

An Unintentional Defense of the Indeterminacy of Meaning?

Manfred Kupffer

Abstract Markus Werning attempts to refute Quine's thesis that meaning is indeterminate. To this purpose he employs Hodges' theorem about extensions of cofinal meaning functions. But the theorem does neither suffice to solve Quine's problem nor the problem Werning mistakenly identifies with Quine's. Nevertheless it makes sense to employ the methods used in Werning's paper with regard to Quine's thesis, only that they tell in favour of the thesis instead of against it.

1 Introduction

1.1 Background: Quine

In § 15 of Chap. 2 of "Word and Object", Quine argues for his famous thesis of the indeterminacy of translation. The argument builds on the findings of earlier paragraphs which suggest that the process of radical translation, translation up from scratch of a native language, may only yield determinate translations for a subset of all the sentences of the native language.

We have had our linguist observing native utterances and their circumstances passively, to begin with, and then selectively querying native sentences for assent and dissent under varying circumstances. Let us sum up the possible yield of such methods. (1) Observation sentences can be translated. [...] (2) Truth functions can be translated. (3) Stimulusanalytic sentences can be recognized. [...] (4) Questions of intra-subjective stimulus synonymy of native occasion sentences even of non-observational kind can be settled if raised, but the sentences cannot be translated. (Quine 1960, p. 68)

M. Kupffer (✉)
Institut für Philosophie, Universität Frankfurt, Grüneburgplatz 1, Frankfurt, Germany
e mail: kupffer@em.uni-frankfurt.de

According to Quine, little more than observation sentences¹ possess a determinate translation, and hence, a determinate meaning. For reasons of simplicity, and in accordance with the target paper, we will assume that undisputed cases of a determinate translation are confined to observation sentences, Quine's category (1). If the field linguist goes beyond these by his method of *analytical hypotheses*, the outcome of the enterprise is rather free-floating and not determined by the available data (stimulus/response pairings), witness the conclusion of § 15.

There can be no doubt that rival systems of analytical hypotheses can fit the totality of speech behavior to perfection, and can fit the totality of dispositions to speech behavior as well, and still specify mutually incompatible translations of countless sentences insusceptible of independent control. (Quine 1960, p. 72)

Since observation sentences seem to be susceptible of independent control, “[s]entences insusceptible of independent control” plausibly mean sentences that are not observation sentences. Thus, Quine's problem in § 15 is the problem how, given an interpretation of all the observation sentences, a determinate interpretation for the rest of the sentences of the language could be found. And Quine's skeptical conclusion is that this problem could not be solved.

1.2 Werning

The first quote above continues as follows:

And how does the linguist pass these bounds? In broad outline as follows. He segments heard utterances into conveniently short recurrent parts, and thus compiles a list of native ‘words.’ Various of these he hypothetically equates to English words and phrases, in such a way as to conform to (1) (4). Such are his analytical hypothesis, as I call them. (Quine 1960, p. 68)

In Werning (2004), Markus Werning takes this passage to motivate the following reading of § 15.

The problem Quine faces in § 15 of *Word and Object* can be described as the puzzle of how to construe a cofinal extension of a meaning function that so far only maps observation sentences onto their stimulus meanings. (Werning 2004, p. 152).²

In less technical jargon, Werning sees Quine's problem in § 15 as the one of finding a determinate interpretation of the syntactic parts of observation sentences. This does not seem to be correct; as I have pointed out above Quine's problem is

¹ *Occasion sentences* are sentences “which command assent or dissent only if queried after an appropriate prompting stimulation” (Quine 1960, p. 35f). *Observation sentences* are cases of occasion sentences where assent or dissent do not vary “under the influence of collateral information”, *ibid.*, p. 42.

² Werning admits that there may be more than one source of the alleged indeterminacy. But once the inscrutability of reference and the indeterminacy of grammar are put aside, then, according to Werning, the above problem remains.

concerned with the interpretation of non-observational sentences. On the other hand, the two problems seem to be intimately related; witness Quine's above description of the practice of the field linguist. And indeed Werning thinks a solution to the problem of finding a determinate interpretation of the syntactic parts of observation sentences would imply a partial solution to Quine's problem. Furthermore he claims to have solved the former problem.

Before I can address these claims in any detail, let me, for further reference, restate the two problems in a conceptual framework that is, but for simplifications of minor importance, the one used by Werning.

I begin with two important clarifications. First, in the following "meaning" should not be understood as "stimulus meaning" in Quine's sense, but in the sense Quine seeks to undermine with his arguments. (According to Quine, his notion of stimulus meaning "captures the notion of meaning for observation sentences" (Quine 1986, p. 428), but only for them.) Second, it is important that Werning does not require there to be a unique assignment of meanings to the expressions of a language in order for meaning to be determinate. For him meaning is already determinate if such assignments are determined *up to isomorphism*; now this is the case exactly if all questions of synonymy in the language possess a unique answer;³ let us accept this notion of determinacy for the purposes of the present paper.

