

Making sense of (in)determinate truth: the semantics of free variables

John Cantwell^{1,2} 

Published online: 18 September 2017

© The Author(s) 2017. This article is an open access publication

Abstract It is argued that truth value of a sentence containing free variables in a context of use (or the truth value of the proposition it expresses in a context of use), just as the reference of the free variables concerned, depends on the assumptions and posits given by the context. However, context may under-determine the reference of a free variable and the truth value of sentences in which it occurs. It is argued that in such cases a free variable has indeterminate reference and a sentence in which it occurs may have indeterminate truth value. On letting, say, x be such that $x^2 = 4$, the sentence ‘Either $x = 2$ or $x = -2$ ’ is true but the sentence ‘ $x = 2$ ’ has an indeterminate truth value: it is determinate that the variable x refers to either 2 or -2 , but it is indeterminate which of the two it refers to, as a result ‘ $x = 2$ ’ has a truth value but its truth value is indeterminate. The semantic indeterminacy is analysed in a ‘radically’ *supervaluational* (or *plurivaluational*) semantic framework closely analogous to the treatment of vagueness in McGee and McLaughlin (South J Philos 33:203–251, 1994, Linguist Philos 27:123–136, 2004) and Smith (Vagueness and degrees of truth, Oxford University Press, Oxford, 2008), which saves bivalence, the T-schema and the truth-functional analysis of the boolean connectives. It is shown that on such an analysis the modality ‘determinately’ is quite clearly not an epistemic modality, avoiding a potential objection raised by Williamson (Vagueness, Routledge, London, 1994) against such ‘radically’ supervaluational treatments of vagueness, and that determinate truth (rather than truth *simpliciter*) is the semantic value preserved in classically valid arguments. The analysis is contrasted with the epistemicist proposal of Breckenridge and Magidor (Philos Stud 158:377–400,

✉ John Cantwell
jcantwell68@gmail.com

¹ Division of Philosophy, Royal Institute of Technology (KTH), Brinellvägen 32, SE-100 44 Stockholm, Sweden

² Swedish Collegium for Advanced Study, Thunbergsvägen 2, SE-752 38 Uppsala, Sweden

2012) which implies that (in the given context) ‘ $x = 2$ ’ has a determinate but unknowable truth value.

Keywords Indeterminate truth · Determinate truth · Indeterminacy · Free variables · Supervaluations · Reference · Supertruth · Arbitrary reference · Arbitrary objects · Plurivaluations

1 The problem

Let x be a whole number such that $x^2 = 4$. Given some elementary arithmetics, it follows that

$$\text{Either } x = 2 \text{ or } x = -2. \quad (1)$$

That is: (1) is in the given context properly assertable. Indeed, letting x be such that $x^2 = 4$, it seems reasonable to say that (1) is *true*. By contrast, the following sentence would *not* be properly assertable:

$$x = 3. \quad (2)$$

Indeed, it seems reasonable to say that (2) is *false* in the present context, and that its negation, $x \neq 3$, is *true*.

But now consider the sentence:

$$x = 2. \quad (3)$$

Clearly the context does not admit that $x = 2$ be asserted; for it does not follow from the posit $x^2 = 4$ that $x = 2$. However, it is less clear what we should say of the truth value of $x = 2$; is it, given the posit $x^2 = 4$, *true* (like 1) or *false* (like 2)? This is perhaps seems a strange question, but *prima facie* it should be perfectly legitimate. For, first, (3) like (1) and (2) involves the use of a well-formed (open) sentence of mathematical English, and so can be expected to express a proposition with a truth value.¹ Second, the assertability of (1) is most naturally explained by appealing to the fact that it is *entailed* by the posit $x^2 = 4$, while the unassertability of (2) and (3) is most naturally explained by appealing to the fact that they are *not* entailed by the posit. Entailment, traditionally, has been characterised as the necessary preservation of truth, so it is not *necessarily* the case that $x = 2$ is true when $x^2 = 4$ is true, but the very fact that one can ask whether $x = 2$ follows from $x^2 = 4$ suggests that it makes sense to ask for the truth value of $x = 2$ in this context, just as much as it makes sense to ask for the truth value of (1).

So raising the question of the truth value of $x = 2$ seems legitimate. Clearly we are not in a position to claim that $x = 2$ is *true*, so maybe we should say that it is *false*? But this will not do. For if the sentence $x = 2$ is false, its negated counterpart

¹ Throughout I will assume that *the truth value of the proposition expressed by a sentence* in a given context is the same as *the truth value of the sentence* relative to that context, so that it makes sense to ask for the truth value of a sentence in a context.

$x \neq 2$ (or $\neg(x = 2)$) is true, yet the sentence $x \neq 2$ is just as unassertable as the sentence $x = 2$.² Given that the negation of a sentence is true if the sentence is false, we are not in a position to claim that $x = 2$ is false. Furthermore, if $x = 2$ is false then, presumably, $x = -2$ is false as well, but their disjunction (1) is true; so we would have a true disjunction where both disjuncts are false.

So maybe $x = 2$ lacks truth value. This is a tempting explanation, however, this too would require a revisionary (non-classical) semantics for either the truth-predicate, the falsity-predicate, for negation or disjunction. For if $x = 2$ lacks truth value, then it is correct to say that it is *not true*, and if $x = 2$ is not true, it is, on standard classical semantics, *false*. Yet we have already seen that $x = 2$ cannot properly be judged false. Furthermore, if $x = 2$ lacks truth value then, presumably, $x = -2$ also lacks truth value, but their disjunction (1) is true; so we would have a true disjunction where both disjuncts lack truth value.

Perhaps, then, the problem is being approached the wrong way. Perhaps $x = 2$ is unassertable because it doesn't express a proposition (in the sense of having truth-conditions, or being truth-evaluable): $x = 2$ is not the *kind* of sentence that can be asserted because it is not the kind of sentence that expresses a proposition. Likewise for its negation $x \neq 2$. The idea being that, as $x = 2$ is not the kind of sentence that expresses a proposition it would be a category mistake to ask for its truth value, much like asking for the truth value of the weather.

