

1.3. Beyond saying: pragmatics

1.3.0. Overview

Our study of logic will be limited to deductive logic; and, even within those bounds, we will consider only the logical forms that are part of first-order logic. These limits imply some others that deserve consideration in their own right: although our study of deductive logic can be seen as the study of meaning, we will not study all aspects of meaning.

1.3.1. A model of language

One simple picture of language sees it as a device for conveying information by way of the proposition expressed by sentences.

1.3.2. Some complications

This simple picture of language is too simple in many respects, but four are especially important for our purposes. Each corresponds to a further way of conveying information.

1.3.3. Speech acts

Questions and commands do not appear to convey propositions, and even declarative sentences may play roles other than assertion.

1.3.4. Implicature

Communication often exploits the assumption that what a speaker says is not only true but satisfies certain other requirements.

1.3.5. Indexicality

When a sentence conveys a proposition, the proposition that is conveyed will usually depend on the context in which the sentence is used, and sentences are sometimes designed to convey information about his context.

1.3.6. Vagueness

The range of application of many terms will have fuzzy boundaries even in a given context, and sentences that apply them to things falling in this gray area may have no determinate truth value.

1.3.7. Presupposition

Another way of conveying information rests on the preconditions for a sentence to have a truth value at all.

1.3.1. A model of language

The idea of truth conditions or of a proposition suggests a simple picture of the way language works. According to this picture, each sentence has truth conditions that are determined by the semantic rules of the language. These truth conditions settle the truth value of the sentence in each possible world, something that is encapsulated in a proposition. The proposition expressed by a sentence is its meaning. The meaning of an expression smaller than a sentence is to be found in the contribution this expression makes to the propositions expressed by sentences containing it. From this point of view, the function of language is to convey propositions.

Just as the information content of a sentence is to be found by considering the range of possible worlds it rules out, the information that a person possesses is to be found by considering the possible worlds that he or she is able to rule out. The more you can rule out, the more information you have; and the kind of information you have is determined by the particular worlds you can rule out. This means that the sum total of your knowledge can be thought of as a proposition.

Anyone's aim in acquiring information could be described as an attempt to distinguish the actual state of the world among the various alternative possibilities—in short, to locate the actual world within the space of all possible worlds. The proposition representing your knowledge goes some distance towards ruling out some possibilities. But it will leave many open, and the actual world could be any of those open possibilities. If someone conveys a proposition to you and you accept it, you are able to rule out a whole region of logical space, a region that can be added to the region ruled out by your existing knowledge. And, in general, this will reduce your uncertainty about the location in logical space of the actual world.

You can generate information to give to others by delimiting a region within the total area you know to be ruled out. Ideally, perhaps, you would simply convey the whole of what you know; but language limits your ability to do this since only a limited range of propositions are expressed by reasonably short sentences. To convey information, you select a sentence that is entailed by what you know and assert it, thereby conveying the proposition this sentence expresses.

This process is illustrated in the following artificial example of sharing information.




Fig. 1.3.1-1. An animation of a conversation in which information is shared. The button > will play the full conversation while the buttons φ, ψ, χ, and θ will each play one of its four stages. The buttons |< and >| move to the initial and final state, respectively.

Initially, the person on the left is able to rule out regions at the left and right of logical space as possibilities for the actual world while the person on the right is able to rule out regions at the top and bottom. The animation then shows a conversation in which each party in turn notices the truth of the one the sentences φ, ψ, χ, and θ and asserts it. The other person accepts this assertion as true and adds its content to the region ruled out by his or her beliefs. At the end of the conversation, the two people share the ability to rule out a region around the boundary of logical space though they still differ in the shape of the region left open in the middle.

In this conversation, each party is depicted as accepting what the other says as true and adding it to his or her own beliefs. The person accepting the assertion could be said to modify his or her beliefs in a way that makes it something he or she might assert. This is an example of a process that the philosopher David Lewis labeled *accommodation*. In this case of accommodation, one's beliefs are altered to accommodate an assertion someone else has made.

Of course, we do not always accept what others say—i.e., we do not always

alter our beliefs to accommodate their assertions—for we may doubt that they are sincere or that they know what they are talking about. But this cannot be the ordinary case. Words can acquire and maintain a conventional meaning only if people usually mean what they say. And the act of asserting a sentence could not have the significance it does unless people were usually willing to accept assertions as well-founded. A critical attitude is important; but, at least practically, it must be the exception. Even when we are critical and ask for the grounds of someone's assertion, our request can be met only if we are at some point willing to accept assertions providing grounds as well-founded; and, when we are willing to do so, this will rarely be because there is no room for further doubt. In short, while we do not always accommodate what others say, accommodation is central to the aspects of language this model captures. We will also see that other forms of accommodation are essential to a number of the aspects of language that are not captured by the model of communication we have been considering.

