

# OUTLINE OF A CONTEXTUAL THEORY OF TRUTH

Keith Simmons

## I. The Liar Paradox

Suppose I say:

This very sentence is false.

If the sentence is true, then what I say is the case. But what I say is that the sentence is false; and so the sentence is false. Suppose, on the other hand, that the sentence is false. But that is just what I say, and so the sentence is true. Either way we obtain a contradiction: we are landed in the paradox of the Liar.

Or suppose that I've just seen a false sentence written in a nearby classroom. Intending to expose this falsehood, I now write on the board the following sentence:

(L) The sentence written on the board in room 213 of Caldwell Hall at 3.05 pm on 5/7/91 is not true.

But I am mistaken about my whereabouts: I am myself in room 213. The sentence (L) that I've just written generates the Liar paradox. For if we assume (L) is true, we are led to the conclusion that (L) is not true; and if we assume (L) is false, then it follows that (L) is not true, and from this we are led to the conclusion (L) is true.

Many ways out of the Liar have been proposed over the last 24 centuries or so. Though little survives of ancient work on the Liar, there is a very rich medieval and contemporary literature on the Liar. It has been suggested, for example, that the Liar sentence (L) is meaningless; or that it expresses no proposition; or that it has no truth value; or that it has both values, true and false. Some have pointed the finger at the self-referential nature of we suggested that we must distinguish between levels of language. Still, it's agreed on all sides that Liar sentences like (L) is semantically sick, semantically pathological.

## II. A Strengthened liar discourse

But this simple observation, that (L) is semantically pathological, generates a new strengthened version of the paradox, which I shall call the Strengthened Liar. We have just established that (L) is pathological. It follows that (L) cannot be true: we may infer that (L) is not true. But now, that's just what (L) says. And now we may conclude that (L) is true, after all.

It's important to notice that this reasoning is quite natural. There is nothing technical or *recherche* here: the argument uses only notions that are available to the ordinary speaker. Such a speaker would find each step perfectly intuitive. You and I can appreciate that (L) is semantically sick, and we can go on to reason that since it's pathological, it's not true, and since that what (L) says, (L) is true. Now, we might try to block this strengthened reasoning in some way or other. But all attempts along these lines seem to me counterintuitive and *ad hoc*. I would recommend a different approach: don't block the argument, but let it go through. It is, I'd suggest, a correct piece of reasoning.

Let's examine the argument a little more closely. We can distinguish two stages of the argument. First, there is the reasoning that leads to the conclusion that (L) is pathological. Second, there is the reasoning that leads to the conclusion that (L) is true, given its pathological status. There are two steps in the reasoning at this second stage: from (L)'s being pathological, we infer (P) (L) is not true, and from (P), together with what (L) says, we infer (R) (L) is true.

Corresponding to the two stages of the argument, there are two distinct evaluative contexts, the original context in which I uttered (L) (using utterance in a broad sense), and the subsequent reflective context in which we evaluate (L) by reflecting on its pathological status. In the original context, the sentence (L) is pathological. In the reflective context, (L) is true, because it is pathological in the original context, and because of what it says.

We may identify a number of contextually relevant differences between the two stages of the strengthened Liar reasoning. Clearly there are differences of time place: for example, (R) is uttered later, and (R) is not written on the board in room 101. And if we suppose that someone other than me carries out the strengthened reasoning, there is a difference of speaker too. Still, the familiar contextual parameters of speaker, place and time do not tell the whole story.

A fourth difference is the relation that each stage of the reasoning bears to the reasoning as a whole. In general, the correct interpretation of an expression or a stretch of discourse may depend on the larger discourse in which it is embedded. It is a familiar fact that surrounding discourse can disambiguate and specify. Now, the reflective reasoning is second not only in temporal order, but in logical order, too. The reflective

reasoning starts out from a subconclusion, that (L) is pathological, established by the first stage of the reasoning. This logical order constitutes a difference in the relation that each stage of the reasoning bears to the reasoning as a whole.

A fifth difference is found in speaker's intentions. When I first utter (L), my intention was to pick out a sentence written in a room nearby. This intention is overridden, given the time and place of my utterance: I have unwittingly landed myself in paradox. Throughout the second stage of the reasoning, we have a very different intention, to treat (L) as a pathological utterance, and evaluate (L) on this basis. This intention is not overridden. On the contrary, the fact that we intentionally take (L) to be pathological leads to our final evaluation of (L) as true.