Now, the sentence $x = 2$ (or at least its equally unassertable sibling 'x is equal to 2') is a grammatically well-formed sentence in English. So why would it not express a proposition? One explanation draws on a purely syntactic feature: the sentence $x = 2$ contains the term x and this is a *variable* which is not bound by any quantifier and so is *free*. Such open sentences are in formal semantic treatments often viewed as having a different kind of semantic status than closed sentences (where all variables are bound by some quantifier). So here is a suggestion: sentences with free variables (open sentences) are not the kinds of things that express propositions, and so are not the kind of things that can be properly asserted.

But if sentences with free variables do not express propositions, how come that *either* $x = 2$ *or* $x = -2$ is properly assertable? It also contains a free variable, indeed the same variable. Just read any work on mathematics, physics or any other more formal discipline from Euclid onwards and you will find that it is standard practice to assertively utter sentences with free variables: it is just plain obvious that utterances of sentences with free variables are used to make assertions. So *some* sentences with free variables are properly assertable; which means that the fact that $x = 2$ contains a free variable won't by itself explain why it is not assertable.

So here is a possible solution. Standard lore holds that when a sentence with free variables is used, it is used *elliptically*: it is a truncated sentence with an implicit universal quantifier. This would explain why both $x = 2$ and $x = -2$ are

² By convention I take the phrase 'the sentence $x = 2$ ' to *mention* (not use) the sentence ' $x = 2$ '. So 'The sentence $x = 2$ is true' is just a stylistic variant of "' $x = 2$ ' is true" or "The sentence ' $x = 2$ ' is true". Sometimes, when it is very clear from context, I will just use the phrase ' $x = 2$ is true' to mean 'The sentence $x = 2$ is true', but in later sections when the object-language/meta-language distinction gets more involved I will be more careful in the usage.

unassertable, for both their universally quantified counterparts $\forall x(x = 2)$ and $\forall x(x = -2)$ are *false*. But standard lore is wrong. For (1) was assertable, yet its universally quantified counterpart $\forall x(x = 2 \text{ or } x = -2)$ was not (as it was false). So we cannot explain the unassertability of $x = 2$ by appeal to an implicit quantifier.

What does hold, of course, is the following (where \supset is the material conditional):

$$\forall x(x^2 = 4 \supset x = 2 \text{ or } x = -2).$$

We can treat the initial posit that $x^2 = 4$ as the antecedent of a conditional with (1) as its consequent and then add a universal quantifier to the whole construct to bind the variable x . Quite generally:

B is properly assertable/true within a context in which all that has been assumed about x is that A_1, \dots, A_n if and only if $\forall x(A_1 \& \dots \& A_n \supset B)$ is properly assertable/true (in a context in which no assumptions on x have been made).

This usage is clearly explicated in Gentzen's (1969) system of natural deduction, encoding inferential practices that have been in place since antiquity. Two inferential rules in particular come into play. First of all, if a sentence *B* can be properly asserted within the context of an assumption *A*, then the conditional 'If *A*, then *B*' can be properly asserted. Second, if there is warrant for the assertion of a sentence with a free variable $C(x)$ and no assumptions about x are in force, this warrants assertion of the universally quantified sentence ' $\forall x C(x)$ '. We thus get a systematic way of turning the proper use of sentences involving free variables within the context of an assumption into context-free use of sentences involving conditionals and universal quantifiers.

What this suggests is that we should not ultimately be worried that the use of free variables or assumptions should involve some semantic mystery or weird metaphysics or that the logic of such sentences will contain some surprise.³ But this does not 'explain away' the use of sentences with free variables. The whole point is that sentences with free variables *are* used to make assertions and that the propriety of making such assertions depends on what has been previously assumed.⁴

³ I take it that this is the point Quine makes in his *Variables explained away* (1960). It should be noted, however, that he in the same paper also 'explains away' the use of singular terms in general (including proper names and definite descriptions): "There cease to be singular terms at all" (p. 347). This only illustrates that 'explaining away' the use of a linguistic expression is not the same as saying that it is pointless to ask for its semantic value in the contexts in which it is actually used (even though one could have used some other expression instead). For clearly English sentences containing 'ordinary' singular terms (proper names, definite descriptions) have a semantics in their own right, even if such sentences would be translatable into a (hypothetical) language that contained no such expressions. The same for sentences containing free variables.

⁴ One should add that *inferentialists* like Gentzen, Dummett and Prawitz, far from taking the universal quantifier to explain the use of free variables, instead take the introduction rule for the universal quantifier to be the determinant of its meaning: what holds for an *arbitrary* object with certain properties holds for *all* objects with those properties—it's our use of sentences with free variables that is to account for the meaning of the universal quantifier, not the other way around. However, one need not take such a position on the explanatory order of things in order to see that we have a complex interplay between different kinds of language use and that this interplay has a semantic component.

The practice of positing, of making assumptions, of hypothesising, etc., is a rich linguistic practice in its own right that cannot be explained by grammatical sentence forming rules alone. What systems of natural deduction have established is that the practice can be characterised with formal rigour. But in so doing we have left any pretence that the assertability of (1) and the unassertability of (3) can be traced to some simple syntactic feature of the sentences. So to the extent that a goal of semantic theory is to explain the propriety of use of sentences by assigning truth conditions to sentences, we need a semantic explanation of why some sentences with free variables are properly assertable in some contexts but not in others.

I shall proceed on the assumption that a sentence such as (1) can express a true proposition on its 'literal' reading. Indeed even the sentence $x = 2$ can be used to express a true proposition. Let x be a positive number such that $x^2 = 4$; we can now properly assert that $x = 2$, and we have thereby made a true assertion. The sentence $x = 2$ can also be used to express a false proposition: just let x be a negative whole number such that $x^2 = 4$. Whether (1), (2), (3) or any other non-tautological sentence with free variables expresses a true or false proposition depends on context; if the context only contains the right kind of assumptions about x , then sentences containing unbound instances of x can express true or false propositions.