For further reference, let us also introduce a bit of (rough and ready) linguistic terminology.⁴ A language L is, in the following, a set of expressions that possesses a grammatical analysis in terms of a set of basic expressions (the vocabulary) and a set of syntactic operations;⁵ if an expression e in L is the result of the application of a syntactic operation to some expressions of L , then these are called the *immediate constituents* of e . Expression a is a constituent (or *part*) of expression b iff a is either b , or an immediate constituent of b or a constituent of a constituent of b . Any part of b that is different from b , will be called a *proper* part. We will also assume that every language L has a designated subset, the set of sentences, and, importantly, that the set of sentences of L is cofinal in L , i.e. that every expression in L is a constituent of some sentence.

Quine distinguishes sentences that are observational like "Lo, a a rabbit" from sentences like "Snow is white" or "Bachelors are unmarried men", which are not. We call any sub-expression of an observation sentence an *observation expression*, with OE being the set of observation expressions; the *observational language* (OL) is the closure of the set of observation expressions under syntactic operations. Clearly, OE is a subset of OL , and the latter is a subset of L . Werning does not make any specific assumptions about whether OE and OL are different, but if they weren't his discussion on pp. 163–164 would be pointless; he explicitly wants to leave it open whether there are terms in L that are not in OE ,⁶ so let us assume for the purpose of our discussion that $OE \neq OL \neq L$.

³ Compare Werning (2004, Theorem 11, p. 172).

⁴ For versions of these notions in a more refined and linguistically adequate framework compare Hodges (2001).

⁵ A syntactic operation is, for some n , a partial function from L^n into L .

⁶ Werning (2004, p. 164).

Given all this we may now put the two problems as follows. (For ease of reference, I call the second problem “Quine’s Problem”; to put it more precise, the formulation below is the result of trying to express Quine’s concerns within the same conceptual framework.)

Werning’s Problem Given an interpretation of the observation sentences, how can we determine which observation expressions are synonymous?

Quine’s Problem Given an interpretation of the observation sentences, how can we determine which sentences of the native language are synonymous?

Now finally, I am able to express the main claims of Werning’s paper in terms of these two problems. Werning claims (i) that a theorem by Wilfrid Hodges shows that if a context principle accepted by Quine holds, then Werning’s Problem possesses a solution; and (ii) that a solution to Werning’s Problem also entails there is a solution to what amounts to a restricted version of Quine’s Problem. Additionally (iii) he claims that these considerations back a certain (neuroscientific) conception of meaning. In the following I will show that Werning’s premises do not suffice to show his claims. Indeed, instead of devising a refutation of Quine’s argument for the indeterminacy of meaning, Werning provides us, unintentionally, with a line of defense for Quine’s thesis.

1.3 Leitgeb

Werning’s solution to Werning’s problem has recently been criticised by Hannes Leitgeb (2005) in this journal. First, Leitgeb doubts that it is consistent with Quine’s intentions to take stimulus meanings to be the meanings of observation sentences. Second, Leitgeb attacks Werning’s particular version of Frege’s context principle. While Leitgeb does not attack (i) above, if successful, his critique would already suffice to establish that Werning fails to achieve his objective to refute Quine by using Quinean premises. Let me therefore briefly say why I find Leitgeb’s critique not compelling.

First, in order to argue that Werning is wrong in equating the meanings of observation sentences with stimulus meanings, Leitgeb is trying to exploit Quine’s distinction between meanings of sentences taken *holophrastically* and of sentences taken as *analysed* in Quine (1986, p. 428). The first variety of meaning, Quine says, might be identified with stimulus meaning, the second not. Leitgeb writes: “Thus, according to Quine, while the stimulus meaning of an observation sentence can be identified with its meaning if the sentence is taken holophrastically, i.e. as a single-word expression ..., *this is no longer possible if the observation sentence is not taken holophrastically*, i.e. if it is taken analysed as being syntactically composed of meaningful parts” (*ibid.*, p. 418; emphasis is mine). The consequence of this interpretation of Quine’s distinction is that Werning cannot equate the meaning of observation sentences with stimulus meanings, simply because he has to presuppose a syntactic analysis of these sentences. Let me point out that Quine’s distinction allows for a different interpretation which does not have that consequence.