And sixth, there is, in our present case, a difference of relevant information. When I first utters (L), the information that (L) is pathological is not available to me. But of course the information is available throughout the second stage of the reasoning: it is made available by the reasoning of the first stage. The reasoning of the second stage should be interpreted as incorporating this information.

So we distinguish two distinct evaluative contexts. Let 'true<sub>L</sub>' abbreviate 'true in the original context of (L)', and let 'true<sub>R</sub>' abbreviate 'true in the reflective context'. Then we may represent my pathological utterance as:

(L) (L) is not true<sub>L</sub>.

The conclusion of the first stage of the Strengthened Liar argument is that (L) is neither true<sub>L</sub> nor false<sub>L</sub>: (L) is pathological in the original context. At the second stage, we argue that since (L) is pathological in the context of (L), it follows that:

(P) (L) is not true<sub>L</sub>.

Given (P), and given that (L) says that (L) is not true<sub>L</sub>, we infer:

(R) (L) is true<sub>R</sub>.

Corresponding to the occurrences of 'true' in (L) and in (R), there are two truth schemas:

(T<sub>L</sub>) ... is true<sub>L</sub> iff ...

and

(T<sub>R</sub>) ... is true<sub>R</sub> iff ...

If we insert my Liar sentence into T<sub>L</sub>, we obtain a contradiction. This is just what is going on at the first stage of the Strengthened Liar argument. But if we insert (L) into T<sub>R</sub>, the right side of the biconditional is true (since (L) isn't true in the original context), and so

(L) is true (true<sub>R</sub>, that is). And this is just the conclusion (R) of the Strengthened Liar reasoning. So the first stage of the Strengthened Liar argument corresponds to the failed attempt to evaluate (L) via the (T<sub>L</sub>) schema, and the second stage corresponds to the evaluation of (L) via the (T<sub>R</sub>) schema.

Let us now focus on two occurrences of the predicate 'true': the occurrence in (L), and the occurrence in (R). The occurrence of 'true' in (L) does not have (L) in its extension because, in the original context, (L) is not true. But the occurrence of 'true' in our final evaluation (R) does include (L) in its extension. (L) is in the extension of true<sub>R</sub> but is not in the extension of true<sub>L</sub>. So there is a shift in the extension of 'true' according to context. In other words, 'true' is a context-sensitive term.

So, on my view, Strengthened Liar discourses bring out the contextsensitivity of truth (and falsity). They also point up a special kind of evaluation, of which (R) is an example. (R) evaluates a pathological sentence on the basis of its pathologicality. (R) evaluates (L) as true in (R)'s context of utterance, because (L) is not true (since pathological) in its original context of utterance, and because (L) says it's not true in its context of utterance. I shall call evaluations like (R) reflections. The intuitiveness of the Strengthened Liar argument shows that we have semantic intuitions that produce reflections. For a given pathological utterance, there is a reflective context in which that utterance receives a truth value, a truth value determined by its pathologicality and by what it says. In the appropriate contexts, 'This sentence is not true' is true, while 'This sentence is false' and 'This sentence is true' are each false.

### III. Tarskian approaches

What I have said so far is enough to place my proposal in the category of contextual solutions to the Liar. In recent years, contextual approaches to the Liar, in various forms, have been offered by Parsons, Burge, Gaifman, and Barwise and Etchemendy.<sup>1</sup> Now all these views are not only contextual, they are hierarchical too. In

<sup>1</sup> Charles Parsons, "The Liar Paradox", *Journal of Philosophical Logic* 3 (1974), pp. 381-412, reprinted with a postscript in R. L. Martin (ed.), *Recent Essays on Truth and the Liar Paradox*, Oxford 1984, pp. 9-45; Tyler Burge, "Semantical Paradox", *Journal of Philosophy* 76 (1979), pp. 169-198, reprinted with a postscript in R. L. Martin (op. cit.), pp. 83-117; Haim Gaifman, "Operational Pointer Semantics: Solution to self-referential puzzles I", *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, ed. M. Vardi and M. Kaufman, Los Altos, California 1988, 43-60; Jon Barwise and John Etchemendy, *The Liar: an essay in truth and circularity*, Oxford 1984.

some form or other, these views incorporate a Tarskian hierarchy. The view I shall propose does not: my proposal will be strongly anti-Tarskian. But this is to anticipate. Let us turn first to Tarskian approaches.