So far the concern has been *truth*. But there is a closely related issue concerning *reference* or *denotation*. Syntactically a variable plays the role of a singular term, like a proper name or a definite description. Singular terms might or might not refer to or denote some object. Variables, obviously, can denote different objects in different contexts, if they denote at all. So: what object does 'x' refer to in the context in which it has been assumed that $x^2 = 4$? As the sentence ' $x = 3$ ' is false in such a context we can safely conclude that 'x' does not refer to the number 3. But does it or doesn't it denote the number 2? Or the number -2 ? Should we say that 'x' denotes *nothing*, that it's value is undefined, that it has no value? Or that 'x' denotes *both 2 and -2* ? Or perhaps that it denotes the set $\{2, -2\}$? Or should we say that it makes *no sense* to ask what 'x' denotes? All these answers must accommodate the fact that an apparently reasonable answer to the question *What number does 'x' denote?* is: *Either 2 or -2* , as witnessed by the fact that either $x = 2$ or $x = -2$. So the question apparently both makes sense and has an answer. If 'x' denotes nothing, how come that it denotes either 2 or -2 ? If it denotes *both 2 and -2* , how do we avoid a sleuth of troublesome inferences like the one from $x = x$ to $2 = -2$? The variable x clearly doesn't denote the set $\{2, -2\}$, as we know that $x^2 = 4$ whereas the construct $\{2, -2\}^2 = 4$, if syntactically well-formed, is just plain false. There are thus a variety of answers that simply do not seem to fit the bill.

2 Indeterminacy: outline of an analysis

The thesis that will be argued for in this paper is that, given only that $x^2 = 4$, the truth value of ' $x = 2$ ' is *indeterminate*. It doesn't *lack* a truth value or have *many* truth values. ' $x = 2$ ' has one and only one truth value (giving us bivalence), but it is indeterminate whether it is true or false: ' $x = 2$ ' is neither determinately true nor

determinately false. There is a reason for this. The truth value of ' $x = 2$ ' is indeterminate because the reference of ' x ' is indeterminate. The variable x definitely denotes either the number 2 or the number -2 ; so it definitely does not denote the number 3, but it is indeterminate which of the numbers 2 and -2 that it denotes. Accordingly ' $\text{Either } x = 2 \text{ or } x = -2$ ' is determinately true, and ' $x = 3$ ' is determinately false, while both ' $x = 2$ ' and ' $x = -2$ ' lack a determinate truth value.

What of these mysterious properties (*in*)*determinate truth* and (*in*)*determinate reference*? The source of the indeterminacy should at least not be any mystery. The posit that $x^2 = 4$ constrains the possible values of ' x ' but fails to uniquely determine a particular value. The indeterminacy lies in the relation between the variable x and its possible referents and spreads to an indeterminacy of truth value of the sentences in which x occurs. There is no suggestion that the indeterminacy of the truth value of ' $x = 2$ ' is due to an indeterminacy in the identity relation or its extension nor is there any suggestion that the variable x or the objects involved (2 and -2) are in any way indeterminate,⁵ it's the referential relation between x and the involved objects that is indeterminate. The indeterminacy in question is thus *purely semantic*.⁶

Mere talk of indeterminacy won't work as a magic wand, however. We need a proper analysis. The starting point will be what I take to be two constitutive properties of determinate truth:

DT. If A is determinately true, then A is true.

DK. A can be known true only if A is determinately true.

By themselves these obviously fall short of as a definition or analysis of the concept of determinate truth. Determinate truth could, given (DT + DK) simply be truth *simpliciter*. Determinate truth could also be knowable truth, as this too would satisfy (DT) and (DK). The properties only guarantee that determinate truth is conceptually connected to both truth and knowability, but it becomes an interesting concept only if it can be shown to fill some important gap between truth *simpliciter* and knowable truth.

Williamson (1994) is sceptical of the prospects of such an account, at least in the domain of vagueness. Attacking treatments of vagueness [as in McGee and

⁵ Fine (1985) offers an account according to which free variables denote a particular kind of object, *arbitrary objects* [the idea is older, going back at least to medieval times, see also Price (1922)]. On such an account the indeterminacy involved would not be semantic as the semantic relation between the variable and the object it denotes would be determinate; instead it is the object itself that is indeterminate. I do not think such metaphysically rich extensions to our conceptual apparatus are required in order to understand the phenomena exhibited by free variables.

⁶ Cases of purely semantic indeterminacy stand in contrast to domains where there are grounds for talking about *metaphysical* indeterminacy [see, e.g. Barnes and Williams (2011)]. For instance, it would seem that in the context of quantum mechanics, properties like *position*, *spin* and *momentum* can have indeterminate extension. In such cases we plausibly are dealing with an indeterminacy 'in the world', an indeterminacy that ultimately gives rise to semantic indeterminacy (e.g. the sentence 'The electron has z -spin up' can have indeterminate truth value) but which doesn't originate in the semantic machinery. In the case of vagueness it is perhaps more difficult to locate the source of the indeterminacy—is it 'in the world' or a mere side-effect of language use?—but the properties, say, of being bald, tall or a heap are clearly non-semantic properties, and so if their extensions are indeterminate there is at least one sense in which the indeterminacy is not purely semantic.

McLaughlin (1994)] that invoke the idea that the extension of vague predicates are indeterminate (or ‘indefinite’, I take these to be equivalent) and that thereby the truth value of a sentence like ‘Harry is bald’ is indeterminate, Williamson argues that there is no important gap to be filled:

If we cannot grasp the concept of definiteness by means of the concept of truth, can we grasp it at all? No illuminating analysis of ‘definitely’ is in prospect. Even if we grasp the concept as primitive, why suppose it to be philosophically significant?

The alternative to equating determinate truth with truth *simpliciter*, he suggests, is to equate it with something like knowable truth. Indeed Williamson, famously, argues for an *epistemicist* analysis of vagueness: vague expressions do not have ‘indeterminate’ semantic values in any interesting non-epistemic sense, their semantic values are just (partially) unknowable.

An epistemicist analysis along these lines for the semantics of free variables has been offered by Breckenridge and Magidor (2012). They propose the thesis of Arbitrary Reference:

Arbitrary Reference (AR): It is possible to fix the reference of an expression arbitrarily. When we do so, the expression receives its ordinary kind of semantic-value, though we do not know and cannot know which value in particular it receives.