There is one simplification in the picture above that is not an essential feature of the model depicted but is worth mentioning because it concerns an important use of entailment. Entailment appears in the picture in one way by setting bounds on the range of sentences that you can sincerely assert: if what you assert is to be something you believe, it must be entailed by your beliefs. But entailment also plays a role in your acceptance of what is asserted to you. Even when you do not doubt what has been asserted, you often add only some of its content to your beliefs. While, ideally, you might like to add the full content of what you hear to your beliefs, your ability to store information is limited, and what you do store is determined by your interests. And, if what you store is to be really part of what was asserted it must be implied by that assertion. That is, a fuller picture of the way a proposition is conveyed is the following:

$$\Gamma \models \varphi \models \psi$$

speaker's asserted proposition
beliefs sentence accepted

The first entailment turnstile marks one aspect of the process of determining what to assert (“invention” in the terminology of traditional rhetoric) while the second marks one aspect of the process of interpretation.

Glen Helman 03 Aug 2010

1.3.2. Some complications

Probably no one ever believed that the simplified model of language we have been considering was entirely accurate. But it, or something like it, was until recent decades the working model most logicians used for thinking about the function of language. Around the middle of the 20th century, philosophers became interested in a number of features of language that suggest this picture is inadequate; and these features have been incorporated into a number of richer models of language. The norms of deductive logic that we will study do not rest on the richer structure of these new models, so we will not consider them in detail. But some of the further features of language that they attempt to capture are intertwined with those we will study, so we need to take some time now to disentangle ourselves from a few of these features once and for all and to lay the groundwork for disentangling ourselves from others at later points in the course.

The complicating phenomena that we need to consider have come to be studied under the rubric of *pragmatics*. This term was originally introduced (by Charles Morris) as an alternative to *semantics* in order to distinguish issues concerning the relation between language and its users from the issues concerning the relation between language and what is spoken of. The use of the term *pragmatics* is no longer closely tied to this definition, and I know of no definition that really captures the way it is now used. Probably the best way to understand current usage is to consider some commonly agreed examples of pragmatic phenomena. The following ones are the most important for our purposes.

1) *Sentences are not always used to express propositions.* When a sentence is used to express a proposition, the question of its truth value is a significant one. But not all sentences have truth values or raise questions of truth value. And even when a sentence does have a truth value, its truth value may not be its most important feature. There are many ways of using sentences, many *speech acts*, besides assertion, and the way a sentence is used is one aspect of its meaning. The term *force* is often used for this aspect of meaning.

2) *The information we derive from the use of sentences is not limited to what follows from accommodating them as true.* Assertions can be expected to have properties other than truth, and there can be forms of accommodation associated with these other properties. In particular, the assumption that an assertion has a given property can be the basis for deriving information from the assertion. This produces the phenomena of *implicature*, in which a sentence suggests more than it says. Even when everything a sentence literally says is

true, an additional false suggestion can make it misleading.

These two complications suggest that propositions are not quite as central to the use of language as the simple model suggests: sentences do not serve merely to convey the propositions they express. Several further complications concern the relation between language itself and propositions: saying simply that sentences express propositions is at best a rough approximation to their meaning.

3) *The proposition expressed by a sentence (and thus its truth value) may vary with the context in which it is used.* For example, there is no way to judge the truth value of a sentence like **I put that here yesterday** when it is taken out of context. This dependence on context is due to various phenomena known collectively as *indexicality* or *deixis*. Both terms are etymologically related to terms for pointing, and the functions of words **this** and **that** are paradigm examples. The term *character* has been used for the way the proposition expressed depends on the context.

4) *Even with regard to a given context, a sentence may not have a definite truth value.* The meaning of *vague* terms like **small** and **hot** will vary with the context; and even in a given context there will be no sharp delineation of the cases where they apply truly. We can continue to speak of the character of a sentence containing such terms but only if we allow the proposition expressed to be depend on factors that are not fully determined by actual contexts of use.

5) *Sentences may have truth values in some possible worlds and not in others.* There can be preconditions for a sentence to have a truth value at all. Anything implied by these preconditions counts as a (*semantic*) *presupposition* of the sentence, and it constitutes another way in which information can be derived from it.

The force, implicatures, character, and presuppositions of a sentence are parts of its meaning in the fullest sense of the term. We will consider each at least briefly to distinguish it from the narrower sense of meaning that will be our focus. It is easy to disentangle our topic from some of these phenomena but others require more detailed consideration, and some forms of entanglement are more likely to trip us up than others. As a result we will consider some of these sorts of meaning only to dismiss them quickly, and we will set others aside without completing disentangling ourselves from them. Implicature is the only one of these aspects of meaning that we will need to pay much attention to in later parts of the course.

