consequent that is itself a conditional.

In some of these cases in which a multicontext is pertinent to a conversation, there may be some context in the multicontext that the interlocutors would like to identify as a context more immediately pertinent to their conversation. That is, the goal of a conversation governed by a multicontext may be precisely to identify the member of that multicontext that will govern the interlocutors’ conversation and their actions at the next stage. For example, Ailard and Arno would like to decide what to buy and in that way identify one of the primitive contexts in the multicontext pertinent to their conversation as *the* primitive context pertinent to their subsequent conversation. In such cases, we will say not that there are two contexts that are pertinent to the conversation, both a multicontext and some member of that multicontext. Rather, we will suppose that the multicontext is *the* context pertinent to the conversation at a certain point in time, relative to which assertibility should be evaluated, but that achieving the goal of that conversation might mean moving on to a new conversation for which some member of that multicontext becomes *the*

context, relative to which assertibility in that new conversation ought to be evaluated. (If it seems like a stretch to call this a new “conversation,” then we can associate contexts with conversation *stages* and say that they move on to a new conversation *stage*.)

In the previous section, I said that the primary function of a conditional-free sentence is to narrow down the range of primitive contexts that might be *the* primitive context pertinent to the conversation. We may now add that, when interlocutors have to identify a multicontext pertinent to their conversation out of a range of alternative multicontexts, a *secondary* function of the assertion of a conditional-free sentence may be to narrow down the range of multicontexts that might be *the* multicontext pertinent to their conversation. As will be proved in chapter 5, a conditional-free sentence is assertible in a multicontext if and only if it is assertible in every member of that multicontext. So the assertion of a conditional-free sentence can serve to confine the interlocutors’ choice to multicontexts such that that sentence is assertible in every member of those multicontexts.

Although conditionals too are assertible in primitive contexts, as a limiting case, our interest in them is not that they serve to identify the primitive context pertinent to a conversation. We have an interest in conditionals when, for a reason such as I have illustrated above, the context pertinent to our conversation comprises a whole range of other contexts. We could say that a multicontext is a context that exhibits one or another of the kinds of *indeterminacy* that I have illustrated above. When the context pertinent to a conversation is such a multicontext and the interlocutors have to identify the multicontext pertinent to their conversation out of a range of alternative multicontexts, the assertion of a conditional can serve to narrow the field, namely, to those multicontexts in which the conditional is assertible. That, I suggest, is the function of the assertion of a conditional.

In explaining the function of an assertion in section 2, I explained that an interlocutor’s *take* on the context pertinent to his or her conversation might comprise a whole range of candidates. It is important not to confuse this range of candidates with the multicontext that may be the context pertinent to a conversation. When it is a multicontext, not a primitive context, that pertains to a given conversation, that multicontext may be only one member of a whole range of multicontexts within which the interlocutor expects to find the multicontext pertinent to his or her conversation (and if the interlocutor is mistaken, it may not even be one of

those). The range of contexts in an interlocutor's take on the context pertinent to his or her conversation is an expression of his or her *uncertainty* in the identification of the context pertinent to the conversation. In contrast, the range of contexts (primitive or multi-) in the multicontext that actually does pertain to the conversation is an expression of the *indeterminacies*, as I have just now called them, that characterize the conversation itself.

5 Indicative Conditionals versus Material Conditionals

The *material conditional* is usually defined in terms of truth conditions as a conditional that is true if and only if either the antecedent is false or the consequent is true. Here, where we are formulating our semantics in terms of assertibility in a context instead of truth and allow three semantic values instead of just two, we can say that a sentence of the form “If **p** then **q**” is a material conditional if and only if it can always be replaced by a sentence of the form “Either not-**p** or **q**”. That is, substituting the latter for the former, whether it is free standing or a component of another sentence, does not affect the assertibility conditions of the resulting sentence.

On the present account, indicative conditionals are definitely not material conditionals. As we will see, “Either not-**p** or **q**” logically implies “If **p** then **q**” and, as we will see, there is a special sense in which “If **p** then **q**” implies “Either not-**p** or **q**” (even though the argument is not strictly speaking valid). Nonetheless, indicative conditionals are not material conditionals because they are not everywhere intersubstitutable. To see this, consider the following pair of sentences:

(A1) It is not the case that if **A** then **B**.

(A2) It is not the case that either not-**A** or **B**.

If indicative conditionals are material conditionals, then (A1) and (A2) should be assertible in exactly the same contexts. But here is a simple context in which (A1) is assertible and (A2) is not:

$$\Gamma_1 = \{\{\mathbf{A}, \mathbf{B}\}, \{\mathbf{A}, \text{not-}\mathbf{B}\}\}.$$

(A1) is assertible in Γ_1 because “If **A** then **B**” is deniable there, which is so since Γ_1 contains a context in which **A** is assertible and **B** is deniable. But (A2) is not assertible in Γ_1 because **B** is not deniable in Γ_1 . **B** is not deniable in Γ_1 because it is not deniable in every member of Γ_1 .

Or compare the following two sentences:

(B1) Either if **A** then **C** or if **B** then **C**.

(B2) Either not-**A** or **C** or not-**B** or **C**.

If indicative conditionals are material conditionals, then (B1) and (B2) should be assertible in exactly the same contexts. But here is a simple context in which (B2) is assertible but (B1) is not:

$$\Gamma_2 = \{\{\text{not-}\mathbf{A}, \mathbf{B}, \text{not-}\mathbf{C}\}, \{\mathbf{A}, \text{not-}\mathbf{B}, \text{not-}\mathbf{C}\}\}.$$

(B2) is assertible in Γ_2 because it is assertible in each of the two contexts in Γ_2 . It is assertible in each of the two contexts in Γ_2 because “not-**A**” is assertible in the first-listed one and “not-**B**” is assertible in the second-listed one. But (B1) is not assertible in Γ_2 because neither of its disjuncts is assertible in Γ_2 . “If **A** then **C**” is not assertible in Γ_2 because Γ_2 contains a context in which **A** is assertible and **C** is not, and “If **B** then **C**” is not assertible in Γ_2 because Γ_2 contains a context in which **B** is assertible and **C** is not. Indeed both disjuncts are deniable in Γ_2 . Although (B1) is assertible in each member of Γ , that does not make it assertible in Γ itself. (I will return to this last point in section 8.2, below.)