They consider the posit ‘Let Pierre be an arbitrary Frenchman’, a posit that, while it ensures that ‘Pierre’ receives its ordinary kind of semantic-value, apparently picks out no particular Frenchman, and they explain the resulting ‘indeterminacy’ in purely epistemic terms. While they maintain that ‘*nothing* determines which Frenchman is referred to’ (p. 379) which could lead one to think that they take the reference of ‘Pierre’ to be indeterminate in a non-epistemic sense, they continue: ‘nothing, that is, other than the semantic fact that we have referred to the particular Frenchman in question’ (p. 379). So on their account there is a determinate semantic fact (the semantic fact that we have referred to the particular Frenchman in question) which endows ‘Pierre’ with a determinate (but unknowable) reference. If this sounds circular, it is, and they bite the bullet:

We simply deny that for it to be a semantic fact that some particular Frenchman is being referred to, some other facts need to determine this fact. (p. 379–380)

There is *nothing* in the posit that x is a whole number such that $x^2 = 4$, that could be said to favour one of the values 2 or -2 as the referent of the variable x .⁷ This much the epistemicist will acknowledge. I interpret the epistemicist position to be

⁷ This much is obvious. For instance, one cannot upon letting $x^2 = 4$ and $y^2 = 4$ conclude that $x = y$; for in this situation all that can be concluded is that either $x = y$ or $x = -y$. The epistemicist has to concede this and so has to allow that the way that x and y have acquired their values might differ. Thus whatever it is that determines the value of x in this posit (nothing, it would seem), it need not determine the same value for y .

that the only philosophically interesting claim we can make about this ‘indeterminacy’ is that we cannot know whether the variable x refers to 2 or whether it refers to -2 . If, say, the variable x refers to -2 then the semantic relationship between x and the number 2 is precisely the same as the semantic relationship between the variable x and the number 3. The only philosophically interesting difference is epistemic: we *know* that $x \neq 3$, we do not—indeed *cannot*—know that $x \neq 2$. Treating ‘determinate’ as a non-epistemic (but perhaps philosophically uninteresting) primitive, I take the epistemicist position to be that x has a determinate reference that cannot be known, and so that $x = 2$ has a determinate truth value that cannot be known.

The epistemicist position has its virtues. It is non-revisionary with respect to semantic vocabulary (bivalence is unthreatened, as is the truth-functional interpretation of the connectives), and it is conceptually parsimonious. However, some, like the present author, find the position highly counter-intuitive: the very idea that the variable x on a posit like *let x be a natural number* comes to determinately denote some particular (but unknowable) number seems outlandish. But merely subjecting epistemicism to an ‘incredulous stare’ [c.f. Lewis (1986)] wont do the job, a positive argument is called for. The epistemicist position’s weakest point is its commitment to conceptual parsimony. The most direct way of showing that it is untenable is by providing an analysis of (in)determinate truth that shows it to be a coherent, philosophically significant concept that can be reduced to neither truth *simpliciter* nor unknowable truth, and by showing that the non-revisionary epistemicist is already committed to both the coherence of such a concept (if only by some other name) and its philosophical significance. This is what I propose to do in the remainder of this paper.

3 A supervaluational analysis

Let us turn to a standard truth-conditional semantic theory to make the thesis precise, drawing only on semantic concepts that derive from Tarski. Consider a fragment of our language (English, or ‘mathematical’ English). The fragment needs to be big enough to contain sentences like ‘ $x = 2$ ’ but not so large as to invite paradox. We assume that there is some domain D that is inclusive enough to contain whatever one speaks of in the fragment. The interpretation of singular terms, functions, predicates, and relations is divided up into two components, one that is fixed relative to the model and one that is allowed to vary with context. The fixed component is the interpretation function I . It is assumed that all constants, functions, predicates and relations of the object language are given a fixed interpretation by I (so $I(\alpha) \in D$, where α is a constant, $I(\phi) \subseteq D$ where ϕ is a predicate (and so on for relations and function; I use Greek letters to denote meta-variables that vary over object-language constructs). The interpretation is supposed to be homophonic, so $I(\text{‘is a whole number’}) = \text{the set of whole numbers}$, and $I(\text{‘2’}) = 2$.

To complete our interpretation of terms we need to introduce the notion of an *assignment* g . It will assign values to the variables ‘ x ’, ‘ y ’, etc. of the object

language. For any such variable α , let $g(\alpha)$ be some element of the domain. For any term α we can let $I_g(\alpha) = I(\alpha)$ if α is a constant, and $I_g(\alpha) = g(\alpha)$ if α is a variable.

Holding our model constant (the domain D and the interpretation function I), sentences are attributed truth conditions relative to an assignment g .⁸ For instance, we have:

$$'x = 2' \text{ is true relative to } g \text{ iff } I_g('x') = I_g('2') \text{ iff } g('x') = 2.$$

So if $g('x') = 2$, then $'x = 2'$ is true relative to g , otherwise it is false relative to g .

The boolean connectives and the quantifiers receive a standard interpretation:

' A or B ' is true relative to g iff A is true relative to g or B is true relative to g .

' $\forall xA$ ' is true relative to g iff A is true relative to g' , for any g' that differs from g in at most the value it assigns to x .

And so on. As long as we are dealing with a purely mathematical context (where closed sentences if true, are necessarily true, and if false, are necessarily false) the semantic entailment relation is only sensitive to the assignment function:

A_1, \dots, A_n *semantically entail* B iff B is true relative to every assignment g where all the A_i are true.

This much is standard. The aim is to employ this formal semantic apparatus in a characterisation of what it takes for a sentence to be true in a given context of use and what it takes for a variable to refer to some object in that context (keeping in mind that the semantic values of such constructs can vary with context). Our formal machinery relativizes truth to an assignment, so the task is to explain how such an assignment can be extracted from a context of use in such a way that the used sentence can be evaluated using the formal apparatus [this is also standard, see e.g. Kaplan (1978, 1989)]. For the present purposes a context can be equated with a (finite) set of sentences A_1, \dots, A_n representing the assumptions that have been made in that context (when no assumptions have been made this will be called the 'null' context).⁹ The assumptions made in a context constrain the range of *admissible* assignments:

In a context where all has been assumed is A_1, \dots, A_n , an assignment g is *admissible* if and only if each A_i is true relative to g .