Tarski proved, roughly speaking, that a classical formal language cannot contain its own truth predicate. There is no predicate of the language which applies to exactly the true sentences of the language, on pain of contradiction. The concept of truth for the language cannot be expressed by the language. The concept can be expressed in a metalanguage, an essentially richer language which contains a truth predicate for the first language. And now, Tarski's theorem applies in turn to this metalanguage" the metalanguage cannot express its own concept of truth. For that, a further metalanguage is required. In this way, a hierarchy of formal languages is generated, each containing a truth predicate for the preceding language.

Now a 'Tarskian' approach to the Liar attributes to natural language the same kind of hierarchical structure that Tarski established in the formal setting: a natural language like English is split into a series of increasingly comprehensive language, each containing a truth predicate for the previous one. (Note the scare quotes here: Tarski himself thought that such an approach was misguided.) The Liar is avoided by the distinction between levels of language.

We can now see what a Tarskian would say about our Strengthened Liar reasoning. Consider again our two truth predicates, 'true<sub>L</sub>' (true in the original context of (L)) and 'true<sub>R</sub>' (true in the reflective context). According to a Tarskian analysis, the predicate 'true<sub>R</sub>' is a more comprehensive truth predicate than 'true<sub>L</sub>'. Correspondingly, the shift between the two stages of the strengthened reasoning marks a shift in language level: we move from a semantically less rich language to a semantically richer language. I believe there are a number of serious problems with the Tarskian approach. Perhaps the first that comes to mind is this: the Tarskian approach is highly artificial. Isn't there just a single truth predicate in English, the predicate 'true', rather than an infinite series of truth predicates? And isn't there just a single language English, rather than an infinite series of stratified languages? It was for these reasons that Tarski thought the Tarskian approach was misguided: if we adopt this approach, then natural language will lose its naturalness, and take on instead the character of a formalized language.

Second, consider how we would begin our Tarskian story. The sentence 'Snow is white' would be assigned level 0, since it is a sentence free of semantical predicates. Suppose I now say: 'Snow is white' is true". Since I'm evaluating a sentence of level 0, my utterance is of level 1. So let us assign the subscript 0 to the predicate 'true' in my utterance, to indicate that my predicate applies to just the sentences of level 0. The truth predicate 'true<sub>0</sub>' has in its extension all sentences of level 0, and no others. So there is a vast number of truths (of level 1, and beyond) which are not in the extension of my use of 'true'. My use of 'true' applies only to a tiny fraction of the truths. Gödel remarks of Russell's hierarchy of types that

"... each concept is significant ... for an infinitely small portion of all objects."<sup>2</sup>

Similarly, the Tarskian's hierarchical approach places massive, counterintuitive restrictions on the extension of a use of 'true'.

Third, it is sometimes entirely unclear what level to assign to an utterance involving truth. Suppose you say: "Every sentence is true or not true". What level should I, as a Tarskian, assign to your utterance?

Fourth, the Tarskian approach evaluates utterances as pathological when they are straightforwardly true or false. Suppose Joe says: "Everything Mary says about me is untrue". And suppose Mary says "Everything Joe says about me is untrue". What levels should we, as Tarskians, assign to these utterances? Since, presumably, each utterance includes the other in its scope, there seems to be no reasonable, consistent way of assigning levels. The level of each would have to be higher than the other: and that's impossible. So the Tarskian will have to declare these utterances pathological. And yet, if Joe and Mary have each uttered just one truth about the other, we would want to say that each utterance is straightforwardly false, and not pathological.

Fifth, suppose a Tarskian tells you about her hierarchical account of English. She tells you that the English truth predicate is split into an infinite hierarchy of truth predicates. Now, it seems, you have all you need to form the predicate: 'is true at any level of the hierarchy'. Since she can talk about the hierarchy, so can you. And the predicate itself seems perfectly coherent: for example, you might apply it to the sentence 'Snow is white'. Surely this sentence is true at every level of the hierarchy. But now you

<sup>2</sup> Kurt Gödel, "Russell's mathematical logic", in *The Philosophy of Bertrand Russell*, edited by P. A. Schilpp, La Salle, Open Court 1944, p. 149

have all the resources you need to confront the Tarskian with a new version of the Liar, expressed in the very terms of her solution. For consider the sentence:

This sentence is not true at any level of the hierarchy. If this sentence is true, then, on the Tarskian account, it will be true at some level of the hierarchy. But it says that it isn't true at this level of the hierarchy (or any other level). So suppose the sentence is not true at any level of the hierarchy. Then what it says is true; that is, true at some level of the hierarchy. And again we are landed in contradiction.

