

False Truths

Isidora Stojanovic

ENS “ Institut Jean-Nicod

published in Fernando Martinez Manrique and Luis Miguez Peris-Viné (Eds.), *Actos des V Congreso de la Sociedad de Logica, Metodologia y Filosofia de la Ciencia en Espana*, 2006: pp. 75-80)

One of the most interesting and fruitful applications of logics, classical or other, has been in supplying formal frameworks for the semantics of natural language. In this paper, I discuss the following puzzle: there seem to be arguments that are logically valid - more precisely, that are instances of the rule of universal instantiation, and yet, the utterance of the premise is intuitively true while the conclusion is false. I will discuss two strategies, developed in response to different sorts of problems, that seem immediately applicable to this puzzle. While the so-called contextualist strategy blocks the puzzle at the level of syntax, the index-shifting strategy actually embraces the apparently paradoxical claim that there are logically valid arguments with premises whose utterances are true and a conclusion whose utterance is false, but insists that different points of evaluation come into play in determining the truth values of the utterances involved in the alleged counter-instances to rules of logic.

1. Some Preliminaries

Given any formal framework for natural language semantics, there are two basic assumptions that one ought to make about the framework:

1. Conformity to logic:

Our semantic framework should not assign value False, or 0, to logically valid sentences, that is, sentences derivable in one's preferred deduction system (such as,

say, natural deduction for first order logic). And more generally, if the framework assigns value True, or 1, to each among the sentences S_1, \dots, S_n , and if sentence S logically follows from the set $\{S_1, \dots, S_n\}$, that is, if it is derivable from this set by a finite application of inference rules from the deduction system, then the framework should also assign value True to S.

2. **Conformity to competent speakers' intuitions:**

If the majority of competent speakers are inclined to judge that a given sentence S, on a given occurrence, is true, then the framework had better assign it value True.

Now, of course, both assumptions call for an important qualification. Normally, a semantic framework for natural language does not assign a truth value to sentences on their own. Consider the sentence "He is a philosopher". It would be silly to suppose that our semantics, *even* given a possible world, should decide whether the sentence is true or false. It is crucial to know who is being referred to by the pronoun. Thus if this sentence is uttered in a conversation about, say, John Perry, who is indeed a philosopher, it should be assigned value True with respect to the actual world, and value False with respect to those possible worlds in which Perry is not a philosopher, while if the same sentence is uttered while talking about the Italian singer Toto Cutugno, it should be assigned value False with respect to the actual world, and True with respect to the worlds in which Cutugno happens to be a philosopher.

Since the 70's and the work of David Lewis, David Kaplan and others, it is widely accepted that there are (at least) two parameters that the assignment of truth values to natural language sentences should take into account. Unfortunately, there is not only theoretical disagreement regarding the nature of those two parameters, but considerable terminological divergence as well. Kaplan talks of *contexts* and of *circumstances* (of evaluation), while Lewis calls the first parameter *context* and the second *index*.¹ In the work of Robert Stalnaker, both parameters are *possible worlds*, and what he calls *context* is yet a third thing, namely a certain set of possible worlds (those that are compatible with the conversational background). In his recent work, Stefano Predelli

1 See Kaplan: *Demonstratives*, in Almog et al. (eds.), *Themes from Kaplan*, Oxford UP (1989), and Lewis: *Index, Context and Content* (1980).

expresses concerns, which I share, about calling "context" any among the technical parameters with which the semantic framework operates, because of a heavy load of connotations. He opts for what he takes to be neutral terminology, calling *index* the first parameter (Kaplan's *context*) and *point of evaluation* the second (Kaplan's *circumstance*), so that the objects to which a semantic framework assigns truth values are *sentence-index pairs*, and the truth value assignments are functions of *points of evaluation*.²

Let us use, as usual, the notation $[[\dots]]$ for the semantic value function, that is, for the function that takes an expression, an index and a point of evaluation as inputs, and returns that expression's semantic value as output. In the case of a sentence, the value is either 0 (False) or 1 (True).³ For the sake of simplicity, the assignment function (which assigns values to free variables) will be seen as part of the index parameter. Finally, if i is an index variable (in the meta-language) and w a point of evaluation variable, the two basic assumptions on frameworks for natural language semantics will read as follows:

1. For any i and any w , if $[[S_1]]_{i,w}=1$, $[[S_2]]_{i,w}=1, \dots, [[S_n]]_{i,w}=1$, and if $S_1, \dots, S_n \vdash S$, then $[[S]]_{i,w}=1$.
2. If a natural language sentence s , as uttered in a given context, is judged true by competent speakers (given a suitable description of what the world is like), if sentence S is an adequate translation for s , if index j captures the relevant information about the context of utterance, and if point of evaluation k does so for the (remaining) factual information necessary to determine the truth value of the utterance, then we should have $[[S]]_{j,k}=1$. Similarly, if the sentence is judged false, we should have $[[S]]_{j,k}=0$.

Spelling out some among the preliminaries, as we have just done, already takes us half a way towards the solution to the puzzle that we haven't laid down yet, which is what I am turning to now.

2 See Predelli: *Contexts*, Oxford UP (2005). In fact, Predelli is also reluctant to talking of *sentences* and prefers to talk of *clauses*.

3 For the time being, let us leave it open whether $[[\dots]]$ may be a partial function, that is, undefined for certain inputs.

2. A Failure of Universal Instantiation?

Suppose that, talking of a dinner that took place during some philosophy conference, I say:

1) Everyone was a philosopher.

Since, let's assume, there were no non-philosopher at that dinner, (1) is intuitively true. But now consider the following, as uttered right after (1), hence presumably in the same context:

2) Toto Cutugno was a philosopher.

By the rule of universal instantiation, (2) logically follows from (1). Yet (2) is false.

The puzzle may be laid down more explicitly as follows:

i. "Everyone was a philosopher. Therefore, Toto Cutugno was a philosopher." expresses a logically valid inference. In fact, it expresses a straightforward instance of the rule of universal instantiation: $\forall xFx \mid - F[x/b]$, where F may be any predicate, and b may be any constant.⁴

ii. (1) is intuitively true. Hence, if S_1 , in the formal language, is a translation for (1), and if i_1 and w_1 are respectively the index representing the context relevant for the interpretation of (1) (presumably, the dinner of the philosophy conference) and the point of evaluation relevant to determining the truth of (1) (presumably, the state of affairs at that dinner), then, given assumption 2, we ought to have $[[S_1]]_{i_1, w_1} = 1$.

iii. (2) is intuitively false. Hence, if S_2 , in the formal language, is a translation for (2), and if i_2 and w_2 are respectively the index representing the context relevant for the interpretation of (2) (presumably, that same dinner) and the point of evaluation relevant to determining the truth of (2) (presumably, the state of affairs at that dinner), then, given assumption 2, we ought to have $[[S_2]]_{i_2, w_2} = 0$.

claim. Since it is plausible to assume that (1) and (2) were uttered in the same context, and that their truth depends on the same state of affairs, we have it that i, ii and iii lead

⁴ Let us put aside the question of how the past tense, and tenses in general, ought to be represented. Note that nothing important hinges on this.

to contradiction, given assumption 1.

The obvious ways of dispelling the puzzle would be either to take issue with one of the three clauses, or to take issue with the claim itself that those clauses are inconsistent. It is the latter strategies “ those that do reject none of the clauses, but make room for their mutual compatibility, that are most promising. But let us mention, and discard right from the outset, a strategy that would deny ii: the so-called *literalist* strategy.