⁸ Tarski avoided speaking of a sentence being *true* relative to an assignment and instead spoke of a sentence being *satisfied* relative to (a vector of values corresponding to) an assignment. Sentence truth he defined as satisfaction by all assignments (what is here called determinate truth). However he explicitly restricted the term 'true' to closed sentences (sentences with no free variables) and he did not consider the effect of posits on the set of admissible assignments. Tarski's semantic analysis was not intended to be applied to the use of stand-alone sentences with free variables.

⁹ In the highly idealised setting of, say, a natural deduction it is easy to keep track of the posits made. However, in a more general account the posits form only one aspect of the semantically important contextual features to keep track of and should be seen as involving a more general linguistic 'score-keeping' (e.g. Lewis 1979). Stalnaker (2014)'s notion of the 'common ground' is obviously also relevant here.

Focusing only on the posits and assumptions that have been made in a context there is a clear sense in which these may under-determine the referent of a variable and, consequently, the truth value of sentences in which it occurs; for in general there will be more than one admissible assignment.¹⁰

With the notion of an admissible assignment at hand one can then proceed to give an analysis of what it means for reference and truth to be (in)determinate:

A variable x *determinately refers to* a in a context iff $g('x') = a$ for all admissible assignments g in that context.

A variable x has *indeterminate reference* in a context iff there are admissible assignments g and g' in that context such that $g('x') \neq g'('x')$.

A sentence A is *determinately true (false)* in a context iff A is true (false) in every admissible assignment in that context.

A sentence A has *indeterminate truth value* in a context iff A is neither determinately true nor determinately false in that context.

One can easily (assuming the language of first order logic) show that B is determinately true in a context where only A_1, \dots, A_n have been assumed iff $A_1 \wedge \dots \wedge A_n \supset B$ is determinately true in the null context (the context where no assumptions have been made) iff $\forall x(A_1 \wedge \dots \wedge A_n \supset B)$ is determinately true in the null context.

So far this is a relatively standard supervaluationist analysis of (in)determinate truth and reference (see, e.g. Mehlberg 1958; Fraassen 1966; Lewis 1972; Fine 1975), albeit applied rather non-standardly to the semantics of free variables. How does this analysis fare? If B is true in all assignments in which the sole posit A is true, then A entails B , and so B will be properly assertable given the posit A . As B is properly assertable (given A), it is, in this context, true. This gives us (DT): determinate truth implies truth. If, on the other hand, B is *not* true in some admissible assignment, this means that the sole posit A does not entail B which means that B is not properly assertable within a context in which only A has been assumed and so—given the posits at hand— B cannot be known true. This gives us (DK): only determinate truths are knowable. In particular, (DT) provides a sufficient semantic criterion for the assertability of 'Either $x = 2$ or $x = -2$ ' in a context where it has been posited that $x^2 = 4$, while (DK) provides a necessary semantic criterion that blocks the assertability of ' $x = 2$ ' in the same context.

The account so far thus provides an explanation for why 'Either $x = 2$ or $x = -2$ ' is assertable while ' $x = 2$ ' is unassertable in terms of their semantic status: the former is determinately true the latter has an indeterminate truth value. But what of

¹⁰ Note that it is quite possible to make assumptions that are not jointly satisfiable, in which case there will be no admissible assignments. I will call a set of posits *proper* if there is some assignment that makes them all true. Thus the posit *let x be such that $x + 1 = x + 2$* would not count as a proper posit, even though it is quite legitimate to *assume* that $x + 1 = x + 2$ and proceed to perform a *reductio*-argument on that basis; *posits* are in this respect taken to be different from assumptions. To *assume* that $A(x)$ where x is free, can be viewed as equivalent to *assuming* that there exists some x such that $A(x)$ ($\exists xA(x)$) and then letting x be such an individual. A posit is an on-the-fly baptism, a naming of an individual with some property or properties, and only existing individuals can be thus named. Throughout, the topic will be *proper* posits.

the two questions that prompted the present discussion? The questions: *What is the truth value of the sentence $x = 2$?* and *What is the value of the variable x ?*

According to the perhaps dominant tradition within supervaluationism (e.g. Mehlberg 1958; Fraassen 1966; Fine 1975; Keefe 2000, 2008) truth *simpliciter* just is determinate truth and falsity *simpliciter* just is determinate falsity. Indeed, van Fraassen—who coined the term—defined a supervaluation as a function that assigns the value *true* (*false*) to all and only to those sentences that are true (*false*) in all admissible valuations. Supervaluationism is on this view a framework for the analysis of our ordinary semantic concepts (truth and reference in a context of use) in cases where the context of use underdetermines the value of the parameters invoked by the apparatus of the formal semantics (the assignment g). The phenomenon of semantic indeterminacy on this view is not so much an indeterminacy of truth and reference but an indeterminacy in how to apply our technical semantic apparatus (truth/reference relative to g) due to an indeterminacy in how the relatum (in our case: the assignment g) is fixed by the context. As the sentence $x = 2$ is neither determinately true nor determinately false, it is, according to this tradition, *not true* and *not false*; rather than saying that it has an *indeterminate* truth value, it is more correct to say that it has *no* truth value (the indeterminacy lies elsewhere). On the same way of thinking indeterminate reference becomes *no* reference: the variable x has no value—there is nothing that it denotes, it fails to refer. The problem with this ‘traditional’ way of applying the supervaluationist framework is that it becomes revisionary in its analysis of the ordinary semantic concepts (it forces us to give up bivalence and the truth functional analysis of the standard boolean connectives) which many, including myself, feel is a hefty price to pay. In addition, I would add, this form of analysis has some, in my mind, deeply counter-intuitive consequences: e.g. if ‘ x ’ refers to nothing, how come that we are in a position to assert that either $x = 2$ or $x = -2$?