According to the interpretation I have in mind meanings of expressions taken as analysed are analogous to *structured meanings* in possible-worlds semantics (cf. Cresswell 1985), while, within the same framework, *intensions* correspond to holophrastic meanings. Now it is of course possible to consider the (unstructured) proposition expressed by a sentence while, at the same time, the syntactic analysis of the sentence is taken into account. If my interpretation is correct, we may likewise apply Hodges' theorem to what Quine calls "holophrastic meanings", while, at the same time, the syntactic analysis of sentences is taken into account.⁷

Second, Leitgeb tries to show (by way of an example) that Werning's version of the context principle yields counterintuitive predictions and can be used to derive contradictions. In this paper, I sketch a straightforward interpretation of such principles that makes them come out consistent and helps to explain away the apparent oddity of some predictions. Such principles should best be understood as providing only relativised notions of synonymy, or so I will argue. This applies to the context principle in the version used by Werning, as well as to the simplified version I will work with in the bulk of the paper, for ease of exposition.⁸

2 Does Werning Solve Werning's Problem?

Suppose you know the interpretation of all the sentences of a language L , so you know exactly which sentences are synonymous but you do not know for other expressions of L whether they are synonymous. In order to find out the latter, you have to know more about L . It turns out that you only need to assume a certain principle that goes back to Frege. Every expression of the language is part of some sentence. So, if the behaviour of expressions in sentences would determine whether they are synonymous, then we could come to know the whole pattern of synonymy in the language, simply by learning the interpretation of sentences. The following principle would do.

Context Principle Any two expressions of L are synonymous in L if, and only if, for every occurrence of either one in some sentence S of L , replacing the occurrence with an occurrence of the other expression results in a sentence of L that is, in L , synonymous to S .

The following case may show how the Context Principle can be applied to determine the relations of synonymy in a language, given some interpretation of its sentences.

Suppose L has three syntactic categories, the set of names $\{a, b\}$, the set of predicates $\{P, Q\}$ and the set of sentences $\{Pa, Pb, Qa, Qb\}$. Every sentence is

⁷ Of course, this application of Hodges' theorem would yield, at best, holophrastic meanings. But since for atomic expressions holophrastic meanings and analysed meanings coincide, if the application of Hodges' theorem were successful, we could then combine the meanings of atomic expressions to form all other meanings of expressions taken as analysed, in the same way structured meanings are composed out of atomic intensions.

⁸ For a discussion of Werning's principle and Leitgeb's objection, see footnotes 9 and 10 below.

the result of the syntactic operation of the concatenation of a predicate and a name, in that order. Suppose we know that the sentences of the language are interpreted by assigning them the two pieces of fruit in the bowl on the cupboard:

Pa: apple

Pb: apple

Qa: pear

Qb: pear

This interpretation, obviously, settles all questions of synonymy in *L* between sentences. But it also may be used to determine the answer to these questions for arbitrary expressions of the language, as follows. *If the Context Principle holds good of L*, we may conclude that *a* and *b* are synonymous in *L* but *P* and *Q* are not, because only the first pair may be exchanged in sentences without changing their interpretation in *L*.

Summing up, *if the Context Principle holds for a language L now our assignment of meanings in L to the sentences of L indeed completely determines the relation of synonymy in L*. That is (roughly) the content of Hodges' theorem; the latter is slightly more general, it replaces the set of sentences of *L* with some *cofinal* set of expressions of *L* and uses a context principle that is modified accordingly. As Werning puts the theorem,

if you [...] have already compiled a meaning function for some subset of the expressions of the language such that every expression of the language is a syntactic part of some element of this subset, then there is one and, up to isomorphism, only one extension of this meaning function to the entire language, provided the language satisfies the Frege-Husserl triad. (Werning, p. 162)

The "Frege-Husserl triad" means a version of the context principle together with the category principle, and the principle of compositionality. The latter two, however, are built into Werning's version of the context principle, cf. Hodges (2001, p. 17). None of the subtle differences of the technical version of Hodges' theorem to my above presentation of it is important for the points I am going to make, so let me return to the simplified presentation.

It is important to note that our above context principle may only be used to determine the relation of synonymy *in a language*, it makes synonymy language-relative. If you want to assume the context principle both for a language *L* and a sub-language *L'* and avoid contradictions, this relativisation is vital, for it is not guaranteed that two expressions that are synonymous in *L'* will stay synonymous in *L*, even if the interpretation of the sentences of *L'* stays the same and *L* and *L'* satisfy the context principle. This is so because the two expressions might fail to be substitutable *salva interpretatione* in some additional sentences. We only have a guarantee that non-trivial synonyms in *L'* stay synonyms in *L* if *L* contains already all sentences of *L'* that consist of *L*-expressions.