1.3.3. Speech acts

Although we have been speaking of sentences as if they all had truth values, there are some sentences that not only do not have truth values but cannot have them. It would be crazy to respond to a question like **What time is it?** by saying **True enough** or **You're wrong!** And these responses would be equally out of place in the case of an imperative sentence like **Please shut the door.**

Questions and imperatives are clear cases of sentences where truth values are irrelevant. But truth values may be beside the point in the case of some declarative sentences, too. Saying **True enough** or **You're wrong!** would be out of place in response to a sentence like **I promise to be here tomorrow** or **I apologize for what I said**, but the reasons they would be out of place are different here than in the case of questions and commands. The verbs **promise** and **apologize** can be used to describe certain sorts of actions that can be performed in using language; that is, they express *speech acts*. And, when they are used in the first person present tense (as in the sentences above) by the right person under the right circumstances, they can be used to perform the sort of actions they describe. That is, by saying **I apologize for what I said**, I can do something that can be described truly by the sentence **He apologized for what he said**; that is, given the right circumstances, I apologize simply by saying I do. Verbs that may be used in this way to perform the actions they describe were labeled *performative* by J. L. Austin, the philosopher who did the most to call attention to the variety of speech acts. When I use a performative verb correctly, what I say is true; but the fact that it is true is not very interesting because my saying it is what made it true.

Austin estimated that the performative verbs in English number “on the order of the third power of 10.” If this estimate is accurate, there are thousands of kinds of speech act besides assertion and thousands of varieties of force beyond the sort of force we will focus on. Of course, much of this vocabulary marks only subtle differences of force between speech acts, but the fact that we have vocabulary for making such subtle distinctions indicates how important it is to us to know the specific force of an utterance. Moreover, we need not use performative verbs to perform the acts that these verbs describe. I can apologize without saying **I apologize** and I can make a promise without saying **I promise**. So we can expect that, even when we use declarative sentences, many, and perhaps most, of things we say are not simply assertions. The statement **I will be there** might be a simple assertion predicting the speaker's future location, but it will often (perhaps most often) be a promise.

In spite of this, we will not consider speech acts other than assertion, and

our interest in assertion itself will be limited to one aspect of its force: the expression of a proposition. Although this will cut us off from much of the richness of language, it will not cut us off from much that is central to deductive reasoning. Of course, there is a sense in which conclusions can be drawn from apologies and promises, but such inferences will tend to be matched by conclusions drawn from ordinary assertions using performative verbs to describe apologies and promises (rather than make them). Moreover, many accounts of speech acts generally treat propositions as important components of their meaning, and this gives the study of assertions a central place in the study of all speech acts.

Glen Helman 03 Aug 2010

1.3.4. Implicature

As we have been using the term **imply**, a sentence implies anything whose content is included in the proposition it expresses. Thus we can say that the sentence **My class was taught this morning** implies **A class was taught**. The philosopher H. Paul Grice employed the term **implicates** to capture a different idea that is sometimes expressed by the ordinary use of the term **implies**.

It is not uncommon for information to be suggested by a sentence even though it is not entailed and thus is not part of what the sentence literally says. For example, my assertion of the sentence **My class was taught this morning** would, in most contexts, suggest that I did not teach the class myself. However, this is not part of what I said since my statement would be perfectly true if I taught the class, so **My class was taught this morning** implicates **I did not teach my class this morning** but does not imply it.

The contrasting vocabulary **say** and **suggest** was used in passing in the previous paragraph, and it is a convenient way of expressing the difference between implications and implicatures. Still, it makes a difference how the term **suggest** is understood. In particular, it is not intended in this use of it to convey the idea of subjective association. What a sentence implicates can be as much the product of rules of language as what it implies. The difference between the two lies in the fact that the rules leading to implicature are not (or are not only) rules assigning truth conditions.

To see what sort of rules they might be, let us consider an extension of our simple model of language use that incorporates implicature; in its outlines, it is due to Grice. To account for implicature, we extend the scope of accommodation to include not only the truth of assertions but also certain other features assertions ought to have. The maxim *Speak the truth!* is no doubt the key rule governing assertions, but other maxims, such as *Be informative!* and *Be relevant!*, also play a role. Someone who assumed I was obeying all maxims of this sort when I said, “My class was taught this morning,” might reason as follows:

Although Helman’s assertion **My class was taught this morning** would have been perfectly true if he had taught his class, it would have been a strange thing to say in that case because the proposition expressed by **I taught my class this morning** would have contained more relevant information. So I can best accommodate his use of language if I assume he did not teach the class.

Let us say that an assertion is *appropriate* when it is in accord with all

maxims governing language use and otherwise say that it is *inappropriate*. An assertion could be inappropriate even though true, so we go further when we assume it is appropriate. At that is something we usually do; that is, we usually accommodate our beliefs about the world to the assumption that the assertions others make are not only true but appropriate for the context in which they are made.

These ideas can be used to state contrasting definitions for implication and implicature. First let’s restate our definition of implication in a way that will make the comparison easier:

ϕ implies ψ (in a given context) if and only if ϕ cannot be true (in that context) when ψ is false (in that context).

To define implicature, we follow the same pattern using the concept of appropriateness instead of truth.

ϕ implicates ψ (in a given context) if and only if ϕ cannot be appropriate (in that context) when ψ is false (in that context).