There is an alternative way of applying the supervaluational framework. Its starting point is that our practices of language use seem to give rise to cases where truth and reference themselves are truly indeterminate, and uses the supervaluational framework in the analysis of this indeterminacy. Semantic indeterminacy on this view involves an indeterminacy in our ordinary semantic concepts, and the task for the analyst is to make sense of this. Within this tradition we have McGee and McLaughlin (1994, 2004) arguing that the truth predicate inherits the indeterminacy of that to which it applies: if ‘Harry is bald’ has an indeterminate truth value, so does “‘Harry is bald’ is true”. Meanwhile, Lewis (1993) argues that as it plausibly is indeterminate exactly what collection of atoms that make up the cat sitting on the mat, it is indeterminate what physical object that ‘the cat on the mat’ refers to, yet ‘the cat on the mat’ refers to *some* physical object (thus it refers, but its reference is indeterminate). In his extensive discussion Smith (2008) locates the origin of the semantic indeterminacy of vague predicates in an indeterminacy of *intended meaning* and gives it a supervaluational analysis. Though the details differ, these authors concur in treating semantic indeterminacy as a phenomenon that strikes at our ordinary semantic concepts and in using supervaluational methods for analysing this indeterminacy. I will call this ‘radical’ supervaluationism, to keep it apart from ‘traditional’ supervaluationism (Smith dubs this form of analysis

‘plurivaluationism’, a term I find somewhat awkward). The cited authors all apply such a ‘radically’ supervaluational analysis of semantic indeterminacy to the phenomenon of vagueness; I propose that it be applied also to the semantic indeterminacy that arises in the use of free variables.¹¹

The answer to the question *What is the truth value of the sentence $x = 2$?* thus becomes: *Its truth value is indeterminate, but it’s either true or false.* Similarly, the answer to the question *What is the value of the variable x ?* becomes: *Its value is indeterminate, but it’s either 2 or -2 .* I think these answers strike the exactly right balance between the specificity allowed by the posit $x^2 = 4$, and the indeterminacy it induces. ‘ $x = 2$ ’ has a truth value, but it is indeterminate whether it is true or false, and ‘ x ’ refers to some object, but it is indeterminate which object it refers to. The standard non-modal semantic terminology will, on this kind analysis, ‘inherit’ the indeterminacy of the expressions that they are applied to, so “‘ $x = 2$ ’ is true” will have indeterminate truth value, and “The value of the variable ‘ x ’” will be a singular term with indeterminate reference (but it is determinate that it refers to either 2 or -2).

By itself, however, the supervaluational framework provides us with a reductive analysis of the modal notions *determinate truth*, *determinate falsity* and *determinate reference*, not of unqualified *truth*, *falsity* and *reference*. To complete the analysis one needs show that it is possible to make supervaluational sense of the latter concepts in a non-revisionary way that connects the ordinary non-modal semantic concepts to the semantic apparatus of supervaluationism.

4 Non-modal semantic concepts

4.1 The T-schema and its semantic derivation

So how can unqualified truth (truth *simpliciter*) enter the picture? This is an issue just as pressing for the supervaluationist as the epistemicist. It is worth keeping in mind that when a language contains expressions with context dependent semantic values, the truth value of sentences can also vary with context, and so the extension of the non-relational truth-predicate is context dependent. On positing that $x^2 = 4$, ‘ $x = 3$ ’ is false. On positing that $x = 4 - 1$, ‘ $x = 3$ ’ is true. This much is common ground for the supervaluationist and the epistemicist.

¹¹ The two contenders ‘traditional’ and ‘radical’ supervaluationism do not exhaust the logical space. For instance, Belnap (2009) adopts a *quietist* stance and holds that it makes *no sense* to speak of unqualified truth and reference in cases where there is no determinate truth or reference. Again, I think this has counter-intuitive consequences. For in my mind it makes perfect sense to hold that ‘ x ’ either refers to 2 or -2 just as it makes perfect sense to hold that either $x = 2$ or $x = -2$, furthermore, either the value of ‘ x ’ is 2 in which case ‘ $x = 2$ ’ is true, or the value of ‘ x ’ is -2 in which case ‘ $x = 2$ ’ is false. Quietism deprives us of these, in my mind, natural ways of connecting basic semantic vocabulary to ordinary mathematical language (and to other domains where free variables are employed). Quietism on these issues to me suggests that there is a mystery about ordinary non-modal semantic concepts where I do not think there is one.

The obvious place to start when characterising truth is the *T-schema*; the account of unqualified truth should be such that it commands acceptance in all contexts of every instance of the sentence-schema:

(T-schema) ‘#’ is true if and only if #.

Is there some reason why the supervaluationist in particular (as opposed to the epistemicist) should not abide by the T-schema? I think the answer is *no*.¹²

Given the T-schema (and an underlying classical logic) we get bivalence ($'x = 2'$ is not true if and only if $'x \neq 2'$ is true) and a truth-functional analysis of disjunctions: $'x = 2$ or $x = -2'$ is true iff $'x = 2'$ is true or $'x = -2'$ is true. The supervaluationist abiding by the T-schema would not need to be revisionist on these matters.

The idea, of course, is not new. Fine, in his *Vagueness, Truth and Logic* (1975), proposes a supervaluationist analysis of vague sentences and considers using a distinct notion of truth, $'true_T'$ in such a way as to satisfy the T-schema. The result, in Fine's words, is that “The vagueness of $'true_T'$ waxes and wanes, as it were, with the vagueness of the given sentence” (p. 298); in the present context one could rephrase this as: the *indeterminacy* of truth waxes and wanes...with the indeterminacy of the given sentence. Fine, however, does not wish to equate $'true_T'$ with $'true'$ (*simpliciter*), as it would violate the requirement that “the meta-language not be vague or, at least, not so vague in its proper part as the object-language”. McGee and McLaughlin (1994), in articulating their ‘radical’ supervaluationism, reject this line of reasoning:

‘True’ is likewise vague. Because “‘Harry is bald’ is true if and only if Harry is bald” is definitely true—it is, on to the disquotational conception of truth, analytic—“‘Harry is bald’ is true” will be definitely true, definitely false, or unsettled according as ‘Harry is bald’ is definitely true, definitely false, or unsettled. ‘True’ inherits the vagueness of other vague terms, like ‘bald’. (p. 228)

The latter approach will be adopted here. But can the supervaluationist *make sense* of a notion of truth governed by the T-schema? I take this question to mean: can one devise a supervaluationist theory of a fragment of our language in such a way that the T-schema is entailed by the theory? The worry here is that somehow the semantic framework adopted by the supervaluationist would make it impossible to also adopt the T-schema. One could, of course, have a similar worry for the epistemicist. So let us explore whether the semantic framework itself could deliver the T-schema, on an appropriate interpretation of the truth-predicate.