A literalist will insist that (1) is actually false, though it may *convey* something true. There are two ways of running the literalist strategy. One would be to take issue with the claim that (1), as uttered at the philosophy conference, at which, by assumption, there were only philosophers, is intuitively true. And it is doubtlessly the case that there can be competent English users whose intuitions on the truth value of (1) might be “literalistfi in this sense, and who would take (1) to be false. However, the considerable literature on quanti er domain restriction appears to agree on this much, that the majority of speakers intuitively do take (1) to be true in the situation at stake. I will take those linguistic data for granted, and leave out of consideration the strategy that will dispute this claim. But the other, better known way of being a literalist is to say that, while accepting that (1) is *intuitively* true, the truth value that *semantics* should assign to (1) ought to be false. In other words, one would deny the idea that speaker's intuitions have a say on the semantic content.⁵ But this is tantamount to rejecting assumption 2, laid out at the outset as a basic and non-negotiable assumption about natural language semantics. For this reason, the literalist strategy cannot be considered as offering a solution to our puzzle.

3. The Contextualist Move

The tradition that goes under the name of (*indexical*) *contextualism* accepts all the three clauses above (with some reserve with respect to i, as we shall see), but disputes the claim that those clauses lead to a contradiction. The contextualist thinks that there is

⁵ Defenders of this sort of literalism include Kent Bach, Emma Borg, Herman Cappelen, Ernie Lepore, Scott Soames, relevance theorists, and so on.

some equivocation in clause i. For him, "Everyone was a philosopher. Therefore, Toto Cutugno was a philosopher." expresses a logically valid inference *only if* quantification is unrestricted, that is, if the formal translations for the two sentences are respectively ' $\forall x(\text{Philosopher}(x))$ ' and ' $\text{Philosopher}(\text{Toto})$.' But this is not, he would say, the correct translation for the utterance in (1). Rather, the correct translation is something like ' $\forall x(\text{Dom}(x) \rightarrow \text{Philosopher}(x))$ ', where Dom is a higher-order variable that takes as its values either predicates or sets, depending on the version adopted. With respect to index i_1 , Dom will presumably receive as its value the predicate 'people at that dinner', or the corresponding set of people. And from this we cannot derive that Toto Cutugno is a philosopher, unless we have the additional premise $\text{Dom}(\text{Toto})$ (w.r. to i_1), which we don't. Hence, there is no contradiction.

Now, I have just depicted the contextualist solution using a very old-fashioned account of natural language quantifiers. Even though this is not crucial, let me move, for the sake of accuracy, to representations in terms of generalized restricted quantifiers. Here is what we would then have instead of the original clause i:

iv. "Everyone was a philosopher. Therefore, Toto Cutugno was a philosopher." expresses a logically valid inference. In fact, it expresses an instance of a valid derivation that uses the rule of universal instantiation and the rule of modus ponens: $[\forall x:Gx]Fx, G[x/\text{Toto}] \mid \text{---} F[x/\text{Toto}]$, where, applied to our example, G is the predicate encoded by '-one' in 'everyone' (presumably, something like 'animate human'), F is 'philosopher'. Note, though, that the premise $G[x/\text{Toto}]$, though not explicitly stated, follows from the lexical knowledge that Toto Cutugno is the name of a person (rather than, say, of a city).

Now, once we reformulate in this way our clause i, the contextualist's point will be that $[\forall x:Gx]Fx$ is not yet the correct translation for (1). Rather, the correct translation is $[\forall x:Gx \wedge \text{Dom}(x)]Fx$. Again, since we don't have the premise that Toto Cutugno belongs to Dom, we don't have a counter-instance to any logically valid inference.

4. Index-shifting

Indexical contextualism, as its name indicates, makes crucial use of the parameter of index (or context) in the semantics that it provides for natural language quantifiers, because the value for variable Dom (which restrict the quantifier's domain) is provided by the index. But the role of providing values for indexicals, whether overt (such as pronouns and demonstratives) or hidden (such as quantifier domains) is only one of the two roles that the index plays in Kaplanian theories. The other role is that of determining the point of evaluation that will establish the truth value of the sentence (relative to that index). To acknowledge this double role of indices points to another way of making the semantics of quantifiers index-sensitive, a way that does not have to stipulate any phonetically unrealized variables in the syntactic form of sentences containing quantifiers. For, if we suppose that the domain of quantification may vary across points of evaluation, then the same sentence containing a quantifier may be true with respect to one point and false with respect to another, even if the two points agree and all the facts about the world; and the truth value of the sentence will depend on the index with respect to which the sentence is evaluated.

