

# The Liar

PHIL 20229  
Jeff Speaks

April 24, 2008

1	The Liar paradox . . . . .	1
2	Grounding and truth . . . . .	2
3	The strengthened liar . . . . .	3
4	Hierarchical solutions . . . . .	3
5	Ungrounded sentences and meaninglessness . . . . .	5

## 1 The Liar paradox

It seems that some sentences refer to themselves. For example, “This sentence is the one that I am uttering now” seems to make sense, and be true.

But other sentences which refer to themselves are not so harmless. Consider the following sentence:

L1. L1 is false.

This sentence says of itself that it is false. But this seems to lead to the following line of argument:

Suppose that L1 is true. Then what it says must be the case; so it must be the case that L1 is false. So, if L1 is true, then it is false.

Suppose that L1 is false. This is what L1 says is the case; since if what a sentence says is the case is in fact the case the sentence is true, if L1 is false, then L1 must be true.

So L1 has the peculiar property that it is true if false and false if true. Since nothing can be *both* true and false, it follows that L1 is *neither* true nor false.

This is not by itself paradoxical. But it may seem somewhat surprising — it can seem as though every declarative sentence must either describe the world accurately or not describe it accurately; it is true in the former case, and false in the latter. So what we

might want is a theory about how ‘true’ works which explains how sentences like L1 can turn out to be neither true nor false.

## 2 Grounding and truth

One attempt to do this is via the notion of a sentence being *grounded*. Suppose we want to figure out how the word ‘true’ works — you can imagine that you need to explain it to someone who understands English, but just doesn’t understand the word ‘true.’ You might give them the following general rule: just in case they’re willing to endorse a sentence ‘*S*’, then they should also be willing to endorse the sentence ‘*S* is true.’ They might use this rule to decide what sentences involving the word ‘true’ to accept. For example, if they are willing to endorse, or accept

Grass is green.

they will now also be willing to endorse, or accept,

‘Grass is green’ is true.

However, since they reject

Grass is red.

they will not be willing to endorse, or accept,

‘Grass is red’ is true.

In fact, they will know to reject it. They might also then master the rule which tells them that if they are willing to reject a sentence ‘*S*’, they should be willing to endorse the sentence ‘*S* is false.’

This will give them a lot of information about how to use the words ‘true’ and ‘false.’ But it will not tell them which to apply to every sentence of the language. Consider, for example, the ‘truthteller’ sentence:

T1. T1 is true.

If our imaginary language learner looks at T1 to decide whether to accept or reject it, he first notices that the word ‘true’ occurs in the sentence; so, following his rule, he looks at the sentence to which ‘truth’ is applied and asks whether he accepts *that* sentence. But this sentence itself involves the notion of truth . . . and so we never get an answer to the question of whether we should accept or reject T1. When this is the case, we say that a sentence is *ungrounded*. Ungrounded sentences, on this view, are neither true nor false.

It comes as no surprise that L1 is also ungrounded. This may explain why it is neither true nor false, and help to make plausible the solution to the Liar paradox discussed above.

### 3 The strengthened liar

This looks great as far as it goes; but we are not out of the woods yet. Consider the following sentence, which is called the ‘strengthened liar’:

SL1. SL1 is not true.

SL1 is ungrounded, so on the above view it is neither true nor false; in other words, it is not true and not false. So, in particular, SL1 is not true. But this is just what SL1 says is the case, so it must be true. So it looks like the attempt to solve the Strengthened Liar via the view that ungrounded sentences are neither true nor false is a failure.

Maybe we should change the view in order to provide an account of SL1 as follows: perhaps we should say that ungrounded sentences are not just neither true nor false, but also neither true nor not true.

But this modified solution itself faces two apparently insuperable problems:

- To say that SL1 is neither true nor not true is so say that it is both not true and not true. But that is a contradiction.
- To say that SL1 is neither true nor not true is so say that it is both not true and not true. But that means that it involves saying that SL1 is not true; which is just what SL1 says, which means that SL1 is true.

### 4 Hierarchical solutions

These problems might suggest that we need a completely different sort of solution to the Liar paradox. An alternative to the grounding approach is the view of truth defended by the midcentury Polish logician, Alfred Tarski.

Tarski thought that a language  $L$  would give rise to the Liar paradox if it had three features:

1.  $L$  contains the resources for stating facts about its own semantics, such as the term ‘true’ applying to sentences of  $L$ . Tarski calls this  $L$  being ‘semantically closed.’
2.  $L$  contains the capacity to refer to its own expressions.
3. For every meaningful declarative sentence  $S$ , the ‘T-sentence’ formed using  $S$  and a name of  $S$  as follows

‘ $S$ ’ is true if and only if  $S$

is true.

Since these three assumptions about a language lead to a contradiction, we must reject one of them. The problem is that each of the three seem very plausible. In particular, it is very plausible that English meets all three conditions.

Tarski's conclusion is that we must reject the first of the three assumptions: "Accordingly, we decide *not to use any language which is semantically closed* in the sense given." According to Tarski, no language can contain a word 'true' which can apply to its own sentences.

So how should we understand claims involving the word 'true'? Strictly speaking, we should think of them as belonging to a different language. Let's call the set of sentences which can be formed by English words other than 'true' and 'false' the language English<sub>1</sub>. Then we can introduce a word, 'true<sub>2</sub>', which applies to the true sentences of English<sub>1</sub>. However, 'true<sub>2</sub>' is not itself a word of English<sub>1</sub>, so it cannot be a part of sentences of English<sub>1</sub>. Can we talk about the truth-values of sentences involving 'true<sub>2</sub>'? Yes; but only by bringing in another predicate, 'true<sub>3</sub>.'