Here, the Tarskian bumps up against a feature of natural language that Tarski calls the universality of natural language. Tarski says this:

A characteristic feature of colloquial language (in contrast to various scientific languages) is its universality. ... it could be claimed that 'if we can speak meaningfully about anything at all, we can also speak about it in colloquial language.'

In particular, natural languages are what Tarski calls semantically universal: they can talk about their own semantics. When the Tarskian tells you about the semantics of English, everything she says is expressible in English. And so English has the capacity to express a version of the Liar in the very terms of the Tarskian's purported solution.

#### IV. A Singularity Proposal

I want to suggest an alternative to the hierarchical account. Let's return to my pathological sentence (L) on the board. We have seen that the occurrence of 'true' in (L) is restricted. At least (L) is excluded from its extension, since (L) is not true in the original context. Are there are other restrictions, and, if so, what are they?

At this point, let me articulate what I take to be a basic intuition about our uses of 'true'. Intuitively, we take our uses of 'true' to apply to all things that are true. An ordinary use of 'true' catches in its net anything which is a truth, just as a use of the predicate 'human' catches in its net anything which is a human. When we go on to reflect upon the Liar, we may come to see that this basic intuition requires revision. I have argued that this intuition does require some revision; for example, the occurrence of 'true' in our Liar sentence (L) is restricted. But the more restrictions we place on occurrences of 'true', the more we are at odds with our basic intuition. And the more we are at odds with this intuition, the less plausible the solution. This was one problem with the Tarskian approach.

This suggests a pragmatic principle of interpretation, a principle I shall call Minimality. Minimality says: keep restrictions on occurrences of 'true' (and 'false') to a minimum. Within a given context, if a sentence may be counted as true, then we are to count it as true. We give a sentence truth conditions if we can, and declare semantic pathology only when we have to.

For example, suppose you say "Snow is white" is true". What should be included in the extension of 'true' in your utterance? By Minimality, we include not only truths of which you are aware, but those of which you're quite unaware. For example, truths of chemical theory will be in there. And what about my pathological utterance L? By Minimality, that will be included too. And this is surely plausible. As we've seen, within the original context of utterance, (L) is pathological; but outside that context, (L) is true. Since your utterance is quite unrelated to mine, there is no reason to interpret your utterance as in some way pathologically linked to mine. It would be a mistake to interpret your utterance as implicated in paradox where there is no need.

In general, then, Minimality says that we should place minimal restrictions on occurrences of 'true'. To put things the other way around, the evaluative reach of a given context is to be maximally comprehensive. In this way, we preserve as far as possible our basic intuition about our uses of 'true'.

Sometimes, of course, we are forced to restrict 'true'. We must, for example, restrict the occurrence of 'true' in (L): (L) itself must be excluded from its extension. But Minimality requires us to exclude from a given occurrence of 'true' only those pathological sentences which cannot be included.

So my proposal locates what I shall call singularities of truth and falsity. For example, the sentence (L) is a singularity of the concepts of truth and falsity. Attempts to evaluate (L) within its context fail. In general, a semantically pathological sentence is, in its context of utterance, a singularity of truth and falsity.

Such sentences are singularities only in a context-relative way - there is an appropriate reflexive context in which the singularity is in the extension of 'true' or the extension of 'false'. In the case of our Strengthened Liar discourse, we can, from our subsequent reflective context, evaluate (L) as true. It is only within its context of utterance that (L) is a singularity of truth and falsity.

When we fix extensions of occurrences of 'true' in accordance with Minimality, we fix those extensions in a way that does not depend on what the speaker, or we as

interpreters, know. Such ignorance does not affect the general way in which extensions are determined by Minimality. The relevant empirical and semantical fact may not be available to the speaker or interpreter; nevertheless, these facts in part determine the extension of 'true'.