Although the indicative conditionals that I seek to explicate in this book are not material conditionals, the conditionals that I myself *use* in precise contexts—in particular, whenever I am stating definitions—will be material conditionals. On the present account, indicative conditionals behave, logically, as material conditionals in any situation in which the contexts over which we quantify in declaring an argument to be valid are exclusively primitive contexts that are *maximal* in the sense that for every atomic sentence, either it or its negation is a member. Perhaps the situation in which we should locate this book, its writing and its being read, can be understood as such a situation. But if not, then let us simply stipulate that the conditionals of the language in which this book is written (as distinct from the languages that the book is about) are material conditionals.

6 Subjunctive Conditionals

The category of conditionals that I am calling “subjunctive conditionals” is often characterized as a category of *counterfactual* conditionals. For example, Quirk et al. (1985), who call them “hypothetical conditionals,” write that hypothetical conditionals “convey the speaker’s belief that the

condition will not be fulfilled (for future conditions), is not fulfilled (for present conditionals), or was not fulfilled (for past conditions)” (1091, sec. 15.35). Similar statements can be found in countless textbooks of English grammar.

One may, if one wishes, draw a distinction between conditional *statements* that do not convey disbelief in the antecedent and those that do, but as an account of the semantic distinction between indicative and subjunctive conditional sentences, this one is certainly mistaken. My favorite counterexample involves a union leader who calls up a union member and asks him, “If there were a strike, would you honor the picket line?” In answering “yes,” the respondent is saying, in effect, “Yes, if there were a strike, then I would honor the picket line.” But neither the questioner nor the respondent is presuming that there will be no strike. The whole point of taking the poll is to decide whether to call a strike. Nonetheless, it is often true that a speaker who uses a subjunctive conditional believes that the antecedent is false, and that fact deserves an explanation.

The first step toward defining the assertibility conditions for subjunctive conditionals has to be to supplement our account of contexts. A multicontext is a set of contexts, primitive or multi-. Each of the contexts in the multicontext pertinent to a conversation is in some way relevant to that conversation in light of the features of the situation in which the conversation takes place. In some situations, we may find that we can identify several different multicontexts, each of which comprises the contexts having at least a certain degree of relevance to our conversation. *The* multicontext pertinent to our conversation is the one comprising those contexts that have the highest degree of relevance, but in addition we may countenance another multicontext that comprises also contexts having a lesser degree of relevance. In addition to those, there may be multicontexts comprising contexts that are, to a degree, less relevant still. We can define the assertibility conditions for subjunctive conditionals as relative to such a hierarchy of multicontexts embedding *the* multicontext pertinent to a conversation.

For example, in the story of Ailard and Arno, Ailard and Arno were working under the assumption that they were going to build a three-sided fence, thirty-six feet long on each side, for a total of 108 feet. In anticipation of adding some contexts to the list, let us represent the three contexts immediately relevant to their situation as follows:

$$\Delta_1 = \{\text{We build a 108-foot-long fence. We spend more. Lumber is cedar. We buy thirteen posts. We buy nine-foot rails.}\}.$$

$\Delta_2 = \{\text{We build a 108-foot-long fence. We spend little. Lumber is pine. We buy nineteen posts. We buy six-foot rails.}\}$.

$\Delta_3 = \{\text{We build a 108-foot-long fence. We spend very little. Lumber is pine. We buy thirteen posts. We buy nine-foot rails.}\}$.

Let us say that the multicontext immediately pertinent to their conversation, comprising the contexts immediately relevant to their conversation, is $\Gamma_0 = \{\Delta_1, \Delta_2, \Delta_3\}$. But another prospect for them would be to build a slightly longer fence. In that case, they would have to limit themselves to nine-foot rails and fewer posts per unit length to save money. In this scenario, we might characterize the context immediately pertinent to their conversation as still comprising just the three primitive contexts in Γ_0 , but we might think of Γ_0 as a member of a larger structure of contexts, and in that larger structure of contexts we might find a further multicontext containing all three of those contexts but also the following primitive context:

$\Delta_4 = \{\text{We build a 126-foot-long fence. We spend more. Lumber is pine. We buy fifteen posts. We buy nine-foot rails.}\}$.

In addition, Ailard and Arno might consider spending even more money on their project. In that case, they could buy the cedar lumber for a shorter fence with six-foot rails, or could build the longer fence in pine with six-foot rails:

$\Delta_5 = \{\text{We build a 108-foot-long fence. We spend a lot more. Lumber is cedar. We buy nineteen posts. We buy six-foot rails.}\}$.

$\Delta_6 = \{\text{We build a 126-foot-long fence. We spend a lot more. Lumber is pine. We buy twenty-two posts. We buy six-foot rails.}\}$.

In terms of these primitive contexts, we can define a structure Θ containing not only the context Γ_0 , but also several other multicontexts representing various expanded ranges of contexts:

$\Theta = \{\Gamma_0, \Gamma_1, \Gamma_2\}$.

$\Gamma_0 = \{\Delta_1, \Delta_2, \Delta_3\}$.

$\Gamma_1 = \{\Delta_1, \Delta_2, \Delta_3, \Delta_4\}$ (including the prospect of a long, weak fence in pine).

$\Gamma_2 = \{\Delta_1, \Delta_2, \Delta_3, \Delta_4, \Delta_5, \Delta_6\}$ (including the prospect of a short, strong fence in cedar and the prospect of a long, strong fence in pine).

The assertibility conditions for subjunctive conditionals may be understood in terms of these expanded ranges of contexts. For comparison's sake, consider first the indicative conditional,

- (*) If we build a 126-foot-long fence, then we buy fifteen posts and nine-foot rails.

It is true that for every context Δ in or identical to Γ_0 if the antecedent is assertible in Δ , then so is the consequent; but that is so just because there is no context in Γ_0 where the antecedent is assertible; so (*) is vacuously assertible in Γ_0 . But the structure Θ provides a “least expansion” of Γ_0 that includes a context where “We build a 126-foot-long fence” is assertible, namely, Γ_1 , and for every context in or identical to Γ_1 where “We build a 126-foot-long fence” is assertible, “We buy fifteen posts and nine-foot rails” is assertible as well. So the indicative conditional (*) is assertible in Γ_1 nonvacuously. For just that reason, the following subjunctive conditional is assertible in Γ_0 relative to the structure Θ :

- (**) If we were to build a 126-foot-long fence, then we would buy fifteen posts and nine-foot rails.

