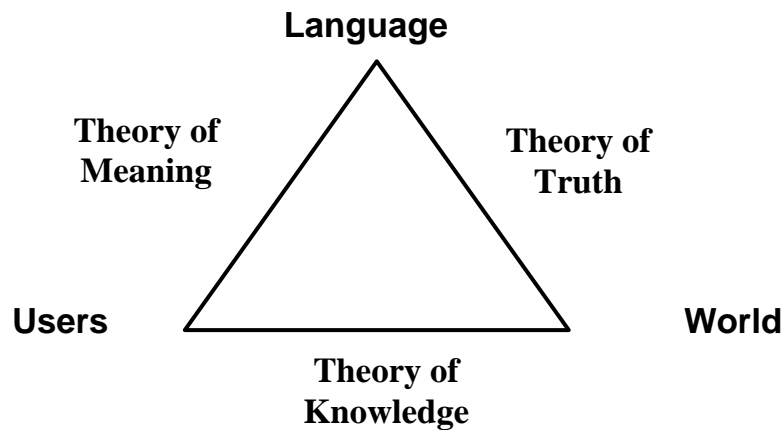


Philosophy of Language: Preliminaries

OVERVIEW

Philosophy of language is concerned with a cluster of problems involving the relationships among these three elements:

1. Language
2. Language users (speakers, hearers)
3. World



Complication: the elements overlap: both language and speakers of language are things in the world. Users can use language to say things about themselves, and about language.

PHILOSOPHY OF LANGUAGE VS. LINGUISTIC PHILOSOPHY

Should not be confused (although they often are). The first is the philosophical study of the phenomenon of language; the second is a philosophical methodology—the attempt to use linguistic methods to solve philosophical problems, or to reduce philosophical problems to problems about language.

Two prominent types of ‘linguistic philosophy’ have been **logical positivism** and **ordinary language** philosophy.

SYNTAX, SEMANTICS, AND PRAGMATICS

Syntax

Deals with the properties of and relations among symbols. The syntax of a language is its **grammar**.

Semantics

Deals with the relations between language and extra-linguistic entities. These extra-linguistic entities include the **meanings** of the expressions and the things in the world that the expressions **stand for** or are about. **Reference, truth and meaning** are the central semantic concepts.

Semantics is what will be of primary concern to us in this course.

Pragmatics

Deals with the relations between language and the users of the language, or the contexts of use of the expressions of the language.

Examples

Chomsky's illustration of the difference between syntactic and semantic correctness:

1. Furiously sleep ideas green colorless.
2. Colorless green ideas sleep furiously.

(1) is syntactically incorrect: it is not a grammatically well-formed sentence of English. (2), on the other hand, is grammatically correct ($S = NP + VP$, $VP = V + Adv$, $NP = Adj + NP$) but semantically deviant. Neither (1) nor (2) is **meaningful**, but (2) is at least a **sentence**. (1), by contrast, is just a "verbal salad."

SENTENCES AND PROPOSITIONS

These three **sentences**:

1. Snow is white.
2. Der Schnee ist weiss.
3. La neige est blanche.

all express the same **proposition**: viz., the proposition that snow is white. Propositions are theoretical entities that are appealed to in attempts to explain the concepts of truth and meaning.

Propositions as bearers of truth-value:

The proposition is what is true (or false) – the “bearer of **truth-value**.” A sentence is in a **language**, but a proposition is not. It is what a sentence “expresses” or “contains”. Propositions are thought of as being language-neutral entities.

Different sentences in the **same** language can also express the same proposition. Arguably, these two:

1. Bill is Chelsea’s father.
2. Bill is Chelsea’s male parent.

both express the same proposition. (Exactly which pairs of sentences have this feature is a matter of dispute.)

Propositions as meanings:

Propositions are also thought of the meanings of sentences. Roughly: for two sentences to have the same meaning is for them to express the same proposition.

Problems with propositions:

Identity-conditions:

Under what conditions are propositions p and q the same? It has been argued that it is impossible to provide identity-conditions for propositions (Quine). For example, how do we know that (1) and (2) above express the same proposition?

Usual test: is it possible to believe that p but not believe that q ? But this does not yield definitive results. A person who doesn’t know that ‘father’ means ‘male parent’ might believe (1) but not (2), or vice versa.