We are now in a position to see the strongly anti-Tarskian nature of the singularity proposal. The singularity proposal is non-hierarchical. This is a consequence of Minimality. There is no hierarchy of language-levels, but just one language (English), and a single, context-sensitive truth-predicate. Consider again our Strengthened Liar discourse. As we saw, on a Tarskian the shift from (L) to our reflection (R) would involve a shift in language level. But according to the singularity proposal, there are no such levels. Indeed, by Minimality, (R) is not excluded from the extension of 'true' in (L). (R) is a truth in any context: in any context, it is true that (L) is true in (R)'s context of utterance. In particular, (R) is a truth in the context of (L). This means that in the extension of 'true,' there is a sentence in which the 'true<sub>R</sub>' predicate appears. For the Tarskian, this would amount to an unacceptable mixing of language levels. On the singularity approach, there are no such language levels to be mixed.

It seems to me an unwarranted Tarskian prejudice to deny that (R) can be in extension of 'true<sub>L</sub>'. Perhaps this prejudice is encouraged by consideration of the utterer's epistemic situation. As I set up the original case, in my original context of utterance of (L), I do not know or believe that (L) is pathological, and so (R) does not express anything I know or believe. I certainly may come to know or believe (R), but only by transcending the epistemic situation I am in when I utter (L). But as we have already said, what is true or false in a given situation is not determined by what speakers know. Moreover, the Tarskian cannot really help herself to this epistemic justification of the levels. For it would be quite possible for me to produce intentionally a pathological token on the board, in full knowledge that I will go on to reflect on that utterance, qua pathological. We may produce pathological utterances in full knowledge that in another (reflective) context, these utterances will be true, because pathological in the present context. In this case, there is no shift in what is known, yet still the Tarskian will discern a shift of language level.

Gödel makes the following suggestion regarding the paradoxes:

It might even turn out that it is possible to assume every concept to be significant everywhere except for certain 'singular points' or 'limiting points', so that the paradoxes would appear as something analogous to dividing by zero. Such a

system would be most satisfying in the following respect: our logical intuitions would then remain correct up to certain minor corrections, i.e. they could then be considered to give an essentially correct, only somewhat 'blurred' picture of the real state of affairs.<sup>3</sup>

There are some disanalogies between my singularity proposal and Gödel's suggestion. First, according to Gödel, Liar sentences are meaningless; on the singularity proposal, pathological sentences are meaningful, and ultimately true or false. Second, though Gödel does not expand on the passage quoted in the text, on the natural way of developing Gödel's suggestion, it is semantical predicate types simpliciter that have singularities: semantical concepts have standing singularities, so that there are systematic errors in our understanding of semantical concepts. But on my proposal, minimal restrictions are placed on semantical predicates in a context. We make a minimal mistake about what the concept word picks out on certain occasion of use -- singularities of truth and falsity shift with the context. Our mistake is relative to the context of utterance: we do not, in general, go wrong in a principled way, about standing singularities of semantical concepts. Indeed, as we have seen, on my proposal statements that are pathological in their context of utterance fall under the concept of truth or falsity in a suitably reflective context.

Despite these disanalogies, I take my singularity proposal to be very much in the spirit of Gödel's remarks. And we can claim for it the same satisfying feature: our logical intuitions about 'true' and 'false' are almost correct. Our intuition that every sentence is true or false is essentially correct. It is only in pathological or paradoxical contexts that we mistakenly suppose that certain statements are true or false, when they are neither -- and in such cases our applications of 'true' and 'false' require only minimal corrections. And, ultimately, bivalence is upheld: any sentence is true or false, in an appropriate reflective context.

A second intuition that requires revision is that 'true' is a predicate constant. Strengthened Liar discourses suggest that 'true' shifts its extension according to context. But these shifts are kept to a minimum. In correcting both these intuitions, we retain a single truth predicate which undergoes minimal changes in its extension-according to context. There is no wholesale revision of the notion of truth (or falsity); no division of


<sup>3</sup> Kurt Gödel, *op.cit.*, p. 150.

'true' into infinitely many distinct predicates; no splitting of everyday English into an infinite hierarchy of languages; and no absolute rejection of bivalence.