At least relativising synonymy is one of two systematic options one has in dealing with the context principle. The other option is to drop the qualification "in

L ” to “synonymous” and then never to apply the resulting context principle unless you are sure that nothing could be added to your language which would break synonymies that are already established. That also implies that, but for a few exceptions, we should refrain from applying the context principle to sub-languages. It is important in this respect that Werning’s version of the context principle is indifferent between our two options above, and might be read either way.⁹

Hodges seems to have had the same two options in mind in the following passage from Hodges (1998, p. 21). (Our above context principle equals roughly the combination of his principles of strong compositionality and full abstraction.)

If you believe in both strong compositionality and, say, full abstraction in sentences, then certain things follow. Given two expressions e and f which have the same meanings in your language E , you have to allow that either (a) there could be a larger language in which these same expressions e and f have different meanings, or (b) there couldn’t. Each possibility leads into problems. The problem for (a) is that one thinks of the meaning of an expression as something about that expression and not about the language as a whole. How could we learn the language if we had to know the meaning of every expression in order to be sure of knowing the meaning of a single expression? The problem for (b) is that it is hard to legislate what might in principle be added to a language. Suppose we had no concept of belief. Then it might simply not occur to us that one might add expressions to the language which distinguish one Sinn from another.

If there could be a larger language F that contains E and in which e and f have different meanings while they are synonymous in E , then, clearly, synonymy and meaning in general have to be understood as language-relative.

I will not compare the *pros* and *cons* of the two options in any detail. But at least the difficulty Hodges raises for the second option is by no means less severe than that for the first one. To say “never apply the context-principle to sub-languages unless you are sure that no additional vocabulary introduces new distinctions in meaning” might be tantamount to the advice never to apply the principle at all, because any language could become a sub-language of the troublesome kind if new concepts emerge during the evolution of the language.

For the purposes of his argument, Werning has to assume a version of the context principle, too. But it would be ad hoc to assume that the principle holds just of this

⁹ The following is Werning’s principle. “[F]or every natural language, if the meaning function ν , defined on a set of grammatical terms Y of the language, is a cofinal extension of the meaning function μ , defined on a set of grammatical terms X of the language, then ν is a cofinal Fregean extension of μ .” Werning (2004, p. 155). The principle is at least closely related to principle F in Hodges (2001, pp. 16–17).

ν and μ above might either be understood as assigning meanings in the unrelativised sense or as assigning meanings *in* Y and *relative* to X . The need for the latter relativisation may be illustrated by Leitgeb’s example, see footnote 10. It may facilitate the discussion of that example to introduce also an informal version of Werning’s principle in its relativised reading.

For any natural language L and subsets X and Y of L such that X is cofinal in Y , two expressions of Y are synonymous in Y relative to X if, and only if, for every occurrence of either one in some expression $e' \in X$, replacing the occurrence with an occurrence of the other expression results in some $e' \in X$ that is synonymous in Y relative to X .

case. If the assumption is based on principled reasons and fit to be applied to sub-languages, one should at least assume that our context principle holds of arbitrary sub-languages of natural languages. (Indeed Werning's own context principle is even stronger, but this means it is stronger than it need be.) Now, as pointed out above, to assume that the context principle holds of arbitrary sub-languages of natural languages would immediately lead to contradictions, if a non-relativised version of the context-principle were employed. Since it wouldn't be charitable to construe Werning as basing his argument on a contradictory premise, let us, in the following, explore the relativised reading of the principle.¹⁰

Let's come back to Werning's Problem, the problem of how to determine which observation expressions are synonymous, given an interpretation of the observation sentences. Now, because synonymy is relativised, we have to be more precise. We have to say what is meant by "synonymous", here. Is it synonymy in *OE*, or is it synonymy in *L*? In this respect it is most important that Werning really aims at Quine's problem. He claims to defend, by the use of Quinean premises, the doctrine of *semantic realism*.

¹⁰ Already Leitgeb (2005) argues that Werning's principle is contradictory. He construes a case where the application of Werning's principle leads to counterintuitive and indeed contradictory predictions. Now above we have seen that contradictory predictions about synonymy may be avoided by recourse to a relativised notion of synonymy. Similarly, if Leitgeb had employed a suitably relativised reading of Werning's principle, no contradiction would have arisen from his counterexample. And if predicted synonymy clashes with intuitions about synonymy, this may be resolved by pointing out that there are different choices of parameters involved.