Truth and meaning collide:

We have now placed propositions on two sides of our original triangle: as **meanings** they are on the user-language interface as part of the theory of meaning, and as **bearers of truth-value** they are on the language-world interface as part of the theory of truth.

The two roles singled out for propositions – meanings and bearers of truth-value – may come into conflict. We will defer consideration of this point until later.

TYPES AND TOKENS

1. Snow is white.
2. Snow is white.

These are two sentence-**tokens** of the same sentence-**type**.

Truth-values appear to attach primarily to sentence **types**, and only secondarily or derivatively to sentence **tokens**. Each of the two tokens above is true because it is a token of a type which is true.

A sentence-type is true only if it expresses a true proposition. But suppose two different tokens of the same sentence type express different propositions? Then different tokens of the same sentence type may have different truth-values. (We will see how this can happen after we have introduced the topic of indexicality.)

INSCRIPTIONS AND UTTERANCES

Some terminology: since sentences can be either written or spoken, it will be convenient to speak of **inscriptions** (of written sentences) and **utterances** (of spoken sentences).

The type-token distinction applies here as well. (1) and (2) above are two different inscription tokens of the same inscription type. When I utter “Snow is white” twice, we have two utterance tokens of the same utterance type.

USE, MENTION, AND QUOTATION

We normally **use** words to talk **about things** that are not words. Thus, in:

1. Cicero was an orator.

the word ‘Cicero’ is used to say something about a Roman orator. But we can also talk **about the words** we use. In:

2. The name of the most famous Roman orator begins with a ‘C’.

the subject is not the orator Cicero, but his **name**. (The man himself does not begin with a ‘C’ or with any letter at all.) The name of the orator is not **used** in this sentence; rather, it is **mentioned**.

A common way of mentioning a name (or other linguistic expression) is to place it inside quotation marks. Instead of (2) we could say:

3. ‘Cicero’ begins with a ‘C’.

Here we mention the name of the orator by using its **quotation-name**. That is, by placing a name inside quotation marks, we create a name of that name.

In common practice (e.g., especially, in newspapers) the quotation marks are not used. So instead of (3) we might see:

4. Cicero begins with a ‘C’.

But it is best to regard this practice as either careless or elliptical. We certainly don’t want to conclude from (1) and (4) that:

5. Some orator begins with a ‘C’.

The correct conclusion, of course, is:

6. There is an orator **whose name** begins with a ‘C’.

In general, by placing **any** linguistic expression (string of symbols) inside quotation marks, we create a quotation-name of that expression. E.g.,

7. ‘Dances clumsily’ is in English.

‘Dances clumsily’ is a verb phrase, but its quotation-name, “dances clumsily”, is a noun-phrase. Since the quotation-name of any verb phrase is a noun phrase, we can produce curious sentences such as:

8. ‘Is in English’ is in English.

in which the subject is the quotation-name of the verb phrase (the predicate). The importance of this curious phenomenon will become apparent shortly.

USING LANGUAGE TO TALK ABOUT LANGUAGE

As the quotation convention makes clear, we can use language to talk **about** language. For the most part, this causes no problems, as the quotation convention enables us to avoid possible confusions of use and mention: we **use** one expression to **mention** another. For example, we use ‘Cicero’ to talk about (mention) the orator Cicero, and we use “Cicero” to mention the name ‘Cicero’.

Similarly, there is generally no problem when we are talking in one natural language (e.g., English) about another (e.g., French), or talking in a natural language about a formal or artificial language (e.g., a logical calculus, or a computer language).

The sentence:

1. 'Neige' is a French word.

is a sentence of English and mentions (is about) the French word for snow. But the French word for snow is not used in that sentence; it is mentioned by its quotation-name, which can perfectly well be used in an English sentence.

But compare (1) with the following:

2. 'Neige' is white.
3. Neige is a French word.
4. Neige is white.

(2) is false, since the French word for snow is not white (or any color).

(3) and (4) are not sentences in either English or French, for each combines French and English words in a single string.

Now consider cases in which we use a given language to talk about expressions of that very language:

5. 'Snow' is an English word.

is a sentence of English that is used to say something about an English word. And (5) is true, for 'snow' is indeed an English word. In most cases, such sentences are not problematic. But problems quickly emerge.

SELF REFERENCE AND THE LIAR PARADOX

Suppose our language allows us to mention its own expressions (including sentences) by using quotation-names and other such devices. We will quickly discover that this is sufficient to allow for **self-reference** — for a sentence to say something about **itself**.