We are going to build a new ‘object-meta-language’ on top of our original object language. To this effect we add to our original object language a term $'\phi'$ for each sentence ϕ in the object language. In addition we add one variable, A , that is allowed to vary over object language sentences. In addition we add the truth- and

¹² The symbol # is to be replaced by some sentence. The difference between a *sentence-schema* like “‘#’ is true” and a sentence with a free variable, like “ A is true” is that the former doesn't become a sentence until # is replaced by a sentence, whereas the latter is already a sentence, a sentence in which the free variable ‘ A ’ has the syntactic role of a singular term.

falsity-predicates ‘is true’ and ‘is false’ and syntactically restrict these predicates so that sentences of the form ‘ α is true’ are well-formed only when α is a term that is allowed to denote a sentence. Thus “‘ $x = 2$ ’ is true” and ‘ A is true’ will now be sentences of our extended object language (the object-meta-language), but “‘ A is true’ is false” will *not* be a sentence of our new object-meta-language.

Semantically we add a new domain D_M alongside the original domain D . D_M contains all the linguistic expressions of the object language. The original interpretation function I remains unchanged as far as object-language constructions are concerned (so only employs the original domain). The new terms of the form ‘ ϕ ’ are given a homophonic interpretation so $I(\phi) = \phi$. Assignment functions as before assign values to the variables of the object language, but now also assign a value to the sentence variable ‘ A ’. So in our new model the interpretation function I and an assignment function g will have to assign values to a portion of the language that is not in the original object language. Let I^o and g^o denote the restriction of I and g , respectively, to the original object language.

The extension of the truth and falsity predicates is given by the following:

$$I_g(\text{‘is true’}) = \{\phi : \phi \text{ is true relative to } g^o\}.$$

$$I_g(\text{‘is false’}) = \{\phi : \phi \text{ is false relative to } g^o\}.$$

That is, the extension of the predicate ‘is true’ assigned by I_g will be the set of object language sentences that are true relative to g^o .

Note that the extensions of the truth- and falsity-predicates come to depend on the assignment g . So, whereas predicates in the object language proper are given a fixed interpretation (by the non-contextual assignment I), the extensions of the truth- and falsity-predicates are allowed to vary. This respects the fact that their extensions are context dependent.

The semantics yields the following truth-conditions (where α is a name of a sentence):

‘ α is true’ is true relative to g iff $I_g(\alpha)$ is true relative to g .

‘ α is false’ is true relative to g iff $I_g(\alpha)$ is not true relative to g .

Our enriched object language and its accompanying semantics constitutes a theory of how we are to assign relational truth values to English sentences like “‘ $x = 2$ ’ is true”. In adopting a particular theory this will regulate how we use such sentences beyond cases that might not be dictated by the pre-theoretical intuitions we wish to respect. It will regulate usage in the same way that assignment relative truth conditions regulate usage of sentences like ‘ $x = 2$ ’.

Let us now derive an instance of the T-schema. Note that “‘ $x = 2$ ’ is true” is true relative to g iff ‘ $x = 2$ ’ is true relative to g . So:

“‘ $x = 2$ ’ is true if and only if $x = 2$ ” is true in every admissible assignment.

That is:

“‘ $x = 2$ ’ is true if and only if $x = 2$ ” is determinately true.

From (DT) we have:

If “‘ $x = 2$ ’ is true if and only if $x = 2$ ” is determinately true, then “‘ $x = 2$ ’ is true if and only if $x = 2$ ” is true.

So:

“‘ $x = 2$ ’ is true if and only if $x = 2$ ” is true.

When a sentence has been proven true one is in a position to assert it, which is what I now do:

‘ $x = 2$ ’ is true if and only if $x = 2$.

This is an instance of the T-schema. Every other relevant (object language) instance of the T-schema can be derived in a similar way.

One can thus see that on this theory of the meaning of our (unqualified) truth-predicate (‘ A is true’ is true relative to g iff A is true relative to g), supervaluationism dictates that one should accept every instance of the T-schema that falls under the object-language.

4.2 The R-schema and its semantic derivation

Consider the R-schema for reference. It states that every instance of the following schema is properly assertable (where # is to be replaced by a singular term):

R-schema ‘#’ refers to #.

Or, equivalently:

R-schema The value of ‘#’ = #.

Just like the T-schema it provides an implicit definition of ‘refers to’ or ‘the value of’. And just like the T-schema it provides a reasonable starting point for our discussion about how to speak of the value that has been assigned to free variables.

Now, there are cases where one might want to deny an instance of the R-schema, for instance, one might want to deny that ‘Superman’ refers to Superman, on the basis that ‘Superman’ fails to refer. Is there, however, some reason why the supervaluationist in particular should deny some instance of the R-schema when the term is a free variable subject to proper posits like the posit that $x^2 = 4$, where we know that there exists a number fitting the description? I think not. I think the supervaluationist should accept the R-schema and so (in the context in which it has been assumed that $x^2 = 4$) reject any suggestion to the effect that the variable x has no value, or that it has many values or that it is meaningless to ask for the value of x . The variable x refers to exactly one number, it has exactly one value, but its value is indeterminate. Let us see where this leads.

One instance of the R-schema is:

The value of ‘ x ’ = x .