Θ also provides a further expansion of Γ_0 containing a context in which “We build a 126-foot-long fence” is assertible, namely Γ_2 , and it is not the case that for every context Δ in or identical to Γ_2 , if “We build a 126-foot-long fence” is assertible in Δ , then “We buy fifteen posts and nine-foot rails” is assertible in Δ (see Δ_6). So (*) is not assertible in Γ_2 . However, Γ_2 is not a *least* expansion of Γ_0 provided by Θ containing a context in which “We build a 126-foot-long fence” is assertible, since Γ_1 is a subset of Γ_2 . And therefore the presence of Γ_2 in Θ is nothing against the assertibility of (**) in Γ_0 relative to Θ .

In general, we may say that a subjunctive conditional “If it were the case that **p**, then it would be the case that **q**” is assertible in a context Γ relative to a structure of contexts Θ if and only if, for every least expansion Γ^* of Γ in Θ such that for some context Δ in Γ^* **p** is assertible in Δ , the indicative conditional “If **p** then **q**” is assertible in Γ^* . A subjunctive conditional “If it were the case that **p**, then it would be the case that **q**” is deniable in a context Γ relative to a structure of contexts Θ if and only if there is a least expansion Γ^* of Γ in Θ such that for some context Δ in Γ^* **p** is assertible in Δ and the indicative conditional, “If **p** then **q**” is deniable in Γ^* . (The “official” formulation of these conditions is in chapter 7, section 2.)

So subjunctive conditionals have a role to play in situations where we may distinguish between a multicontext immediately pertinent to our conversation, comprising a certain range of contexts immediately relevant to the conversation, and other multicontexts that include this entire range but also contain other contexts as well. We can say that the contexts in the multicontext immediately pertinent to a conversation are those that are the *most* relevant, and the contexts that belong only to less pertinent expansions of that multicontext are *less* relevant. But in what *way*, one might ask, are the contexts that belong only to less pertinent expansions less relevant?

One possibility (among others, as we will see) is that the contexts in the multicontext most immediately pertinent to a conversation are just those that fall within agreed-upon boundaries of some kind. For example, it may be that Ailard and Arno have agreed on an upper limit to the amount of money they will spend on the fence, and contexts Δ_4 , Δ_5 , and Δ_6 all fall beyond that limit. In all of the contexts that fall within the agreed-upon limit, “We will build a 108-foot-long fence” is assertible. So that sentence is assertible as well in the context Γ_0 , which is the context comprising all and only the contexts (of those here considered) that fall within the agreed-upon limit. Nonetheless, the subjunctive conditional “If we were to build a 126-foot-long fence, then we would spend more” will be assertible in the context that governs their conversation, Γ_0 , relative to the larger structure of contexts to which it belongs, because among the contexts in Γ_1 , which includes as well the contexts that surpass that limit without surpassing any other, there is a context, Δ_4 , in which “We will build a 126-foot-long fence” is assertible, and in all of the contexts in or identical to Γ_1 in which “We will build a 126-foot-long fence” is assertible, “We will spend more” is assertible. And the subjunctive conditional “If we were to build a 126-foot-long fence, then we would spend just a little” is deniable in Γ_0 , relative to the larger structure of contexts to which it belongs, even though the indicative conditional “If we build a 126-foot-long fence, then we will spend just a little” is vacuously assertible in Γ_0 .

A distinction between the contexts in the multicontext immediately pertinent to a conversation, on the one hand, and the contexts in various pertinent extensions of that multicontext, on the other, might also be drawn in terms of what we can *claim to know*. The contexts in the multicontext Γ immediately pertinent to our conversation might comprise only contexts, each of which is compatible with everything relevant we can *claim to*

know (more precisely: with each relevant knowledge claim that we can express with a conditional-free sentence). But in addition, we might countenance a range of less relevant contexts, forming a multicontext Γ^* , that includes all of the contexts in Γ but contains also contexts that are not compatible with what we can claim to know. In this way we can explain the famous disparity between the following two sentences:

- (1) If Oswald did not kill Kennedy, then someone else did.
- (2) If Oswald had not killed Kennedy, then someone else would have.

Let Γ be the context pertinent to a discussion of those fateful events on November 22, 1963. We collectively claim to know that Kennedy was killed and even that Oswald did it. So it may be that in every context in or identical to Γ in which Oswald did not kill Kennedy (namely, none), Kennedy was killed. So (1) is assertible in Γ . But (2) might nonetheless not be assertible in Γ relative to a structure Θ of expansions of Γ because the least expansion Γ^* of Γ in Θ such that Γ^* contains a context in which it is assertible that Oswald did not kill Kennedy might be a context Γ^* in which the indicative conditional (1) is not assertible. (Having said this, though, I should note that we might be reversing the proper order of explanation if we said that the contexts in the multicontext immediately pertinent to our conversation were those compatible with what we can claim to know rather than that what we can claim to know is only what is compatible with those contexts that belong to the multicontext immediately pertinent to our conversation.)

Further, in some cases we may find that a difference between the contexts in the multicontext Γ pertinent to a conversation relative to a structure Θ and the contexts in each of the expansions of Γ in Θ is that each of the contexts in Γ is compatible with the speaker's beliefs, while in each of the expansions of Γ in Θ , there is at least one context that is incompatible with the speaker's beliefs. (This is not to admit that the content of that context pertinent to a conversation may be *determined* by what a speaker believes.) In that case, if any subjunctive conditional is assertible in Γ relative to Θ though the corresponding indicative conditional is only vacuously assertible in Γ , the antecedent of that conditional will be incompatible with the speaker's beliefs. Consequently, it will be, from the point of view of the speaker, a *counterfactual* conditional.

Suppose, for example, that the interlocutors take for granted both that Kennedy was killed and also that Oswald killed him, and consider the following subjunctive conditional:

- (3) If Oswald had not killed Kennedy, then Kennedy would not have been killed.