And now suppose our language also contains a **truth predicate**. That is, we can use English to say about an English sentence that it is (or is not) true. These two assumptions seem quite innocuous. But, as we will see, they lead to paradox.

Start with a simple English sentence, such as:

1. Snow is white.

This is an English sentence, which fact is duly noted as:

2. Sentence (1) is in English.

Now we bring in the truth-predicate. Since (1) is an English sentence, and (2) says that (1) is an English sentence, what (2) says is true. That is:

3. Sentence (2) is true.

Notice that (2) is an English sentence that says something about another sentence, (1). But what (2) says about (1) it should be possible for a single sentence to say about **itself**.

4. Sentence (4) is in English.

This is an English sentence which, like (2), claims that a certain sentence is in English. The only difference here is that the sentence that (4) is talking about is (4) itself. So we have **self-reference**, but there does not seem to be anything problematic about it. Indeed, the self-referential sentence (4) seems to be true. It claims that a certain sentence is in English, and that sentence is, indeed, in English. So we can write a sentence that states precisely this:

5. Sentence (4) is true.

(4) is harmlessly self-referential (it says of itself that it is in English); and (5) contains the truth-predicate (it says that sentence (4) is true). What happens if we put the truth-predicate into a self referential sentence?

6. Sentence (6) is true.

This is a puzzling sentence. In the case of the self-referential (4), we could tell by inspection that it was true (it said it was in English, and indeed it is in English). But we cannot tell by inspection whether (6) is true or false. Since (6) says that a certain sentence is true, we should ask ourselves whether that sentence is true or not. If it's true, then so is (6); if it's false, then so is (6).

The problem is that the sentence (6) is talking about is (6) itself. In order to assess (6) for truth or falsity, we must first assess the sentence it is talking about. But that sentence is just (6) itself. So we cannot begin.

If the sentence under discussion, (6), is true, then the sentence making a claim about it, (6), is true. So, if (6) is true, then (6) is true.

If the sentence under discussion, (6), is false, then the sentence making a claim about it, (6), is false. So, if (6) is false, then (6) is false.

Notice that nothing paradoxical has yet emerged. The only puzzle is that we don't know whether (6) is true or false. Still, if it's true, it's true, and if it's false, it's false. Nothing says we have to know the truth-value of **every** sentence.

But with one small change, we confront the **liar paradox** (what Tarski calls 'the antinomy of the liar'):

7. Sentence (7) is not true.

What are we to say about (7)? Is it true or false? Neither assumption is satisfactory.

If (7) is true, then what it says is the case. And what it says is that (7) is not true. So if (7) is true, then it's not true.

If (7) is not true, then what it says is not the case. Since what it says is that (7) is not true, it is not the case that (7) is not true. But then (7) is true. So if (7) is not true, then it's true.

So (7) is true iff it's not true. This is a contradiction. And we have reached this contradiction merely by using English to make a statement about the truth of a certain English sentence. This is the liar paradox: we can construct a sentence that is true iff it isn't true. So it can't be true, and it can't be false, and yet we seem to understand perfectly well what it says.

Tarski (§§1-11 on pp. 85-93) shows that this result is the inevitable consequence of a semantic theory of truth for (what he calls) a **semantically closed language**.

- A semantically closed language is one which contains both a truth predicate and names for all its expressions.
- A semantic theory of truth is one which explains 'is true' in terms of the following schema (T):

X is true if, and only if, *p*

where '*p*' is replaced by any sentence, and 'X' is replaced by a name of that sentence.

The semantic theory of truth holds that an adequate account of truth is one that entails every instance of schema (T). That is, it should entail all of the following, among others:

'Snow is white' is true if, and only if, snow is white.

'Snow is green' is true if, and only if, snow is green.

and so forth. So, any adequate semantic theory of truth will entail:

8. ‘Sentence (7) is not true’ is true iff sentence (7) is not true.

But sentence (7) = ‘Sentence (7) is not true’. So, substituting ‘sentence (7)’ for ‘‘Sentence (7) is not true’’ (via Leibniz’ Law) in (8) we obtain:

9. Sentence (7) is true iff sentence (7) is not true.

And (9) is a contradiction. So, if we apply the semantic theory of truth to a semantically closed language, we get a contradiction.