Leitgeb's counterexample to Werning's principle consists in providing different cofinal subsets *X* of *L* and natural meaning assignments μ for any such *X*, such that (i) for any such sets *X*₁ and *X*₂, sentences *S*₁ and *S*₂, and meaning assignments μ_1 and μ_2 of *X*₁ and *X*₂, respectively, it holds that $\mu_1(S_1) = \mu_2(S_2)$, and (ii) if we apply Werning's principle any two expressions *a* and *b* in *L* are synonymous [in *L* relative to *X*] if, and only if for every occurrence of either one in some expression *e* ∈ *X*, replacing the occurrence with an occurrence of the other expression results in some *e'* ∈ *X*, for details see Leitgeb (2005, p. 415f.). The crucial ingredient is the part of (ii) that is put in brackets, here. If we read Werning's principle unrelativised (as Leitgeb does), the part in brackets has to be deleted; whereas it is vital for our above relativised reading of Werning's principle.

Leitgeb draws two conclusions from this kind of example. In connection with (ii) he complains that the resulting relations of synonymy only reflect "*syntactic issues concerning X*, which is absurd" (*op. cit.*, p. 416). Now Leitgeb's notion of synonymy is non relativised. Once we replace it by the doubly relativised relation above, this complaint can at least no longer be taken to establish the falsehood of Werning's principle. If the complaint is that the resulting relations of synonymy are counterintuitive, Werning could reply that the intuitions pertain to synonymy in *L* relative to the set of sentences of *L*, and not to synonymy in *L* relative to *X*. And if the complaint is that *any* relation of synonymy which reflects syntactic issues concerning *X* is weird, Werning may even agree. E.g. synonymy in *L* relative to some set *X* of the above sort may be weird, but not because Werning's principle gets synonymy in *L* relative to *X* wrong (what is the correct prediction here, anyway?), but because *X* itself is a weird instance. Now in order to show Werning's principle is *false* if understood relativised, it does not suffice to point out that it overgenerates, or that it has weird instances, one has to show it has a *false* instance.

Second Leitgeb derives a contradiction from his example. Let *X*₁ and *X*₂ be two sets *X* as above. Now suppose *S* is a sentence which is in *X*₁ but not in *X*₂. Because of (ii) *S* and arbitrary members of *X*₁ are synonymous in *L* relative to *X*₁, while *S* and arbitrary members of *X*₂ are non synonymous in *L* relative to *X*₂. Because of (i) the meaning arbitrary members of *X*₁ have in *L* relative to *X*₁ equals the meaning that arbitrary members of *X*₂ have in *L* relative to *X*₂. This is not a contradiction, but immediately becomes one if we drop the relativisations. Hence, the relativised reading of Werning's principle is compatible with Leitgeb's example, while the unrelativised reading is shown to be inconsistent by the example.

The paper defends semantic realism against Quine's (1960) sceptical challenge. This is done not by drawing into question any of his premises, but from within his own philosophical set-up. Semantic realism, hereby, is regarded as the assertion that the expressions of natural languages, up to isomorphism (see below), have determinate meanings, which themselves are metaphysically independent from expressions they are neither identical with nor essentially constituted of expressions. (Werning 2004, p. 146)

Since Quine challenges semantic realism wrt. meaning in L , Quine's problem should be understood as: how can we determine which sentences of the native language are synonymous in L ? Therefore, since Werning claims to present an "argument in defense of semantic realism" against Quine's challenge (Werning 2004, p. 162), see also p. 145) his problem should be understood as a problem about synonymy in L likewise.

Unfortunately, if we understand his problem this way, his attempted solution fails. The solution is summarised in the following passage.

In order to apply this result to Quine's project of radical translation, we take X to be the set of all Native observation sentences and construe an observational sub-language that contains only observation sentences and all their syntactic parts such that X is cofinal in the set of expressions of the sub-language. Presupposing the Frege-Husserl triad, Hodges' Theorem assures us that the neuronal or translational meaning function for observation sentences, which the field linguistic should be able to compile by inductive methods, is up to isomorphism uniquely extendable to the entire sub-language. (*ibid.*, p. 162f.)

This is an argument that Hodges' theorem may be used to show that meaning in OE (he calls it the "observational sub-language") is determinate. In other words, if the context principle holds for OE , then the assignment of stimulus meanings to observation sentences yields a determinate relation of synonymy in OE .

But, while Hodges' theorem may be applied to establish that the meaning of observation sentences determines a relation of synonymy in OE , it cannot be applied to establish that they determine synonymy in L . The set of observation sentences is neither the set of all sentences of L nor the set of all sentences of OL ; there are non-observation sentences in OL , namely standing sentences composed of observational vocabulary.

One might still think that, as far as observation expressions are concerned, whether they are synonymous in L already follows from whether they are synonymous in OE . And indeed, if two observation sentences are synonymous in OE , they are also synonymous in L , simply because in both cases to be synonymous comes down to the same, viz. to have the same stimulus meanings.¹¹ Furthermore, if a and b are non-synonymous in OE , they will stay so in L . But if a and b are synonymous in OE it does not follow whether they stay so in L , simply because L might contain additional pairs of sentences that distinguish them. On the other hand it might not. That this is left open shows that the context principle alone does not

¹¹ Thanks to the referees for pointing that out to me.

suffice to establish that there is a determinate relation of synonymy in L among observation expressions, given an interpretation of the observation sentences.