Where Γ is the context pertinent to their conversation, and Γ^* is the least expansion in Θ of Γ such that for some Δ in Γ^* “Oswald did not kill Kennedy” is assertible in Δ , the indicative conditional “If Oswald did not kill Kennedy, then Kennedy was not killed” may be assertible in Γ^* , which means that the subjunctive conditional (3) is assertible in Γ relative to the structure Θ . Insofar as we assume, in this case, that Γ comprises all relevant contexts compatible with the interlocutors’ beliefs and we have to reach up to Γ^* in order to find a context in which “Oswald did not kill Kennedy” is assertible, we may infer that the interlocutors do not believe that Oswald did not kill Kennedy.

For a contrast, consider again the case of the union member who says, in effect, “If there were a strike, I would honor the picket line.” In this case, what distinguishes the contexts in the multicontext Γ pertinent to the conversation may not be that they are all the relevant contexts compatible with the interlocutors’ beliefs. Rather, what distinguishes the members of Γ may be that the differences between them do *not* depend on the outcome of the current negotiations between the union and management. In other words, we cannot think of them as describing prospects whose realization depends on the outcomes of the current negotiations. So one such context might include “The Chinese market for our products expands” (inasmuch as this affects the health of the company) and another might include “The Chinese market for our products contracts.” Since the interlocutors may have beliefs about what will happen that go beyond what they believe will happen *regardless* of the negotiations (for instance, they may believe that there will be a strike), Γ does not represent their beliefs. So where Γ^* is the least expansion in Θ of Γ such that “There is a strike” is assertible in some context in Γ^* , we cannot assume that every context that is in Γ^* but not in Γ is contrary to the interlocutors’ beliefs.

7 Strong Assertibility

In some respects, the present theory of conditionals is exceedingly generous in ascribing assertibility to conditionals. Suppose that we are having a conversation about whether we should go on a picnic. In this conversation, the philosopher Socrates is simply not at issue. In particular, the sentence “Socrates is a man” is neither assertible nor deniable in any context

in or identical to the multicontext pertinent to our conversation. So for every context in or identical to the context pertinent to our conversation in which “Socrates is a man” is assertible (namely, none), “We will go on a picnic” is assertible. So, by the present theory of conditionals, the sentence “If Socrates is a man, then we will go on a picnic” is assertible; it will be *vacuously* assertible in the sense defined in section 4 above. There is no danger that from the fact that this is assertible we will go on to infer, without taking heed of the weather, that we will go on a picnic, because the sentence “Socrates is a man” is not assertible in our context. Nonetheless, it is strange that such a sentence in such a context should qualify as assertible. Similarly, even the sentence “If Socrates is a man, then Socrates is not a man” qualifies as assertible in such a context.

One possible attitude to take toward this result would be that while, yes, such conditionals may be vacuously assertible in a context, that is not as bad a result as it might seem. All it shows, one might say, is that assertibility is not all that we expect from an assertion. In addition, one thing we expect from the assertion of a conditional is that the antecedent be relevant to our conversation. In section 2, I pointed out that some assertible sentences may go without saying in the sense that their assertion would do nothing to narrow the range of contexts in which the other interlocutors expect to find the context pertinent to the conversation. So we could say that this conditional about Socrates, though it is assertible, is not one that should be asserted because inevitably it will go without saying.

But there is reason not to just leave it at that. Consider a context Γ such that “We will go on a picnic” is deniable in every context in or identical to Γ , so that “We will not go on a picnic” is assertible in Γ . Suppose moreover, that “Socrates is a man” is not assertible in any context in or identical to Γ so that “If Socrates is a man, then we will go on a picnic” is assertible (vacuously) in Γ as well. So by Modus Tollens we should be able to conclude that “Socrates is not a man” is assertible in Γ as well. But in fact that sentence may fail to be assertible in Γ ; for it may be that “Socrates is a man” is neither assertible nor deniable in Γ . So if we want to preserve the validity of Modus Tollens, we will have to define a special kind of assertibility that a sentence like “If Socrates is a man, then we will go on a picnic” must lack in a context which is such that “We will not go on a picnic” is assertible in it although “Socrates is not a man” is not assertible in it.

Thus we have reason to define a kind of *strong* assertibility such that an indicative conditional “If **p** then **q**” is strongly assertible in a multicontext

Γ if and only if the following holds: (i) for every context Δ in or identical to Γ , if \mathbf{p} is assertible in Δ , then \mathbf{q} is assertible in Δ , and (ii) for every context Δ in Γ , \mathbf{p} is either assertible or deniable in Δ . (Conditional-free sentences will be strongly assertible if and only if they are assertible. Indicative conditionals will be strongly assertible only in multicontexts.) Of course, we are not likely to find ourselves in a context in which the conditional “If Socrates is a man, then we will go on a picnic” is strongly assertible in this sense. But we will find that in any context in which “We will not go on a picnic” is assertible and in which “If Socrates is a man, then we will go on a picnic” is strongly assertible, the conclusion “Socrates is not a man” is assertible as well. For if a conditional is both vacuously assertible in a context Γ and strongly assertible in Γ , then the antecedent must be deniable in every context in Γ . So this instance of Modus Tollens is in a sense valid. (As we will see, there is reason to deny that Modus Tollens is valid even in this sense when the consequent of the conditional is itself a conditional.)

The introduction of the concept of strong assertibility thus serves two purposes. First, it allows us to identify what is wrong with a sentence “If Socrates is a man, then we will go on picnic” in a case where it is assertible but not something we would want to assert: It is not strongly assertible. Second, it allows us to explain what is right about instances of Modus Tollens when they are right: The premises imply the conclusion in the sense that in any context in which the premises are strongly assertible, the conclusion is at least assertible.

As I will explain in detail in chapter 6, the concept of strong assertibility is not just an ad hoc device designed to excuse away some uncomfortable consequences of the theory. In light of the job we expect a conditional to do, and in light of the sorts of choices we have to make in identifying the context pertinent to our conversation, we can show that what we may expect of a counterexample to a form of argument is that the premises be strongly assertible in it (and that the conclusion not be assertible in it).