That is, while implications are conditions necessary for truth, implicatures are conditions necessary for appropriateness. (Notice that the term **implicature** is used here both for the things a sentence implicates and for the relation between a sentence and what it implicates. Our use of the term **implication** follows the same pattern.)

One aspect of the relation between implication and implicature depends on whether we understand truth itself to be one of the requirements of appropriateness. It is convenient to understand appropriateness to include truth because anything that is implied is then also implicated and implicature is a broader relation than implication. However, there is no consensus about using the terms in this way, and many would use *implicature* more narrowly to cover only those conditions necessary for appropriateness over and above those necessary for truth.

Both definitions above refer to the context in which sentences are used. We have ignored this so far in the case of implication though the phenomenon of indexicality means that such a reference is often required. In any case, it is crucial for appropriateness: while the contextual dependence of truth values is tied to specific vocabulary, appropriateness in the wider sense is always dependent on the specific context in which a sentence is used. In the example used above, if it was well known that I had made a bet that I could avoid using the word **I** for the next 24 hours, no one would be misled by my saying **My class was taught this morning** when I had in fact taught it myself.

Even though appropriateness as a whole depends on the context, there are

specific conditions attached to particular words that can lead to implicatures in every context. Consider, for example, this bit of dialogue:

Q: *Was the movie any good?*

A: *Yes. Even John was laughing.*

The assertion *Even John was laughing* has a number of implicatures that depend on the conversational setting (e.g., that John was at the movie and, perhaps, that it was a comedy), but it also has one that derives from presence of the word *even*. This implicature is easier to recognize than to state, but it comes to something like the claim that John doesn't laugh frequently.

Implicature is a form of non-deductive inference that we will not study in its own right, but we will not be able to ignore it because it is often difficult to distinguish from implication. This is especially true for implicatures that attach to particular words because they have the same sort of uniformity across contexts that holds for the sorts of implications we will study.

One test that can be used to distinguish implicatures from implications is to ask a *yes-no question*. When asked *Was even X laughing?* about someone X who had laughed at the movie but who was known to laugh frequently, we would not answer with a simple "No" but rather say something like, "Yes, but he'll laugh at anything." Such *yes-but* answers indicate that the sentence we were asked about is true but inappropriate. Other qualified affirmative answers can play a similar role, and we will refer to them also as *yes-but answers* even when they do not use the term *but*. To simply answer "Yes" in cases where a sentence is true but has a false implicature could mislead our audience into thinking that the sentence is entirely appropriate and thus that the implicature is true. Indeed, a true sentence with a false implicature could be described as true but *misleading*. *Yes-but* answers acknowledge the truth of such a sentence while correcting its misleading suggestions. (There are further tests that can be used to distinguish implicatures and implications, and we will consider some others in 4.1.2.)

Glen Helman 03 Aug 2010

1.3.5. Indexicality

We will give less direct attention indexicality than to implicature, but it would be hard to ignore the phenomenon. Although indexicality is most obvious in sentences with indexical words like *I*, *that*, *here*, and *yesterday*, there are other features of a sentence, most notably its tense, that can make the proposition it expresses vary with context in which it is asserted. The sentence *It's sunny* is as bound to the time of assertion as is *It's sunny now*. And, while not every sentence contains indexical terms, it is only very special sentences that are not indexical in virtue of tense.

If the propositions expressed by sentences vary with the context, it seems that the logical properties and relations of these sentences (which we trace to the propositions they express) may vary as well. Let's look at one example. The proposition expressed by the sentence *I am here* will depend on the speaker, the speaker's location, and the time of utterance. And this sentence may express the same proposition as the sentence *You are there* when the latter is used by a second speaker in an appropriately related context. There are also many contexts in which these sentences might be asserted where they would not express the same proposition. But sentences are supposed to be logically equivalent when they express the same proposition, so it seems these sentences would be equivalent when used in some contexts and not equivalent when used in others. And the same issue arises for deductive properties as well as relations; a sentence that is a tautology when used in one context might not be a tautology when used in a different context.

More broadly it may seem that we really should not speak of sentences as having deductive properties and standing in deductive relations. If a sentence expresses no fixed proposition independent of the context in which it is asserted, we can really only talk about the deductive properties and relations of sentences-in-context, of sentences each taken together with a context of use. The term *statement* has sometimes been used to speak of a particular use of a sentence. If we use this terminology, we can say that certain statements made using the sentences *I am here* and *You are there* are equivalent and that it statements rather than sentences have deductive properties and stand in deductive relations. Something like this approach would be required if we really were to study the phenomenon of indexicality. However, the logical forms on which we will focus do not include indexical elements, so it will be possible for us to ignore this aspect of meaning.