Summing up, on one reading Werning's context principle is contradictory, on the other it does not suffice to show that Werning's problem possesses a solution (in the sense in which it would be relevant for his attempted solution of Quine's problem).

3 Would a Solution of Werning's Problem Help to Solve Quine's?

But suppose Werning's problem were settled. Would this help with Quine's? Werning admits that a solution to his problem does not imply that the project of radical interpretation is completed; maybe OE is only a proper subset of the native language. While Werning himself prefers to stay neutral wrt. the question, he cites Quinean reasons to think that the native language contains no non-observational vocabulary, such that the native language would coincide with OL , the closure of the observational vocabulary under the syntactic operations of the language. Then, in order to solve Quine's Problem it would suffice to show that the solution of Werning's Problem implies a solution of Quine's Problem restricted to OL , i.e. *given an interpretation of the observation sentences, how can we determine which sentences of OL are synonymous in L ?*

Now Werning offers a suggestion how to extend the results he thinks he has established (a determinate relation of synonymy for OE) to the whole of OL . His proposal goes as follows.

Recall that the principle of compositionality warrants that once the meanings of the syntactic parts of a sentence and the syntactic operation by which it has been composed are given, the meaning of the whole sentence is determined. So, the neuronal and translational meaning of all those sentences of the original language should be determinable that are built exclusively from expressions of the observational sub-language. (Werning 2004, p. 163)

The proposal sounds plausible indeed. The principle of compositionality is often expressed as the principle that the meaning of an expression is a function of the meaning of its parts (and their manner of composition). This may be understood as saying that the meanings of the parts *determine*, given the syntactic operation that combines the parts into a whole, the meaning of the whole. And if the latter holds of the observational language, then, indeed, the meanings of observation expressions determine the meanings of all sentences of the observational language, and hence of OL itself. Hence, it seems, Quine's problem is solved, provided Werning's is.

Plausible as it is, the reasoning is flawed. First the output of a solution to Werning's problem does simply not deliver specific meanings for sub-expressions. Werning tries to solve the problem by appeal to principles like the above context principle. These are principles that deliver, at best, determinate synonymies. They do not deliver *the* meanings of expressions. Hence, even if the solution to Werning's problem would work, we would only be in possession of determinate synonymies and there would still be different possibilities as to what are the meanings (in L) of the expressions in OE . And for any two sentences in $OL - OE$ it might turn out that

they receive the same meanings on one such possibility and different meanings on another.

Furthermore, the principle of compositionality *as it is understood in the rest of Werning's paper* (cf. Def. 3 on p. 150), is not the principle that the meanings of the parts *determine*, given the syntactic operation that combines the parts into a whole, the meaning of the whole. There, and in Hodges' work, the principle is a principle about intra-linguistic synonymy. It says that for any two expressions a and b of the same language L , if a and b have the same form, and corresponding sub-expressions are synonymous, then a and b are synonymous as well. Obviously such a principle does not determine the cross-linguistic identity of the meanings of whole expressions. It is consistent with the compositionality principle thus understood that a and b mean the same thing in two interpreted languages L and L' , f is a syntactic operation common to both languages such that $f(a,b)$ is an expression of both L and L' , and yet $f(a,b)$ means different things in L and L' .¹² So, even if a specific meaning assignment for OE were given, the principle of compositionality, as it is understood in the rest of the paper, would still not yield a specific meaning assignment for the whole of OL . Two meaning assignments for OL compatible with the same assignment of meaning (in L) to observation expressions might still differ with respect to the relation of synonymy in L .

Thus, for the above two reasons, even if we assume that OL coincides with the native language, an appeal to compositionality does not ensure that a solution to Werning's problem would entail a solution to Quine's problem.

If we sum up the last two sections it emerges that *Werning fails to show that Quine's problem possesses a solution*.

4 In Defense of Indeterminacy

Of course I only say it does neither *follow* from Werning's premises that Quine's problem possesses a solution nor that Werning's does. This does not preclude that additional premises could serve to fix the arguments. E.g. it would suffice for the purposes of Sect. 2 to assume that no two observation expressions are intersubstitutable *salva interpretatione* in OE . But why should we suppose such a thing?¹³

Or maybe it turns out that every expression of a language L is a syntactic part of some observation sentence, and hence $L = OL = OE$.¹⁴ Unfortunately, even that would not help. The Context Principle still demands that for two expressions to be

¹² For this point, as well as for a notion of compositionality without that feature, see Szabo (2000).