8 Two Stipulations

8.1 Conditionals with Conditional Antecedents

The project of this book is motivated in part by the assumption that conditionals can be meaningfully embedded under other conditionals and other logical constants. Attempts to deny this (e.g., Jackson 1987,

appendix; Edgington 1995, 282–284) by claiming that sentences in which conditionals appear to be embedded under conditionals or other logical constants are either meaningless or can be translated into some other kind of sentence are not plausible. “If p then if q then r ” does not mean the same as “If p then either not- q or r ” and does not mean quite the same as “If p and q then r ”. “It is not the case that if p then q ” does not mean “ p and not- q ” and does not mean “If p then not- q ”. “Either if p then r or if q then r ” does not mean the same as “If p or q then r ” and does not mean the same as “If p and q then r ”.

But, unfortunately, I must postpone until chapter 8 the consideration of one kind of embedding of conditionals, and that is the embedding of conditionals as the antecedents of conditionals. Until we reach that chapter, we will assume that the antecedent of any conditional is not itself a conditional and, indeed, contains no conditionals as subsentential components. The problem is not that conditionals with conditional antecedents cannot be given assertibility conditions on the present theory. The problem is just that conditionals with conditional antecedents are the source of many exceptions when we try to say what sorts of arguments are valid and are the source of many complications in proofs. I am postponing consideration of conditionals with conditional antecedents until chapter 8 because I am concerned that if I tried to accommodate them from the start, the reader might come away with the impression that the theory was incredibly complicated and fragmented, not clearly perceiving that many of the complications could be traced back to this one source.

So until I address the matter again in chapter 8, wherever I write “If p then q ” or “If it were the case that p then it would be the case that q ” or, as I will do later, $(p > q)$ or $(p \gg q)$, one may assume that p is not itself a conditional and does not contain any conditional components. This will be so both for the conditional schemata that occur in premises of argument forms and for the conditional schemata that occur in conclusions of argument forms.⁵

8.2 The Nonrecursion Stipulation for Conditionals

There is an important difference between the formulation of assertibility and deniability conditions for conditional-free sentences and the formulation of assertibility and deniability conditions for conditionals and sen-

5. To avoid unnecessary typographical fuss, I will sometimes use ordinary quotation marks as if they were selective quotes. If I were not using quotation marks as selective quotes, then I would not write: “If p then q ”. Rather, I would use concatenation marks, as follows: “If” \frown p \frown “then” \frown q .

tences containing conditionals. As I noted in section 4 above, for each kind of conditional-free sentence (atomic, negation, disjunction, conjunction), there will be two conditions that certify its assertibility in a context and two conditions that certify its deniability in a context. For example, a disjunction will be assertible in a context if either disjunct is assertible in the context, but, in addition, a disjunction will be assertible in a multicontext if it is assertible in every member of that multicontext. By contrast, there is only one way for an indicative conditional to be assertible in a context: it is assertible in a context if *and only if* for every context Δ in or identical to Γ if the antecedent is assertible in Δ , then the consequent is assertible in Δ too. Conditionals, and more generally, sentences containing conditionals as components, will *not* qualify as assertible in a context just because they are assertible in every member of the multicontext. I will say that a sentence that is assertible in a context if assertible in every member of the context is *recursively* assertible. The stipulation that sentences containing conditionals not be recursively assertible is the *non-recursion stipulation* for sentences containing conditionals.

If we allowed conditionals to be recursively assertible, then we would obtain unacceptable results. Suppose we allowed sentences containing conditionals to be recursively assertible. We find that there is a sentence and a context such that both that sentence and its negation are assertible in that context. But that is surely an unacceptable result. Consider the following context:

Let $\Gamma = \{\Delta_1, \Delta_2\}$, where

$\Delta_1 = \{\{\text{not-}\mathbf{A}, \mathbf{B}\}, \{\mathbf{A}, \text{not-}\mathbf{B}\}\}$, and

$\Delta_2 = \{\{\text{not-}\mathbf{A}, \text{not-}\mathbf{B}\}, \{\mathbf{A}, \text{not-}\mathbf{B}\}\}$.

Allowing sentences containing conditionals to be recursively assertible, we find that “It is not the case that if \mathbf{A} then \mathbf{B} ” is assertible in Γ , because it is assertible in both Δ_1 and Δ_2 . But “If \mathbf{A} then \mathbf{B} ” is also assertible in Γ just because \mathbf{A} is not assertible in any context in or identical to Γ .

However, we cannot justify the nonrecursion stipulation for conditionals solely on the basis of the unacceptable consequences of allowing conditionals to be recursively assertible. If we could justify the nonrecursion stipulation only in that way, then that stipulation would be merely an ad hoc device for avoiding an undesirable consequence. So it is important also to be able to justify this stipulation in terms of a plausible account of the function of conditionals. Basically, my answer will be that it is the function of conditional-free sentences to characterize what is common to

the members of a multicontext, while it is the function of sentences containing conditionals to characterize relations *between* members of a multicontext. However, I think further development of this answer would seem a bit tedious at this point. So I will postpone further discussion to chapter 5, section 5.1.

There are various other details of the present theory that one might wonder about. Why in giving the assertibility and deniability conditions for conditionals do we say “in or identical to” rather than just “in”? Why do we allow conditionals to be assertible in primitive contexts rather than stipulating that they are assertible only in multicontexts? Why do we allow multicontexts to contain multicontexts rather than just confining our attention to sets of primitive contexts? Each of these questions will be answered explicitly in chapter 5, section 5.

9 Foundations

9.1 The Objectivity Requirement

As here conceived, contexts are objective, or mind-independent. The context for a conversation depends on the goals that the interlocutors adopt for their conversation, and to that extent, the content of the context does depend on the mind-set of the interlocutors. But the content of the context is also a matter of what is objectively relevant to the achievement of those goals given the real character of the environment in which the conversation takes place. For that reason, the participants in a conversation may be quite mistaken about the content of the context pertinent to their conversation.