Even when indexical elements are present, we can set aside explicit reference to contexts of use when speaking only of logical properties and

relations that do not vary from context to context. For such deductive properties and relations will hold of sentences in virtue of the specific ways the propositions they express vary with the context of use—i.e., in virtue of the “characters” of these sentences. For example, we can say that sentences are equivalent if their characters lead them to express the same proposition in any context of use, and we can say that a sentence is a tautology if its character leads it to express a tautologous proposition in every context of use. Again, although the propositions expressed by *The package will arrive next Wednesday* and *The package will arrive next week* will vary depending on the time of utterance, the proposition expressed by the first sentence will always entail the one expressed by the second sentence. We will limit consideration to logical properties and relations of sentences that are independent of the context of use in this way. So, even though *I am here* and *You are there* may be used to make statements that are equivalent, we will not count these sentences as equivalent because it is not the case that, in each context, the propositions expressed by these sentences are the same. (Indeed, it is not easy to think of any single context with respect to which the two would express the same proposition since a single context would require that both be spoken by the same person.)

In fact, we can use this approach without explicitly considering the characters of sentences at all. In fact, this was done in the example in 1.2.3 that included the sentences *The package will arrive next Wednesday* and *The package will arrive next week*. There we simply took it for granted that sentences were being compared with respect to some one context, and we spoke freely of the propositions they expressed in that context without bothering to note that they expressed different propositions in other contexts. This procedure is legitimate if we not to assume anything special about the context of use. And it will be easy not to make special assumptions about the context of use because the deductive properties and relations we are interested in do not depend on this context. There is an analogy here to a typical use of variables in algebra. When numerical laws are used to manipulate algebraic formulas, it is assumed that variables appearing in those formulas have been assigned numerical values. But there is often no need to consider what those values are since the laws being used apply to all numbers.

Of course, there are things we will miss by ignoring character and context. The effects of shifting context in the course of a conversation are among the things we cannot deal with. The assertion *I am here* followed by the confirmation *Oh, so that's where you are* is a simple example of this. Another phenomenon we will miss is the exploitation of some sort of

dependence on context to convey information about the context. If I assert *Today is Tuesday*, the proposition expressed may be no more informative than is *Tuesday is Tuesday* since the first sentence, if true, merely tells us about Tuesday that it is Tuesday. But my assertion can still be helpful because someone who tries to accommodate it will need to take it to have been asserted on Tuesday, and will thus know what day it is. In short, even if the proposition expressed by *Today is Tuesday* in a given context is a tautology and conveys no information, the assumption that this sentence expresses a tautology (rather than an absurdity) in that context yields information about the context. And this way of deriving information can support a form of non-deductive inference.

On the other hand, our approach need blind us to all logical properties and relations that derive from indexical terms. We have seen this already in the case of *next Wednesday* and *next week*, but the role of the indexical terms can be less trivial than this. For example, the terms *today* and *tomorrow* are related in such a way that *Tomorrow is the day after today* is true in any context, so we can recognize it as a tautology. And we can also recognize that *Today is Tuesday* implies *Tomorrow is Wednesday*.

It would be too much to say, however, that our limited perspective will not blind us to any logical properties or relations that hold for all contexts of use. For there are relations between the meanings of indexical terms that hold in any context, but only with respect to the actual world of that context; and our approach will miss logical relations that derive from these aspects of meaning. For example, whoever is the speaker in a context will actually be speaking at the time of utterance, so the premise *Today is Tuesday* would justify the conclusion *I am speaking on Tuesday*. But this conclusion is not entailed by the premise—even given the contextually assigned meanings of the terms—since nothing about the day of the week of a given date logically necessitates someone speaking. To get a feel for the issue, it may help to look at a related example: although *I am here now* is true in the actual world of any context, it is not a tautology. That is, the proposition expressed by *I am here now* in a context of utterance is bound to be true in the actual world of that context, but this proposition will also be false in other possible worlds. And the fact that it is false in other worlds can be crucial for the meaning of sentences—such as *I am here now but I almost didn't make it*—that speak of unactualized possibilities.

1.3.6. Vagueness

One way of understanding vague terms is to suppose that their significance varies with the context of use but is not completely determined by it. The meaning of a word like **small** depends on the line to be drawn between what is and what is not small. This line is settled to some degree by features of the context of its use—whether the word appears in a discussion of molecules or of galaxies, for example—and some contexts will pin it down more precisely than others. But there is usually, and perhaps always, some indeterminacy remaining, and the class of things that count as small in a given context will have fuzzy edges.

Although the context dependence of vague terms means that vagueness is somewhat analogous to indexicality, the fact that sentences containing vague terms may not have definite truth values even when the context is specified means that we cannot handle such sentences in quite the same way as we do sentences exhibiting ordinary forms of indexicality. We can understand entailments involving indexical terms—such as

Today is Tuesday \models **Tomorrow is Wednesday**

—to hold because the propositions expressed by the two sentences are related in a certain way in every context of use. But we cannot understand the entailment

Crawfordsville is small \models **Crawfordsville is not large**

to hold for the same reason because the sentences involved may not express definite propositions in any context of use.