¹³ To assume that every observation expression is non synonymous in OE would mean that there are no stimulus synonymous observation sentences; now this assumption could hardly be termed Quinean. You might still think that for any sentence S , "A says: ' S '" is plausibly an observation sentence. But S is not a syntactic part of "A says: ' S '", only a name for S is, according to the standard account of such contexts.

¹⁴ cf. Müller (1998, p. 274). The author argues that, for any sentences p and q , the conjunction " p and if q , then q " is stimulus synonymous to p . If this is right, then if p is an observation sentence, so is the above conjunction, and any q is part of an observation sentence.

synonymous in *L*, they should be intersubstitutable *salva interpretatione* in all sentences of *L*, and not just intersubstitutable in all *observation* sentences.

Or, finally Werning might want to replace the Context Principle with the following principle.

Verificationist Context Principle Any two *observation* expressions of *L* are synonymous in *L* if, and only if, for every occurrence of either one in some *observation* sentence *S* of *L*, replacing the occurrence with an occurrence of the other expression results in a sentence of *L* that is, in *L*, synonymous to *S*.

If the Verificationist Context Principle holds of the native language, then indeed Werning's Problem would be solved. But does it hold? Not obviously so. At least it does not seem to be a principle that Quine would want to accept. Indeed none of the Quinean reasons in favour of the context principle cited in the text can be applied here. Unlike the Context Principle, the Verificationist Context Principle cannot simply be argued for by appeal to the idea that meaning is use, or semantic holism (cf. Werning 2004, p. 157).

And this, finally, is why Werning's methods not only fail to show that Quinean premises imply that Quine's problem has a solution. These methods could actually be used to argue that the premises imply that the problem has no solution. If I am right, instead of achieving his aim of solving Quine's Problem, Werning really helps to uncover an argument in favour of the indeterminacy of meaning.

Assume, first, that the factual basis of semantics consists *only* of the stimulus meanings of observation sentences; such that any relation of synonymy in *OL* that is consistent with that basis and the Context Principle is admissible. And second, that there is at least one pair of non-observational sentences with the property that (i) neither sentence can be obtained from the other by replacing occurrences of some observation sentence with those of a synonymous observation sentence, and (ii) the pair is itself stimulus-synonymous, i.e. replacing one sentence of the pair for the other in any observation sentence yields a synonymous observation sentence, and (iii) that replacing one for the other in any sentence yields a sentence of *L*. The pair "bachelors are unmarried men" and "bachelors are bachelors" seems to be a plausible candidate, here. It could then be shown that under these conditions, there are both admissible synonymies that treat the pair as being synonymous, and ones that treat the pair as being non-synonymous. But then, for such sentences (and some of their sub-expressions), synonymy in *OL* is indeterminate; but if synonymy is indeterminate, then so is meaning.

5 "Neuronal Meanings"

Werning draws quite a different conclusion from his inquiries.

Quine's naturalist framework, up to isomorphism, uniquely determines a function that maps every expression of a language that may figure as a syntactic part in some observation sentence onto its neuronal meaning, if we

only presume that the language fulfils the principle of compositionality, the context principle, and the category principle. (*ibid.*, p. 169f.)

Above, I have argued against that conclusion. Still, in this part I'd like to comment on the notion of "neuronal meaning". Specifically, I want to point out that even if Quine's Problem had been solved, this would still not allow to conclude, as Werning does, that "[t]he three guidelines Frege and Husserl set up at the beginning of the 20th century can be combined with Quine's naturalistic maxim. When so combined, they lead to a realist attitude towards meanings and advise us to view semantics as a part of neuroscience." (*ibid.*, p. 170)

If Werning had solved Quine's Problem with the above methods, would this warrant any conclusion as to the nature of meaning? None whatsoever! Of course, the observation sentences would have received a definite interpretation, so we can trivially "conclude" semantic facts about them. But the rest of the meanings would still have been only determined *up to isomorphism*; and this kind of determination is very weak; indeed it is slightly misleading to characterise it in a way that suggests that meanings have been determined at all. What has been determined is only their structure, not their content. To say that meaning has been determined up to isomorphism means nothing more than that all questions of synonymy in the language have been settled.

If all these questions have been settled you might take the resultant relation of synonymy as a license to assign any kind of meanings you like to sub-sentential expressions, if that interpretation only respects the relation. If the fruit are enough, and you adapt the interpretation of syntactic operations accordingly, you could assign them other fruit from the bowl on my cupboard. You could take a set of lice to be the realm of meanings for the rest of your language. You could even take neuronal states. But any such pick would be completely arbitrary. It would not justify any conclusions about the nature of meaning. Even if you pick neuronal states, it is wrong to say that "a function that maps every [...] syntactic part in some observation sentence onto its neuronal meaning" is uniquely determined (up to isomorphism); simply because there are lots of other kinds of meaning functions that could have been picked as well.