The conception of contexts as here defined thus contrasts sharply with another popular conception of contexts. Commonly the context for a conversation is conceived as a set of propositions that defines the shared beliefs, or common ground, of the interlocutors.⁶ The rationale for defin-

6. The primary exponent of this conception is Robert Stalnaker; see his 1972 and 1974, reprinted in his 1999, and his 1973. (Actually what Stalnaker usually says is that the context for a speaker is the set of assumptions that that speaker *supposes* are shared.) On Stalnaker’s theory of conditionals, which will be the main subject of chapter 4, conditionals are not to be semantically evaluated as assertible in a context; so he is certainly not guilty of subjectifying conditionals in the way I am warning against in this section. The present account of context contrasts also with that found in David Kaplan’s work (1989), where a context is simply a parameter specifying such things as a world, a time, a speaker, a hearer, and other things that we need to know in order to fix the reference of indexicals, such as “actually”, “now”, “me”, and “you”.

ing contexts in this way is not primarily the theory of conditionals; rather it is the phenomenon of presupposition. Elsewhere I have argued that an objective conception of context is more appropriate even for purposes of explicating presupposition. (Roughly, an interlocutor's presuppositions may be identified with his or her *take* on the objective context. See my 1998 or my 2003b, chapter 5. The common ground conception will be briefly criticized in chapter 8, section 1.2, below.) But in any case, we certainly would not wish to define contexts in this way if we wished to explicate the semantics of conditionals in terms of assertibility in a context, as I have done in this chapter. More generally, for present purposes it would be a mistake to identify the content of a context in terms of the contents of the subjective states of the interlocutors.

To see this, observe first that, for a couple of reasons, we expect a single semantic standard, applicable to conditionals and nonconditionals alike. One reason is just that conditionals combine with nonconditionals to form compound sentences in the same way that nonconditionals combine with nonconditionals to form compound sentences. If **p** and **q** are two nonconditional sentences, then we can form the compound "Either **p** or **q**". Likewise, from the nonconditional **p** and the conditional "If **q** then **r**", we can form the compound "Either **p** or if **q** then **r**". If we evaluate the nonconditional **p** by one standard and the conditional "If **q** then **r**" by some other, then how are we to evaluate the compound "Either **p** or if **q** then **r**"? As I have already noted (at the beginning of section 8), these compounds cannot always be construed as equivalent to sentences that do not embed conditionals. Another reason is that we recognize logical relations between conditionals and nonconditionals. For example, "If **p** then **q**" together with **p** implies **q**. For another example, it is generally accepted that "**p** and not-**q**" implies "It is not the case that if **p** then **q**". If one semantic standard is right for the nonconditional sentences in these arguments and some other is right for the conditional sentences, then how are we to define the logical validity that such arguments are supposed to possess?

In the case of nonconditional sentences, it certainly cannot be the case that the only question of right or wrong is whether the sentence adequately expresses a speaker's state of mind or the shared assumptions of interlocutors. If a speaker utters, "The cat is on the mat," it is not the case that the only question we can reasonably ask is whether this sentence adequately expresses what the speaker has in mind or the assumptions that the speaker shares with his or her interlocutors. We can also ask whether the sentence is objectively right or wrong. Even if what we are

ultimately interested in is only the speaker's state of mind, we have to evaluate the speaker's sentence and not just the speaker's state of mind as a part of the process of determining whether the sentence is an adequate expression of the speaker's state of mind. By what we saw in the previous paragraph, the same must go for conditionals. If a speaker says, "If you turn left at the next corner, then you will see a blue house at the end of the street," then it is not the case that the only question we can reasonably ask is whether this conditional is an adequate expression of the speaker's state of mind or of the assumptions that the speaker shares with his or her interlocutors. For whatever sort of objectivity we can expect in a nonconditional, we can expect the same sort of objectivity in a conditional.

Likewise, if the evaluation of conditionals is inevitably context-relative, then in the same way the evaluation of nonconditional sentences must be context-relative. So both conditionals and nonconditional sentences are both objective and context-relative. Consequently, the claim of context-relativity must not undermine the claim of objectivity. The content of the context relative to which we evaluate sentences of either kind must be objective. That is, the content of the context must not be simply determined by the attitudes of the interlocutors but must depend on how the world really is apart from what they may believe or believe in common.

And yet, to say that the content of the context is not determined by the contents of the interlocutors' attitudes is not to deny that it is determined by the actual environment of the conversation in a somewhat mind-dependent way. Where a single primitive context pertains to a conversation, its content depends on the goals of the interlocutors, and to that extent even a primitive context is mind-dependent. Moreover, we have seen that there may not be just one primitive context pertinent to a conversation; the pertinent context may be a multicontext comprising a range of alternative primitive contexts, or even other multicontexts. We evaluate indicative conditionals relative to such multicontexts. Further, such multicontexts in turn may be embedded in a structure of multicontexts that contains also larger multicontexts that include multicontexts the members of which are in various ways less relevant. We evaluate subjunctive conditionals relative to multicontexts embedded in such structures. As we have seen, the range of the membership of the multicontext pertinent to a conversation may reflect indeterminacies of practical, epistemic, didactic, diplomatic, and no doubt other kinds. Inasmuch as the pertinent range of alternatives reflects an indeterminacy in individual or collective attitudes, we may allow that the content of a context may be in these

ways mind-dependent as well. But the reason why a context qualifies for membership in the multicontext pertinent to a conversation is never just that the interlocutors *regard* it as belonging. Given the goal of the conversation and the pertinent kind of indeterminacy, it is still the real character of the world in which the conversation takes place that determines which range of contexts best serves that goal relative to that indeterminacy.

9.2 Assertibility versus Truth

The present theory of conditionals illustrates a more general thesis to the effect that we ought to substitute the concept of assertibility in a context for the concept of truth in the formulation of the semantics of natural language. The proposal to substitute assertibility in a context for truth raises at least two deep questions. One is: What do you mean by “semantics”? If the project of semantics is just *defined* as the project of formulating the truth conditions for sentences of natural languages, then the proposal to substitute assertibility in a context for truth has to be a mistake. But in fact the project of semantics can be defined more generally as the project of explicating that which speakers of a language know in common about their language that enables them to enter into productive verbal exchanges with one another. That is how I propose to define it. (Granted, there is some vagueness in this definition in that we may wonder whether some given bit of knowledge is “about” a language.)