V. A glimpse of the formal theory

We can catch a glimpse of the formal theory by way of some examples. Consider the following chain of sentences:


- A. Snow is white.
  - B. A is true.
  - C. B is true and Quayle is not President.
- This chain of sentences can be represented in the obvious way by a tree as follows:



Notice that this tree is finite; that is to say, all of its branches are finite. We can go on to obtain what I shall call the evaluated tree. We have sentences in the tree that do not themselves involve the truth predicate. We assign these the obvious truth values. On the basis of these assignments, we assign truth values to the other sentences of the tree according to obvious rules about truth, negation and conjunction. We obtain the following evaluated tree:



According to this tree, the sentence (C) is assigned the value T, as we would expect. Let us now consider the tree for our Liar sentence (L). We have to introduce a further ingredient into our analysis. As we have seen, where the Liar is concerned, the context of utterance matters. In the previous example I treated sentences simply as sentence types. Now we must be more precise. From now on, we shall treat sentences as sentence types in a context. In particular, the Liar sentence (L) will be represented as an ordered pair: the first member is the sentence type '(L) is not true' (we simplify the denoting phrase), and the second member is the context of (L). The tree for (L) looks like this:



This tree has a single branch, and it is infinite. Pathology is indicated by the infinite branch, on which (L) repeats. Intuitively, the sentence (L) is not grounded in any sentence free of the truth predicate. This is in contrast with the grounded sentence (C); each branch of (C)'s tree terminates in a sentence free of semantical predicates.

At this point we are in a position to identify a singularity of 'true' in (L). We need a little terminology. The sentence T(L) is an evaluating sentence of (L)'s tree, since it appears on the tree, and is of the form T(φ); (L) is the corresponding evaluated sentence of the tree. Now we can give a general (rough rough) characterization of a singularity of an occurrence of 'true' in a sentence σ:

**Definition** The first evaluated sentence on an infinite branch of σ's tree is a singularity of 'true' in σ.

Here, then, (L) is a singularity of 'true' in (L).

Having identified this singularity, we may proceed to obtain the evaluated tree for (L). First, we pair the truth value F with the evaluating sentence associated with the singularity, and terminate the branch at this node. We obtain:

$\langle L, \text{Context}_L \rangle$   
 |  
 $\langle \langle T(L), \text{Context}_L \rangle, F \rangle$

This corresponds to the exclusion of the singularity from the extension of 'true' in (L). Then we pair truth values with every sentence of the tree, in the obvious way, to obtain the evaluated tree. The evaluated tree for (L) is:


$\langle \langle L, \text{Context}_L \rangle, T \rangle$   
 |  
 $\langle T(L), \text{Context}_L \rangle, F \rangle$

This yields a truth value for (L) itself. This value coincides with the value (L) receives in an appropriate reflective context; that is, the value true.

Let us now turn to the reflection (R), in our strengthened Liar reasoning. The tree for (R) looks like this:

$\langle R, \text{Context}_R \rangle$   
 |  
 $\langle L, \text{Context}_L \rangle$   
 |  
 $\langle T(L), \text{Context}_L \rangle$   
 |  
 $\langle L, \text{Context}_L \rangle$

(Notice the switch in contextual subscripts.) Like the tree for (L), the tree for (R) has an infinite branch. But there is a crucial difference. (R) does not repeat on its infinite branch. Though (L) continually repeats, (R) does not. We could represent this tree in a different way:



The Liar sentence (L) is caught in a loop: but (R) stands above the loop. In the formal theory, this is sufficient to characterize (R) as a reflection. And since (L) is assigned the value true, and (R) says that (L) is true, (R) is assigned the value true as well.

#### VI. Theory and object language

This then is a brief glimpse of the formal theory. Let me now make some more general remarks about the formal theory. It is a theory about our ordinary notions of truth and falsity. According to the theory, our ordinary semantical predicates are context-sensitive; and so the theory is a theory about context-sensitive usage. But the language of the singularity theory is a classical formal language that does not itself contain context-sensitive terms. There will be, then, some separation of theory and object language.

And so the question naturally arises: what is the relation between the object language and the singularity theory? I have claimed that the singularity theory is strongly non-Tarskian. But this claim will be severely compromised if it turns out that the language of the theory is a Tarskian metalanguage. I've argued elsewhere that many non-Tarskian theories of truth are compromised in just this way. I've argued that the non-Tarskian accounts of Kripke, Herzberger, Gupta, and Feferman cannot dispense with a Tarskian hierarchy.<sup>4</sup> Each of these theories is expressed in a language which is a Tarskian metalanguage for the object language. But things are different with the singularity proposal.