Quine’s version of the Liar Paradox

The following observation might seem to offer a way out of the liar paradox: it is only **some** tokens of ‘(7) is not true’ that give rise to paradox—viz., those that are numbered ‘7’. To get a **type** that involves paradox seems to require indexical self-reference (i.e., use of context-sensitive words like ‘I’ or ‘this’):

‘I am not true’ ‘This very sentence is false’

It does not seem possible to produce a version of the Liar that does not depend on either accidental naming (numbering) or indexical self reference. But Quine showed otherwise. Here’s his version:

‘Appended to its own quotation is false’ appended to its own quotation is false.

This sentence ascribes falsity to a certain sentence—one that takes a certain sequence of words and appends them to a certain quotation-name. The sequence in question is the following: ‘appended to its own quotation is false’. The quotation name in question is just the very same sequence placed inside quotation marks. So it turns out that the sentence being said to be false is just the very sentence that is doing the saying! So here we have a Liar that avoids both accidental numbering and indexical self-reference. The problem runs deeper.

Can we use the Quine sentence to derive a contradiction from schema (T)? Certainly. According to schema (T), the Quine sentence is true iff what it says is the case; that is, iff ‘appended to its own quotation is false’ appended to its own quotation is false. But notice that the Quine sentence just **is** the result of appending the phrase ‘appended to its own quotation is false’ to its own quotation. That is, the expression ‘‘appended to its own quotation is false’ appended to its own quotation’ refers to the Quine sentence. So that is the expression that we’ll use to replace ‘X’ in schema (T). (The Quine sentence itself will replace ‘p’.) Given these replacements in schema (T), we obtain:

‘Appended to its own quotation is false’ appended to its own quotation is true iff
‘appended to its own quotation is false’ appended to its own quotation is false.

Now back to Tarski.

OBJECT LANGUAGE AND METALANGUAGE

Did Tarski conclude that the semantic theory of truth was untenable? No. His conclusion was that **any semantically closed language is inconsistent**. Since he thought that we could not do without a truth-predicate, his solution was to observe a sharp distinction between the language we are talking about (the **object** language) and the language we use to talk about it — the **metalanguage**.

It is easy to see that by observing this distinction, we can avoid the liar paradox. Consider our “liar” sentence:

7. Sentence (7) is not true.

What language is (7) a sentence of? We are inclined to say ‘English’, but that cannot be quite right. For if that is correct, then English is both the metalanguage and the object language, which violates Tarski’s requirement that no language with a truth-predicate can be its own metalanguage. So (7) fails to be a sentence in any language, and we cannot write a liar sentence in any semantically acceptable (consistent) language.

Hence we must interpret schema (T) in the following way:

X is true if, and only if, p

All instances of (T) are in the **metalanguage**. We replace ‘ p ’ with a translation into the metalanguage of any **object language** sentence, and ‘X’ is replaced by a metalinguistic name of that object language sentence.

For example:

‘Snow is white’ is true iff snow is white.

The entire sentence is in the metalanguage. The subject of the metalinguistic predicate ‘is true’ is the quotation-name (in the metalanguage) of the object language sentence ‘Snow is white’. The right hand side of the biconditional is the translation of that object language sentence into the metalanguage.

This can be confusing, since both object language and metalanguage here look like English. It’s easier to follow where the two languages are obviously different. Here’s an instance of schema (T) using English as a metalanguage for French:

‘La neige est blanche’ is true (in French) iff snow is white.

Here, the distinction between object language and metalanguage is obvious. The sentence is in English (the metalanguage); the subject of ‘is true’ is the quotation-name of the French sentence ‘La neige est blanche’ (which counts as an expression of English!); the right hand side of the biconditional is the translation into the metalanguage of the object language sentence ‘La neige est blanche’.

(The problem with English is that we seem to insist on using it as its own semantical metalanguage, and so Tarski’s solution in terms of levels of language — hierarchies of metalanguages — seems artificial when applied to a single natural language.)

AMBIGUITY, VAGUENESS, AND INDEXICALITY

Ambiguity

An ambiguous sentence is one that can mean two (or more) different things.

Ambiguous words

A sentence can be ambiguous because it contains an ambiguous word:

‘Paul met Ringo at the bank’.

Syntactical ambiguity

But a sentence can be ambiguous without containing any ambiguous words. The ambiguity can be structural (syntactical). Example:

‘Every arrow is not green’.