The claim that assertibility in a context is the property of interest in semantics is plausible just insofar as basic *norms of discourse* can be defined in terms of it. That is so, since what speakers need to know about their language in order to enter into productive verbal exchanges is what the norms of discourse are that pertain to their language. In the next two chapters we will find that the concept of assertibility in a context can be put to use in formulating norms of discourse inasmuch as a plausible definition of logical validity can be formulated in terms of it. Beyond this, the concept of assertibility in a context may play a role in the formulation of the norms of discourse inasmuch as we may plausibly lay it down, as a minimum condition on assertion, that a speaker ought to assert only what is assertible in the context pertinent to the conversation that he or she is engaged in. This necessary condition will not also be a sufficient condition, however. As I explained in section 2 above, some sentences may be assertible in a context and yet *go without saying*. Also, as I will explain in the next subsection, assertions can be in some ways misleading even if the sentence used is assertible in the context. So what we can say is that it is a basic norm of discourse that an interlocutor ought to assert whatever is

assertible provided it does not go without saying and the assertion is not otherwise misleading.

The other deep question is: What is truth? Part of this question, when it is a question of comparing truth to assertibility in a context, would be how we might extend the theory of assertibility in a context to a language containing a truth predicate. What are the assertibility conditions for sentences of the form, “*s* is true”? (For my answer to that, see my 2003b, chapter 9, and my 2005.) But there are other, deeper aspects to the question as well. If we do not need the concept of truth for purposes of formulating the semantic properties of natural language sentences, then do we still have any use for it at all? In what sense, if any, can a sentence be true even if it is not assertible in the context in which it is uttered? If sentences can be true even when they are not assertible in a context, why can we not formulate the semantic properties of natural language sentences in terms of such truth rather than in terms of assertibility in a context? These are indeed difficult questions, calling for a reexamination of both the history of philosophy and ordinary discourse about language. I am simply not prepared to take up these difficult questions here.

In characterizing the present theory of conditionals, it is crucial not to lose sight of the fact that the semantic value of other types of sentences as well is to be assertibility in a context rather than truth. For just this reason it would be a mistake to say that, on my theory, conditionals are a kind of *strict* material conditional. In possible worlds semantics, “If *p* then *q*” is a *strict material conditional* if and only if: that sentence is true at a world *w* if and only if at every world *w'* *accessible* from *w*, if *p* is true at *w'* then *q* is true at *w'*. The present theory of indicative conditionals might seem to differ from this only in allowing that worlds may be partial, that is, in not assigning truth or falsehood to every sentence, and in allowing that the domain of accessible worlds may vary from context to context. (Such a theory was developed in a precise way in Tichý 1984.)

What shows that this is a misconstrual of my proposal is that just as conditionals are to be evaluated as assertible relative to multicontexts, which are themselves sets of contexts, so too every other sort of sentence will be evaluated as assertible or not relative to multicontexts. For example, an atomic sentence will be assertible or not in a multicontext Γ if and only if it is assertible in every member of Γ . A disjunction will be assertible in a context Γ if and only if *either* (a) one or the other of the disjuncts is assertible in Γ , *or* (in case the disjunction is conditional-free) (b) for every context Δ in Γ , the disjunction is assertible in Δ (which might be because for every context Δ in Γ , one or the other of the disjuncts is asserti-

ble in Δ). As a result, we will find that the logic of conditionals is very different from the logic of strict conditionals that we would obtain in the framework of possible worlds semantics. For example, arguments having a premise of the form, “Either not- p or q ”, and a conclusion of the form, “If p then q ”, will be logically valid on my theory, although they are not valid where conditionals are construed as strict material conditionals.

Still, one might ask: Could we not call primitive contexts “partial possible worlds” and describe the contexts in a multicontext as “accessible partial possible worlds”? Not really. If, as here, the semantics is not formulated in terms of relations that are naturally construed as relations between possible worlds, such as *similarity*, and we do not interpret sentences as expressing meanings of a kind that we can identify with *sets* of possible worlds, then surely there is no point in retaining the “possible worlds” terminology. Further, in order to reformulate the present semantics with partial possible worlds in place of primitive contexts, we would have to make sense not only of the idea that conversations may be governed by sets of partial possible worlds (corresponding to multicontexts of the first level) but also of the idea that conversations may be governed by sets of *sets* of partial possible worlds and by sets of *sets of sets* of partial possible worlds, and so on (corresponding to multicontexts of higher levels); but I do not know what that would mean.

9.3 Semantics versus Pragmatics

The prevailing point of view in semantics draws a distinction between semantics and pragmatics. Roughly, what belongs to semantics is everything on which the truth value of a sentence depends (relative to whatever parameters, such as possible world, our semantics takes truth to be relative to). Thus, everything belongs to semantics in terms of which we give a recursive definition of truth in a model and in terms of which we define properties such as the logical validity of arguments. What belongs to pragmatics is our explanations of the ways in which speakers can exploit the truth conditions of sentences, or the capacity of sentences to express propositions in context, to *mean* things by what they say.

Different theorists may draw the line between semantics and pragmatics in different places. Some theorists may bury their heads in the sand and insist that context-relativity has no bearing on semantics at all. What varies from context to context is never the proposition that a sentence expresses but only what a speaker means by it. Others may allow that context-relativity belongs to semantics just to the extent that there is some simple rule that tells us what the denotation of an expression is in

each context. So, for instance, the fact (if it were one) that “now” always refers to the time in which it is uttered, might qualify as a semantic fact. (This seems to be the position of Bach 1994.) Such theorists might go so far as to deny that a language containing demonstratives, such as “this” and “that”, can really have a semantics. Such a theorist might maintain that, strictly speaking, semantics pertains only to the language of thought, and that sentences in the language of thought are all “eternal” sentences entirely devoid of such terms. Others might be much more liberal in their definition of semantics. Not only the reference of demonstratives but also the domains of quantification might fall within the scope of semantics (see Stanley and Szabó 2000).

One problem that creates confusion over where to draw the line is the problem of how to explain the reference of demonstrative expressions. On the one hand, we would like to say that hearers use their understanding of the semantic properties of a speaker’s words to infer the content of the speaker’s underlying thought. The assumption is that there is some proposition that the speaker’s words express in context, quite apart from what the speaker happens to be thinking, and that normally one can assume that the propositional content of the thought that motivates the speaker to speak those words is that very same proposition. On this view, semantics alone ought to assign a proposition to a sentence in context. But then in explaining the reference of demonstrative expressions, such as “this”, many theorists are tempted to invoke the speaker’s state of mind. Roughly, a bare “this” must refer to whatever the speaker intended it to refer to. So then if we want to keep semantics and pragmatics separate, we cannot give to semantics alone the job of assigning a proposition to each sentence in context.