Still, there is a way of extending our approach to indexicality to provide an approach to vagueness. In both cases we can understand deductive properties and relations to hold for sentences because of the propositions that *would be* expressed by the sentences if certain factors were specified. In the case of the first example above, the relevant factor, the time of utterance, is specified by any actual context of use. In the second example, the relevant factors are precise *delineations* of the classes of things that the terms **small** and **large** are true of. These delineations are not fully determined by an actual context of use, but we can still say that the propositions expressed by the sentences in the second example would represent a case of entailment no matter how these delineations were specified. So, just as we will always take for granted an unspecified context of use, we will take for granted but leave unspecified precise delineations of all vague terms. And that means that we will speak of sentences as if no terms are vague.

Of course, ignoring vagueness means that we will ignore yet another important feature of language. The specific logical properties and relations we will study do not derive from vagueness, so ignoring vagueness will not limit our ability to study them. But, as with implicature and indexicality, we will miss certain ways of deriving information from things that are said. The accommodation of vague language can be analogous to accommodation of indexicality and can be an important way of conveying information. While **This is hot** will often be intended to provide information about whatever **this** refers to, it can serve instead to calibrate judgments of hotness. That is, when the audience already knows the temperature of the thing pointed to, **This is hot** can help someone to specify the significance of **hot** in a given context since accommodating this assertion requires that the thing pointed to falls within (and, indeed, some distance within) the range of hot things on any delineation of that range that is allowed by the context.

The fact that we derive information in this way provides one way of explaining a traditional logical puzzle known as the *sorites paradox* (or “paradox of the heap,” after a particular ancient example trading on the vagueness of the term **heap**). The argument

This is hot and that is only a little cooler / That is hot

is not deductively valid because the things referred to by **this** and **that** could well fall on opposite sides of a delineation. But it seems like a reasonable argument; and, if we suppose that we accommodate vague language by considering only delineations on which what has been said is not just barely true, the conclusion will be true on any delineation that accommodates the premise. The paradox comes by imagining a series of things, with each successive thing asserted to be only a little cooler than the one before with the last clearly not hot. Each step in the series could be justified by an argument like the one above, but the final result seems unacceptable.

This result would not be surprising if we understand the displayed argument to be the result of accommodation. Suppose first that we attempted to collect all the steps in the series into a single argument.*

A is hot
 B is only a little cooler than A
 C is only a little cooler than B
 ⋮
Z is only a little cooler than Y
 Z is hot

This would not be reasonable because accommodating the first premise need not place the temperature assigned to A far enough from allowable delineations to support the truth of the conclusion.

On the other hand suppose we were faced with a series of arguments

α is hot
 β is only a little cooler than α
 β is hot

one for each successive pair of terms in the series. If we really were willing to accommodate the premise at each stage, we would end up accepting the final conclusion; but the allowed delineations of **hot** would have shifted also at each stage and the final conclusion would end up acceptable.

Of course, someone who really refused to accept the final conclusion would probably refuse to accommodate the premise of one of the arguments along the way and would begin to be wary of them before that point. That is, these component arguments each stretch our willingness to accommodate a bit further, and it can only be stretched so far. The paradoxical inference can seem to be supported if we forget this, and think of the corresponding way of extracting information from an assertion as if it was like deductive inference in allowing us to link together inferences that are good individually. That is, the sorites paradox shows us that the non-deductive relation associated with this way of deriving information from the use of vague terms is not transitive.

There is terminological curiosity here. An argument like the one above running from **A** to **Z**—i.e., a multiple-premise argument that is associated with a series of two-premise arguments—is traditionally referred to as a *sorites argument*. But a sorites argument need have no connection with a sorites paradox. Although the term *sorites* is derived in both cases from the Greek term for a heap, its application to a sorites argument reflects the piling up of premises rather than any appearance in it of a vague term such as **heap**. A sorites argument constructed for the sorites paradox in its original form would be an argument about heaps that had a heap of premises.

1.3.7. Presupposition

When the **yes** answer to a **yes-no** question would be tantamount to making a true but misleading assertion, it is appropriate to answer **yes** only if we add a qualification. But it is still possible to give an affirmative answer while no qualification would make the answer **no** appropriate. Another of the complications of the simple picture of language appears in connection with **yes-no** questions for which neither answer seems legitimate.

For example, consider the question

Is John's car green?

asked about someone who does not have a car at all. In such a case, we would be at a loss to answer the question directly. This is usually explained by saying that the question *presupposes* that John has a car and has no appropriate direct answer when this presupposition does not hold. And we can say something similar about the following declarative sentences, which correspond to affirmative and negative answers to the question, respectively:

John's car is green
 John's car is not green

That is, just like the question, we can take each of these assertions to presuppose **John has a car**.

We could capture these limits on appropriateness by regarding presupposition as a sort of implicature. That is, we might say that John having a car constitutes a necessary condition for the appropriateness of either of the assertions above. But many have held that in contexts where John has no car, it is not only the case that neither sentence is appropriate but the case that neither is true. Since one would be true if the other was false, this means that neither claim would have a truth value. If this point of view is correct, what is missing in these assertions when John has no car is not some quality like informativeness or relevance that we expect in addition to truth but instead something that is a precondition for either truth or falsity. Something that is a presupposition in this strong sense is said to be a *semantic* presupposition. If John having a car is a semantic presupposition of the two sentences above, it is easy to see why they seem equally inappropriate when John has no car: each would have no truth value so the two would be in the same position as regards truth and falsity.