Does this say of each arrow that it is non-green? Or does it merely deny the universal generalization that every arrow is green?

Sometimes sentences are syntactically ambiguous, but the ambiguity is resolved by semantic considerations. Example:

‘There’s a sucker born every minute’.

The idea that some one person is both a sucker and born every minute is so implausible (each person is born only once) that one immediately grasps the intended meaning: *For any minute, there is some sucker or other born at that minute.*

A more interesting syntactical ambiguity:

‘Visiting relatives can be tiresome’.

What is it that we are being told can be tiresome — *relatives*, or *visiting*? That is, our sentence might be interpreted in either of the following ways:

‘*To visit* relatives can be tiresome’.

‘Relatives *who visit* can be tiresome’.

Vagueness

Vagueness is not the same as ambiguity. ‘Ringo is bald’ is vague, but not ambiguous. It is vague because it invokes a **vague concept**. What precisely is it to be bald, i.e., to satisfy the predicate ‘is bald’? There are not two or more different meanings of ‘bald’. (Perhaps there are: both tires and human heads can be bald, but only the latter because of a lack of hair. So perhaps ‘bald’ is both ambiguous and vague.)

Indexicality

Many sentences contain words whose reference depends in a crucial way on the context in which they are uttered. Such words as:

Personal pronouns:

‘I’, ‘you’, ‘he’, ‘she’, ‘it’, ‘they’, etc.

Demonstratives:

‘This’, ‘that’, etc.

Temporal Expressions:

‘Now’, ‘today’, ‘yesterday’, etc.

Locative Expressions:

‘Here’, ‘there’

When I use the word ‘I’, it refers to Marc Cohen; when Bill Clinton uses the word ‘I’, it refers to Bill Clinton. When someone says ‘yesterday’ on March 24, it refers to March 23, etc.

Thus, the reference of an indexical expression (in a given context) is determined by (among other things) the context. Which day a given token-occurrence of ‘yesterday’ refers to depends upon the day on which that token was uttered.

Terminological note

Different authors have used different technical terms for what I (following David Kaplan) have been calling ‘indexicals’:

Russell:	‘egocentric particulars’
Reichenbach:	‘token reflexives’
Early Kaplan	‘demonstratives’

Warning!

Not all occurrences of pronouns such as ‘he’, ‘it’, ‘you’, ‘that’, etc. are indexicals. Consider these sentences:

He who hesitates is lost.

You can lead a horse to water, but you can’t make him drink.

These sentences contain no indexicals. The pronouns ‘he’, ‘you’, and ‘him’ are cross-referential—they refer back (or ahead!) to another part of the sentence, rather than having their reference established, in part, by the context of utterance. No pointing (even implicit) is required. Test for indexicality: whether you can paraphrase away the pronoun—if you can, it’s not indexical. Here are paraphrases for our sentences above.

Every person who hesitates is lost.

One can lead a horse to water, but one can’t make a horse who has been led to water drink.

Ambiguity vs. Indexicality

The fact that a sentence contains an indexical expression does **not** make the sentence **ambiguous**. The sentence ‘I am hungry now’ is not ambiguous: it expresses, quite unambiguously, that the speaker is hungry at the time of utterance.

Sentences and Propositions: reprise

We can now return to an issue we postponed: a possible collision between the two roles of propositions (as both **meanings** and **bearers of truth-value**).

For sentences containing indexicals, at least, it does not seem that one and the same entity can be both a meaning and the bearer of a truth-value.

Let S = the sentence-type 'I am hungry now'. S is unambiguous: it means that the person who utters it is hungry at the time of utterance. So S should have a single meaning. But if the meaning of S is the proposition that S expresses, what proposition might that be?

The problem is that **different tokens of S express different propositions**. Bill's utterance of S at midnight expresses the proposition that Bill is hungry at midnight. George's utterance of S at noon expresses the proposition that George is hungry at noon. Bill uses it to say something true; George uses it to say something false.

So the meaning of S is not the same as the proposition(s) it expresses. Hence, if we want to preserve propositions as the bearers of truth-value (as I think we should), then we need to find other entities to be meanings (at least for sentences containing indexicals).

A provisional theory for sentences containing indexicals

- Sentence-**types** have meaning, but do not express propositions.
- Sentence-**tokens** express propositions (which can differ in truth-value).
- Residual problem: what can replace propositions as the meanings of sentence-types?