The present work is, in effect, an attempt to broaden the compass of semantics to include much of what others might have chalked up to pragmatics. In particular, the context-relativity of conditionals, which others might have acknowledged but chalked up to pragmatics, now becomes part of semantics. It becomes part of semantics inasmuch as we define logical validity in terms of assertibility in a context, and the property in terms of which we define logical validity certainly belongs to semantics. (Elsewhere I have argued that for the same reason the domain of discourse on which the interpretation of quantified sentences depends belongs to semantics as well. See my 1997 or my 2003b, chapter 7. That argument can be transposed into an argument showing that even the reference of demonstratives belongs to semantics.)

This is not to say that everything others have thought of as pragmatics now gets absorbed into semantics. Here, where semantics is supposed to concern the conditions under which sentences are assertible in a context, rather than the propositions that sentences express in context, pragmatics will not be defined as dealing with the ways in which speakers use words to express propositions other than those their words express in context. Nonetheless, there will be various things to say about the use of language beyond what immediately follows from an account of contexts, assertibility in a context, and the norms of discourse definable in those terms. In particular, there may be ways in which the assertion of a sentence can be helpful that go beyond its assertibility in its context, and there may be ways in which an assertion of a sentence can be misleading even if the sentence is not, in our sense, deniable.

One way in which this can happen is simply that we may draw inferences from our knowledge of what in the past has tended to be the case when speakers spoke in certain ways. For example, if a person says “Some of the data are in,” we may infer that not all of the data are in, and that may be useful information to us. The conclusion that not all of the data are in certainly does not follow by logic from the premise that some of the data are in. Nonetheless, we may know that in many circumstances people do not use a logically weaker sentence when they might just as well have used a logically stronger sentence. So from the fact that we are in one of those circumstances and the speaker said only “Some of the data are in” and not “All of the data are in,” we may infer that not all of the data are in. So this is an example of an assertion that may be helpful in ways that go beyond our learning that the sentence employed is assertible in its context. (In offering this account, I deliberately do not employ Grice’s [1989] theory of conversational implicature. For criticism of that theory, see my 2001 or my 2003b, chapter 6.)

It can also happen that an assertion is misleading in a way that does not yet make the sentence used deniable or even unassertible. For example, the quantification “Everyone who went to the picnic got stuck in the rain” may be assertible in a context just because no one at all went to the picnic. If it is assertible in context that no one went to the picnic, then “Everyone who went to the picnic got stuck in the rain” will be *vacuously* assertible in that context. However, if “Everyone who went to the picnic got stuck in the rain” is only vacuously assertible, it would be misleading to say it without adding “because no one went.” Normally we do not expect people to make assertions that they take to be only vacuously

assertible. If someone takes it to be only vacuously assertible that everyone who went to the picnic got stuck in the rain, then we would expect her to make instead the logically stronger assertion that no one went to the picnic. We know that usually people make the strongest relevant claim they can make unless they have special reasons not to do so (which reasons are easy to come by). So if someone asserted, “Everyone who went to the picnic got stuck in the rain,” we might assume that the speaker did not take it to be assertible that no one went to the picnic and consequently, not failing to have an opinion on the matter, took it to be assertible that someone went to the picnic. Insofar as we are inclined to accept this aspect of the speaker’s apparent take on the context, we may ourselves be led to take it to be assertible that someone went to the picnic. “Everyone who went to the picnic got stuck in the rain” may be assertible in that asserting it helpfully rules out contexts in which it is assertible that someone went to the picnic without getting stuck in rain; and yet asserting it may be misleading in that it leads us, in the manner I have described, to take something else to be assertible that may in fact be deniable, namely, that someone went to the picnic.

Another such case will be of special interest in what follows, namely, the case of a vacuously assertible conditional. In section 4 above, I said that a conditional is *vacuously* assertible in a context if and only if the antecedent is not assertible in any context in or identical to that context, and then in section 7, in the discussion of strong assertibility, I considered a special problem posed by the possibility of such vacuously assertible conditionals. A vacuously assertible conditional may be misleading without being deniable just because we tend to expect that people will not assert conditionals that are assertible only vacuously (although we may recognize exceptions). There may be various good reasons not to assert a conditional that one takes to be merely vacuously assertible; it might go without saying, or one may be in a position to deny the antecedent instead. But in any case, we know that people tend not to do it. So if someone asserts a conditional that is only vacuously assertible in what we take to be the context, we may get the impression that we need to revise our conception of the context to include a prospect in which the antecedent is assertible. A conditional may fail to be deniable, in that taking it to be assertible does not all by itself undermine our pursuit of the goals of the conversation; and yet asserting it may be misleading in that it leads us to misconstrue the context as containing a prospect that it does not contain.

10 My Argumentative Strategy

The argumentative strategy of this book is not to argue from conceptual foundations or to make an inference to the best explanation. There are few plain data on which a theory of conditionals can rest. So there are not enough plain data to allow us to try to show that one and only one theory of conditionals can explain that data. While a conditional, as uttered in given conversation in a given situation, might appear true or acceptable, we might be persuaded that it is not so given a persuasive explanation of why it appears so. While an argument containing conditionals might at first appear invalid, we might be persuaded that it is valid given a persuasive account of the appearance of invalidity.

The method has to be, rather, one of consilience. The theory of conditionals must address several questions. I will try to show that my answers to those questions are at least plausible. If the answers to the several component questions all support one another, then we may be fairly confident that we have a correct theory. In this first chapter, I have presented an account of the assertibility conditions of conditionals that may be plausible on its own apart from the other questions we might ask about conditionals. In the next chapter, I will present a theory of logical validity that may be plausible apart from the treatment of conditionals that it makes possible. In chapter 3, I will use this theory of conditionals and the theory of logical validity to be presented in chapter 2 to draw what I expect to be an independently plausible line between the valid arguments involving conditionals and the invalid ones.

In this way, by means of partially independent, but mutually supportive parts—the account of assertibility conditions for conditionals, the theory of validity, and the logic of conditionals—I hope to build a castle in the air that actually floats.