Semantic presupposition is unlike the phenomena we have considered so far in that it requires fundamental changes to the simple model of language and not merely additions to it. The simple model is built around the assumption

that a sentence has a truth value in every possible world, and dropping that assumption would force radical changes. And because there is no consensus, even among logicians who accept the idea of semantic presuppositions, about the exact form such changes should take, we will not attempt to incorporate failures of truth value in our model of language.

In part, we will treat semantic presupposition as we do the variety of speech acts: by not considering the examples where it may be held to occur. But we cannot avoid all the difficult cases in this way. The classic examples of semantic presupposition are sentences containing phrases employing the definite article **the** to refer to something by way of a description of it. Such phrases, which logicians classify as *definite descriptions*, cause problems because their success in referring depends on the existence of objects satisfying the descriptions they offer. For example, both the sentence **The building between Center Hall and Sparks Center is occupied** and the sentence **The building between Center Hall and Sparks Center is unoccupied** seem inappropriate when no such building exists because then the definite description **the building between Center Hall and Sparks Center** has nothing to refer to. And definite descriptions that refer contingently are so common that we cannot simply avoid all sentences containing them. The use of possessives that we saw in the example of **John's car** are also common, and they represent a closely related sort of case because **John's car** might be paraphrased by the definite description **the car John has**.

The approach we will take to these sorts of semantic presupposition does share two features with our approaches to other complicating phenomena. First, just as we do not attempt to capture relations of implicature in our study of logic, we will not attempt to capture relations of presupposition as such. However, the line between implication and presupposition is controversial, and relations between sentences like **The building between Center Hall and Sparks Center is occupied** and **There is a building between Center Hall and Sparks Center** fall in the disputed area. In 8.4.2 we will consider an account of definite descriptions according to which the first of these sentences implies the second.

Although we will not attempt to capture semantic relations of presupposition as such, we will need to apply our general account of logical properties and relations to sentences that may have such presuppositions. And we can do this only if we do not recognize the failures of truth value that result when semantic presuppositions are false, so we will assume that every sentence has a truth value under all possibilities. But, since we will eventually analyze sentences into units smaller than sentences, an assumption about the meanings

of sentences is not enough.

We will assume in addition that any term which ought to refer does have a *reference value*. We allow this to be either an actual object or an *empty* or *nil* reference value. The latter option is designed for the case of *undefined* terms like **the building between Center Hall and Sparks Center** that do not refer to actual objects. We will need to distinguish these two sorts of reference value only when we consider definite descriptions in the last chapter, so, for the most part, we will merely assume the every term has been somehow given a reference value and every sentence a truth value. The references and truth values we assume for this reason can be regarded as stipulations added to the conventional meanings of these expressions, and we will consider only logical properties and relations that hold no matter how such stipulations are made. Such assignments of supplementary semantic values are usually called *super-valuations*. Both the name and this way of handling failure of presuppositions are due to Bas van Fraassen, and the assignment of precise delineations to vague terms that was discussed in the last subsection is a further application of this idea by David Lewis. As will be case in our handling of vagueness, our assumptions of references and truth values in cases of semantic presupposition will generally stay in the background. However, we will look at the assumptions we make a little more closely in 6.1.3 when we have begun to analyze sentences into expressions that are not sentences.

Glen Helman 03 Aug 2010

1.3.s. Summary

- 1 The idea that the norms of deductive reasoning reflect a system of relations among propositions fits into a simplified picture of the function of language. According to this picture, a person's beliefs amount to a proposition that rules out a certain range of possibilities for the actual history of the universe. The desire to know more is in part the desire to narrow the range of possibilities that are left open. When language is used cooperatively, we share our abilities to rule out possibilities by using assertions to convey propositions. The sentences we can sincerely assert are the ones that are entailed by the sum total of our beliefs, and we accommodate someone else's assertion by adjusting our beliefs so that what they asserted is now entailed by our beliefs.
- 2 This picture is oversimplified and something must be said about several respects in which the actual operation of language is more complex. Each is associated with an aspect of meaning:
 - (i) the force of a sentence that marks it as an assertion or one of the many other speech acts,
 - (ii) implicatures, which convey information that a sentence does not imply,
 - (iii) semantic presuppositions, requirements for the sentence to have a truth value,
 - (iv) the character of a sentence, which reflects the way the proposition it expresses varies with the context of use due to the phenomenon of indexicality, and
 - (v) a greater or lesser degree of vagueness.While an account of how sentences express propositions is the province of semantics, these complicating phenomenon are usually said to be the subject matter of pragmatics.
- 3 Although assertion is the only speech act we will study, not even all declarative sentences have this force. J. L. Austin estimated that **assert** was only one of thousands of performative verbs that can be used to both perform and describe speech acts. Although many of these speech acts do not serve to convey propositions, their force can often be described with reference to propositions.
- 4 We will consider only what is implied by a sentence as part of its truth conditions and not further information that may be implicated as conditions for appropriate assertion beyond the requirements for truth. A false implicature will make a sentence misleading but may leave it true. One

indication of this sort of case is a **yes-but** answer to the **yes-no** question corresponding to the sentence.