The language of the singularity theory is not a Tarskian metalanguage. Let's see why it isn't. According to the singularity theory, paradox is avoided by the identification and exclusion of singularities. Intuitively, the theory works by saying what is excluded from the extension of 'true', not by saying what is included. Now, no truth of the theory is identified by the theory as a singularity. So, by Minimality, and truth of the theory is in the extension of ordinary occurrences of 'true'. To speak metaphorically, our ordinary uses of 'true' arch over the theory of that usage. Our ordinary uses of 'true' include in their extension all true sentences of the singularity theory. By Minimality, our ordinary uses of 'true' include in their extensions truths of the singularity theory, just as they

<sup>4</sup> Keith Simmons, "The Diagonal Argument and the Liar", *Journal of Philosophical Logic*, 19 (1990), pp. 277-303. The papers by Kripke, Herzberger, Gupta and Feferman are collected in R. L. Martin (op. cit.).

include, say, the truths of chemical theory. This could not be so if the language of the theory was a Tarskian metalanguage.

Now, since the language of the theory is a classical formal language, it is subject to Tarski's theorem. In the usual way, we can generate from this formal language a hierarchy of formal languages, each containing the truth predicate for the preceding language. But of course, none of the truths expressible in these languages are identified as singularities by the singularity theory. And so our ordinary uses of 'true' arch over not only the truths of the singularity theory, but also over all the expressed by the languages of this hierarchy. Truths of any language of the hierarchy are in the extensions of ordinary occurrences of 'true'. There is no stratification of our truth predicate: the singularity theory really is anti-Tarskian.

Let me close with some remarks about universality, in particular, semantic universality. There are at least two aspects of the intuition that a natural language is semantically universal. One is that the language contains a universal truth predicate, one that applies to all the truths. The second is that every sentence of the language can be evaluated, as true or as false, within the language. Let me consider each aspect in turn.

First, we should not expect to find a completely unrestricted, universal truth predicate. Suppose there was such a predicate, call it 'true<sub>U</sub>'. Then the Liar returns. For consider the sentence

This sentence is not true<sub>U</sub>.

When we ask whether this sentence is true or not, we get a contradiction.

According to the singularity proposal, no occurrence of 'true' is universal. Every occurrence of 'true' has singularities. This is readily seen by considering anaphoric cases of the Liar. Suppose you innocently say:

(S) 'Snow is white' is true

But now suppose you add, not-so-innocently: "But this very sentence isn't". Because of the anaphoric construction here, your pathological addition shares its context of utterance with the sentence (S). And so your pathological addition is a singularity of the occurrence of 'true' in the sentence (S).

Such anaphoric additions are always possible. It may be that there are no actual statements uttered that force restrictions on a given occurrence of 'true' or 'false'; there may be no actual singularities. But there are always possible anaphoric back-references,

appropriate pairs of sentence types and contexts. And these will be singularities of the given occurrence of 'true' or 'false'.

So every occurrence of 'true' or 'false' has singularities. No use of the truth predicate is universal. But, on the singularity theory, each occurrence of the predicate is restricted as little as possible. Each occurrence of the truth predicate applies to all truths except its singularities. Its application extends to true reflective evaluations of its singularities; and extends even to truths of the singularity theory and its associated hierarchy. I would suggest that, on the singularity proposal, each occurrence of the truth predicate is as close to universal as it can be without contradiction.

What about the second strand of semantic universality? According to this, every sentence of our language can be evaluated, as true or as false, within the language. I would claim that the singularity theory does full justice to this intuition. In the case of our sentence (L), (L) cannot be evaluated from within its context of utterance. But there are other contexts in which (L) can be evaluated; for example, the context in which we produce the reflection (R). As we've seen, the occurrence of 'true' in (R) has singularities too. And though these singularities cannot be evaluated in (R)'s context, there are other contexts in which they can. In general, for every singularity, there is a sentence of the language that evaluates it. We can even evaluate a 'global' evaluation like:

Every sentence is true or not true

This global evaluation has singularities: the sentence itself is one of them. But still we can reflect on this evaluation, and declare it true. No sentence of our language is beyond the evaluative reach of our semantical concepts.

I am claiming, then, that the singularity theory does justice to Tarski's intuition that natural languages are universal. I'm not suggesting that there is a 'universal' context, a single context that embraces all possible truths and all possible evaluations. Every stretch of semantical discourse has its singularities. But each discourse comes as close to semantic universality as it can. What cannot be expressed within 'one stretch of semantical discourse can always be expressed in another. We can say everything there is to say, but not all at once.