You might have other reasons for choosing some specific kind of meaning function, like e.g. an assignment of neuronal states. Unfortunately it is not even clear to me whether that would be consistent with the context principle. Sure, in the form presented above and in Hodges' work, the principle only states a functional dependence of sub-expression synonymy on sentence synonymy and does not rule out that any specific meaning function that displays the correct dependencies may be the "real" one. Things become different when we ask why such a functional dependence should hold at all. The usual answer is, because there is an ontological dependence between the two realms.¹⁵ E.g. one might say that the functional dependence holds because *the meaning of a sub-expression is nothing over and*

¹⁵ E.g. arguments that start from the thesis that meaning is use are arguments in favour of ontological dependence. An ontological dependence of sub expression meanings on sentence meanings is also expressed in the quote from Quine (1961, p. 39). Werning cites on p. 157f. Frege seems to have held the ontological version of the principle, see Reck (1997).

above its contribution to the sentences it occurs in. But if the meaning of a sub-expression is nothing over and above its contribution to the sentences it occurs in it should not be identified with anything that could not plausibly be called such a contribution, like, e.g. a neuronal state. Contributions, if we want to reify them at all, seem to be something like functions from linguistic contexts into sentence meanings.

And what makes Werning think that *stimulus meanings* are neuronal states, anyway? The stimulus meaning of a sentence is defined to be the pair of its affirmative stimulus meaning and its negative one. Now Quine defines the affirmative stimulus meaning as “the class of all stimulations (hence evolving ocular irradiation patterns between properly timed blindfoldings) that would prompt his assent” (Quine 1960, p. 32); negative stimulus meanings are defined analogously. So stimulus meanings are behavioural dispositions, functional states of their bearers. It is at least not obvious that such states are identical to “states of neuronal receptors” (Werning 2004, p. 153). Why, as far as stimulus meanings are defined, individuals could contain no neurons at all, and some other kind of hardwiring, say computer chips, instead and yet have the same stimulus meanings we have.¹⁶

The neuro-scientific conception of meaning seems to be rather desperately in need of noble ancestry if its proponents have to borrow the authority of behaviourists like Quine.

Acknowledgments Many thanks to Markus Werning for generous help with a paper directed against an argument of his own, to Tim Kraft and Tobias Klauk for important hints, and to two anonymous referees for numerous corrections and helpful suggestions that led to substantive improvements.

References

- Cresswell, M. J. (1985). *Structured meanings*. Cambridge (Mass.): MIT Press.
- Hodges, W. (1998). Compositionality is not the problem. *Logic and Logical Philosophy*, 6, 7–33.
- Hodges, W. (2001). Formal features of compositionality. *Journal of Logic, Language and Information*, 10, 7–28.
- Leitgeb, H. (2005). Hodges’ theorem does not account for determinacy of translation. A reply to Werning. *Erkenntnis*, 62, 411–425.
- Müller, O. (1998). *Synonymie und Analytizität. Zwei sinnvolle Begriffe*. Schöningh: Paderborn.
- Quine, W. V. O. (1960). *Word and object*, Cambridge (Mass.): MIT Press.
- Quine, W. V. O. (1961). Two dogmas of empiricism. In *From a logical point of view* (2nd ed., pp. 20–46). Cambridge (Mass.): Harvard University Press.
- Quine, W. V. O. (1986). Reply to Hilary Putnam. In L. E. Hahn & P. A. Schilpp (Eds.), *The philosophy of W. V. Quine* (pp. 427–431). La Salle: Open Court.
- Reck, E. H. (1997). Frege’s influence on Wittgenstein: Reversing metaphysics via the context principle. In W.W. Tait (Ed.), *Early analytic philosophy* (pp. 123–185). Chicago: Open Court.
- Szabó, Z. G. (2000). Compositionality as supervenience. *Linguistics and Philosophy*, 23, 475–505.
- Werning, M. (2004). Compositionality, context, categories and the indeterminacy of translation. *Erkenntnis*, 60, 145–178.

¹⁶ Quine explicitly refuses to take the internal hardwiring into account for the purpose of the definition of stimulus meaning “A visual stimulation is perhaps best identified, for present purposes, with the pattern of chromatic irradiation of the eye. To look deep into the subject’s head would be inappropriate even if feasible, for we want to keep clear of his idiosyncratic neural routings or private history of habit formation. We are after his socially inculcated linguistic usage, hence his responses to conditions normally subject to social assessment.” (Quine 1960, p. 31).