- 5 Indexicality means that the propositions expressed by sentences—and thus their deductive properties and relations—can depend on the contexts in which they are used. It would be possible to compare sentences only when each was associated with a specified (but perhaps different) context—such sentences-in-context are sometimes called statements. However, we will compare sentences only within a single context of use and consider only properties and relations of sentences that hold no matter what that context is. As with implicature and presupposition, accommodating sentences to the rules governing indexical phenomena provides a way of extracting information that goes beyond entailment.
- 6 Vagueness poses problems analogous to those posed by indexicality and presupposition. As with indexicality, we will assume a context of use; and, as with presupposition, we will assume supplementary specifications of truth value (in this case precise delineations of the boundaries of vague terms). Deductively valid conclusions will not rely on information about these factors, but accommodation to vague assertions can support non-deductive inference to extract further information. One way of explaining the sorites paradox is to suppose that it rests on a failure to distinguish this sort of inference from deductive inference.
- 7 Since a semantic presupposition is something that must hold in order for a sentence to have a truth value at all, sentences with non-tautologous presuppositions can fail to have truth values. The pervasiveness of definite descriptions—which can fail to refer to anything if the facts are not right—makes it hard to simply ignore sentences with non-trivial presuppositions. Instead, we will treat all terms as if they refer, simply stipulating reference values and truth values in other cases (eventually distinguishing an empty reference value) but considering only relations between sentences that hold for all such stipulations (the method of super-valuations).

Glen Helman 03 Aug 2010

1.3.x. Exercise questions

1. For each of the following sentences, give a sentence it implies and a sentence it implicates (but does not imply) in the context described:
 - a. **My plate is clean**, as reported by a small boy who has been told to finish his vegetables by a parent saying, “Clean your plate.”
 - b. **There is a cooler in the trunk**, said in reply to someone’s expressed wish to have a beer.
 - c. **I saw the director’s last movie**, said in reply to someone who asked whether the speaker has seen a certain new movie.
2. Many philosophers would argue that the sentence **I’m Adam**, when true, expresses the same proposition as **Adam is Adam**; that is, if it is true at all, it is true in every logically possible world. The phenomenon of indexicality or deixis can help to explain how **I’m Adam** could be informative even if these philosophers are correct and it expresses a tautology when it is true. To see how this might work, ask yourself what information can be derived about a context of utterance by accommodating the use in this context of the sentence **I’m Adam**.
3. J. L. Austin, the philosopher who made people aware of the variety and importance of speech acts, suggested a way of identifying them. Look for verbs that can fit in the context **I hereby ...** (e.g., **I hereby assert that ...** or **I hereby apologize**). That is, look for, verbs that (in grammarians’ jargon) can be used in “first person indicative active sentences in the simple present tense” along with the adverb **hereby**. These are the “performative verbs” mentioned in 1.3.3. Austin suggested that there are such verbs for most speech acts. Find half a dozen as varied in character as possible.

Glen Helman 03 Aug 2010

1.3.xa. Exercise answers

1. The following are perhaps the most likely answers though they are not the only correct ones:
 - a. implies: **No vegetables are on the boy’s plate**
implicates: **The boy has finished his vegetables**
 - b. implies: **The trunk is not empty**
implicates: **There is beer in the cooler**
 - c. implies: **The speaker has seen a movie by the director in question.**
implicates: **The speaker has not seen the new movie** [with further implicatures depending on the tone of voice]
2. The truth value of **I’m Adam** depends on features of the context in which it is uttered—specifically, on the identity of the speaker. So, it is not true in some contexts of utterance. And that means that, if we assume it is used correctly, it can tell us something about the context—who the speaker is. We derive this information not simply by assuming that the actual world is a world in which the sentence is true but by assuming, more specifically, that the sentence has been uttered in a context that makes it express a true proposition. And even if it tells us nothing about the actual world to know that the person Adam is himself, it does tell us something about the context to know that the person Adam is the speaker.
3. If Austin was right, thousands of answers are possible. I will simply note a five-fold classification of speech acts along with examples of performative verbs for each sort of act. (This classification is due to the philosopher John Searle but based on Austin’s ideas.) (1) **representatives** (e.g., **assert** and **conclude**) commit the speaker to the truth of something. (2) **directives** (e.g., **order** and **ask**) are attempts to get the speaker’s audience to do something. (3) **commissives** (e.g., **promise** and **threaten**) commit the speaker to some future action. (4) **expressives** (e.g., **apologize** and **congratulate**) express a psychological state. (5) **declarations** (e.g., **sentence** and **promote**) effect a change in an institution.

Glen Helman 03 Aug 2010