Why does this view that no sentence can contain a truth predicate which can apply to itself stop the Liar paradox from arising?

However, there is something puzzling about Tarski's response to the Liar. Is he saying that no languages do contain their own truth predicates, or that no language *should*? It seems that Tarski had the second view:

"A characteristic feature of colloquial language ... is its universality. If we are to maintain this universality of everyday language in connexion with semantical investigations, we must, to be consistent, admit into the language, in addition to its sentences and other expressions, also the names of these sentences and expressions, and sentences containing these name, as well as such semantic expressions as 'true sentence', 'name', 'denote', etc. ... But it is presumably just this universality of everyday languages which is the primary source of all semantical antinomies, like ... the liar ... These antinomies seem to provide a proof that every language which is universal in the above sense, and for which the normal laws of logic hold, must be inconsistent." (from "The Concept of Truth in Formalized Languages")

But this is puzzling — what does it mean to say that, all this time, we have been speaking an 'inconsistent language'?

A more moderate application of Tarski's idea is that 'true' in English is systematically ambiguous. Sometimes we mean 'true<sub>2</sub>', sometimes 'true<sub>3</sub>', etc., depending on the sort of sentence which we are saying to be true.

However, this approach faces two sorts of problems, both of which were pointed out in a 1975 paper by the philosopher Saul Kripke, entitled 'Outline of a Theory of Truth'. The first is that speakers won't usually be in a position to know *which* of the hierarchy of truth predicates they want to use on a given occasion:

Unfortunately this picture seems unfaithful to the facts. If someone makes such an utterance as (1), he does *not* attach a subscript, explicit or implicit, to his utterance of ‘false’, which determines the “level of language” on which he speaks. An implicit subscript would cause no trouble if we were sure of the “level” of *Nixon’s* utterances; we could then cover them all, in the utterance of (1) or even of the stronger

(4) All of Nixon’s utterances about Watergate are false.  
 simply by choosing a subscript higher than the levels of any involved in Nixon’s Watergate-related utterances. Ordinarily, however, a speaker *has no way of knowing the “levels” of Nixon’s relevant utterances*. Thus Nixon may have said, “Dean is a liar,” or “Halderman told the truth when he said that Dean lied,” etc., and the “levels” of these may yet depend on the levels of Dean’s utterances, and so on. If the speaker is forced to assign a “level” to (4) in advance [or to the word ‘false’ in (4)], he may be unsure how high a level to choose; if, in ignorance of the “level” of Nixon’s utterances, he chooses too low, his utterance (4) will fail of its purpose. The idea

The second is that in some cases it looks like there is *no* level of truth predicate which will do the job:

Another situation is even harder to accommodate within the confines of the orthodox approach. Suppose Dean asserts (4), while Nixon in turn asserts

(5) Everything Dean says about Watergate is false.  
 Dean, in asserting the sweeping (4), wishes to include Nixon’s assertion (5) within its scope (as one of the Nixonian assertions about Watergate which is said to be false); and Nixon, in asserting (5), wishes to do the same with Dean’s (4). Now on any theory that assigns intrinsic “levels” to such statements, so that a statement of a given level can speak only of the truth or falsity of statements of lower levels, it is plainly impossible for both to succeed: if the two statements are on the same level, neither can talk about the truth or falsity of the other, while otherwise the higher can talk about the lower, but not conversely. Yet intuitively, we can often assign unambiguous truth values to (4) and (5). Suppose Dean has made at least one true statement about Watergate [other than (4)]. Then, independently of any assessment of (4), we can decide that Nixon’s (5) is false. If all Nixon’s other assertions about Watergate are false as well, Dean’s (4) is true; if one of them is true, (4) is false. Note

## 5 Ungrounded sentences and meaninglessness

So the appeal to a hierarchy of truth predicates doesn’t seem altogether satisfactory, either.

Could we return to the idea that there’s something wrong with ungrounded sentences, and that this holds the key to solving the Liar? Perhaps we could say that sentences which are ungrounded — like SL1 — simply fail to say anything at all.

Recall above our argument against the application of the ‘ungroundedness’ solution to the strengthened Liar:

To say that SL1 is neither true nor not true is so say that it is both not true and not true. But that means that it involves saying that SL1 is not true; which is just what SL1 says, which means that SL1 is true.

This argument assumes that the proponent of this solution must say that SL1 is neither true nor not true. But perhaps this what we must avoid. Perhaps the proponent of this solution should reject both of the following claims:

SL1 is true.

SL1 is not true.

BUT perhaps rejecting these claims does not involve accepting their negations; perhaps sometimes we should reject both a sentence and its negation. If that is the right attitude in this case, then we cannot move from the rejection of

SL1 is true or SL1 is not true.

to the acceptance of

It is not the case that (SL1 is true or SL1 is not true).

i.e.

SL1 is neither true nor not true.

There are a few further problems here. First, it is hard to understand what it would mean to reject a sentence other than simply claiming that it is not true.

Second, it is very tempting to think of the present view as including the claim that SL1 is ungrounded. This seems to indicate that the predicate 'is ungrounded' is an intelligible expression of our language. But now consider the following revised strengthened liar:

SL2. SL2 is not true or ungrounded.

SL2 appears to be ungrounded; since SL2 just says that it is either ungrounded or not true, that seems to indicate that SL2 is true. But no sentence can be both ungrounded and true.