Now, there is nothing really new in this move. For, consider a standard Kaplanian theory that does not stipulate a unique domain for all possible worlds. In such a theory, the truth value of sentences containing quantifiers will depend on the index, since the index fixes the possible world. So, for instance, the sentence "Some philosophers are French" is true if evaluated at the index of the actual world, but false when evaluated at worlds that are like the actual world except that they do not contain any French philosophers. Now, if we want to make the same move in order to account for the variability in the truth value of sentences considered as all evaluated at the actual world, then either the world parameter in the point of evaluation is to be thought of as what is sometimes called "small worlds", or as situations, or else, the points of evaluation should contain a distinguished parameter for the domain of quantification.

Going back to our example, the sentence in (1), "Everyone is a philosopher", will be true if evaluated at an index whose domain contains among human beings only philosophers, such as the situation of the dinner of the philosophy conference, and

false if paired with an index whose domain contains some non-philosophers.

We now have all the elements to present the *index-shifting* strategy of dealing with the puzzle. This is the strategy that Predelli (2005) deployed to deal with cases that exhibit the same pattern: the premises are intuitively true, the conclusion is false, and yet, the inference appears to be logically valid. But Predelli's cases are significantly different: he is concerned with situations in which one can truly say "This table is 2m long", yet in which an utterance of "This table is more than 1999 mm long" turns out to be false. Predelli's explanation is that the second utterance shifts the standards of precision from low to high, hence the truth of the two sentences depends on points of evaluation that differ on the standards of precision. However, there is no evidence that Predelli would use the same strategy to deal with cases that involve quantifiers.

The index-shifting strategy, as I suggest applying it to our puzzle, will also dispute the claim that the truth of (1) and the falsity of (2) are inconsistent with the validity of the inference $S_1 \vdash S_2$. But, unlike contextualism, the index-shifting strategy respects surface syntax, and admits $[\forall x:\text{one}(x)]\text{Philosopher}(x)$ and $\text{Philosopher}(\text{Totò})$ as true translations for (1) and (2). What it will deny, though, is that (1) and (2) are really uttered in the same context. Hence $i_1 \neq i_2$, and possibly $w_1 \neq w_2$ as well. The idea is that when (2) is uttered, the mere mention of Totò Cutugno makes the index shift, and brings to salience a larger situation, one that besides the philosophers at the dinner also contains Totò Cutugno, and maybe other Italian singers. And, while we have $[[S_1]]_{i_1, w_1} = 1$, it takes little to see that $[[S_1]]_{i_2, w_2} = 0$. So the falsity of (2) does not contradict the truth of (1), because for a contradiction to take place, the two truth values would need to be established with respect to one and the same index.

There is a lingering worry, though. For, what is the value of $[[S_2]]_{i_1, w_1}$? We cannot say that $[[S_2]]_{i_1, w_1} = 0$, because we then get again a contradiction. So presumably, we should say that $[[S_2]]_{i_1, w_1}$ is undefined. But that is almost as much of a problem. For, we have a true premise that leads to a conclusion that is neither true nor false, which is at odds with assumption 1. Fortunately, though, we can step out of the worry by amending our basic assumption with this proviso:

For any i and any w , if $[[S_1]]_{i,w}=1$, $[[S_2]]_{i,w}=1, \dots, [[S_n]]_{i,w}=1$, if $S_1, \dots, S_n \vdash S$, **and if $[[S]]$ is defined in (i, w)** , then $[[S]]_{i,w}=1$.

5. Conclusion

In this paper, I have discussed a problem for any natural language semantics that preserves logical validity and conforms to speakers' intuitions. The problem lies in cases of logically valid inferences whose premises are uttered truly while the conclusion is uttered falsely. I have discussed two strategies of explaining the problem away. While the contextualist strategy manipulates the syntax of sentences containing quantifiers, the index-shifting strategy manipulates the parameters at which those sentences are evaluated for truth. Unfortunately, I lack the space to compare the two strategies, or discuss their applications to further puzzles of the same pattern 'T \vdash F'.