So, the supervaluationist and epistemicist alike can always reply—if anyone asks for the value of the variable x in a given context—that the variable x has the value x . This is a strange answer, I admit, but its strangeness can be explained by the fact

that it is a completely uninformative and so uncooperative: it doesn't mean that the answer is false or lacks truth value. Indeed the following would be an equally uncooperative answer:

$$x = x,$$

yet no one denies, given proper posits, that $x = x$. Of course, in many cases we can say something more informative. Given that $x^2 = 4$, we know that either $x = 2$ or $x = -2$ so we can, given the R-schema and the transitivity of identity, derive the meta-language claim:

Either the value of ' x ' = 2 or the value of ' x ' = -2.

This is as informative a reply as we can offer yet neither more nor less informative than the claim that either $x = 2$ or $x = -2$. No new more specific solution to the equation $x^2 = 4$ will show up just because we use semantic vocabulary in the meta-language, but at least we can specify the possible values of x .

One can make supervaluationist sense of the R-schema. Our object-meta-language already contains terms that denote object-language terms (so the object-meta-language term ' x ' denotes the variable x). Now introduce the function-name 'The value of' into our object-meta-language. It is restricted syntactically only to take terms denoting object-language terms as arguments and will yield a term in the extended object-meta-language. Thus, for instance, "The value of ' x '" will be a term in our extended object-meta-language and its semantic value will be an object in our original object-language domain.

We expand the meta-domain D_M so as to contain every possible assignment function on the original domain (we keep in place the restriction that no term or predicate of the object language can denote an element of the meta-domain, so the quantifiers—defined on object-language variables—will only range over the object domain).

The expression 'The value of' is syntactically a functional term. However, it is also context sensitive. So as opposed to the functions of the object language, it cannot be interpreted by the fixed interpretation function I . Instead we let:

$$I_g(\text{'The value of'}) = g^o.$$

This entails:

$$I_g(\text{'The value of } \alpha \text{'}) = g(\text{'The value of'})(I_g(\alpha)) = g^o(I_g(\alpha)).$$

Given this one can derive every object language instance of the R-schema. For it follows that:

"The value of ' x ' = x " is determinately true.

Combining this with (DT) we get

"The value of ' x ' = x " is true.

So "The value of ' x ' = x " is properly assertable, and so:

The value of ' x ' = x .

This vindicates that the R-schema is accessible to both the supervaluationist and epistemicist. There is no need for the supervaluationist to hold that ‘ x ’ lacks reference or has no value, or has many values in a context where there are multiple admissible assignments of value to x . What the supervaluationist should say is that ‘ x ’ lacks *determinate* reference, that it has no *determinate* value only many different *possible* values. In this way the supervaluationist will differ from the epistemicist who claims that ‘ x ’ has a determinate value, only it cannot be known. When talking about the non-modalised value of ‘ x ’ they can both agree that the value of ‘ x ’ is x , and in a context where it has been assumed that $x^2 = 4$, they can both agree that either the value of ‘ x ’ is 2 or it is -2 .

4.3 Truth-tracking assignments

So, in a radical supervaluational mode of analysis one can ‘make sense’ of ordinary unqualified truth and reference, and of the idea that unqualified truth and reference may be indeterminate. A worry one might have, however, is that these notions, as analysed, are left dangling, quite free from any connection to the semantic notions of the supervaluational framework (truth and reference relative to an assignment g). Why take our ordinary semantic concepts and their potential indeterminacy to be at all important if they do not connect to the semantic theory of our choice?

A further worry is about the concepts of truth and reference *in a context of use* u . As a semanticist one is concerned not only with truth and reference in *this* (one’s own) context of use, one is concerned with *all* contexts of use. Accordingly, truth and reference *in a context of use* u are the target concepts of the supervaluational analysis: by means of the formal/technical notions of truth and reference relative to an assignment g one strives to analyse *truth and reference in a context of use*, with the idea that a context of use constrains but doesn’t uniquely determine the relevant assignment g . All we got from the supervaluational analysis, however, was an analysis of the concepts of *determinate* truth and reference in a context of use. What about their non-modal counterparts? Where do they enter the analysis, what is their relationship to the technical concepts of truth and reference relative to an assignment g ?

Let us tackle the first problem first. In a given context an assignment g is *truth-tracking* iff for every A :

A is true relative to g iff A is true.

In general there can in any given context be several truth-tracking assignments—if there are any at all; this holds in virtue of the fact that when the domain is uncountable we can have two distinct assignments g and g' such that for every A : A is true relative to g iff A is true relative to g' . However, if the domain is countable (e.g. the set of whole numbers) and every element of the domain is picked out by some constant in the language (e.g. the numerals) then there is only one truth-tracking assignment—if there is any truth-tracking assignment at all. To simplify the discussion let me assume then that the domain is countable and that every

element of the domain is picked out by some constant, guaranteeing uniqueness of truth-tracking assignments.

The potentially problematic part is *existence*: will there in every context exist a truth-tracking assignment? The epistemicist will of course answer ‘yes’: as there in any proper context exists precisely one assignment that picks out the referent of each variable, that assignment is truth-tracking; as there can only be one truth-tracking assignment, that one assignment is *the* truth-tracking assignment. I claim that the supervaluationist can also answer ‘yes’: there is always a truth-tracking assignment, *the* truth-tracking assignment, giving us, for every object language sentence A :

A is true iff A is true relative to the truth-tracking assignment.

At first blush it might seem surprising that the supervaluationist can make claim to refer to *the* truth-tracking assignment in a context where there are many admissible assignments: any two different admissible assignments will make different sentences true, and so at most one can be truth-tracking. So which one is *the* truth-tracking assignment? The answer is that this is indeterminate. Of any admissible assignment we can say that it is *possible* that it is *the* truth-tracking assignment; of no admissible assignment can we say that it is definitely not the truth-tracking assignment. The property of being an admissible assignment will in every context have a determinate extension, but the property of being a truth-tracking assignment—like the property of being true—will have an indeterminate extension. Thus we can add ‘the truth-tracking assignment’ to the list of meta-language semantic singular terms and definite descriptions (a list already containing “the truth value of ‘ $x = 2$ ’” and “The value of ‘ x ’”) that have, or can have, indeterminate reference. The truth-tracking assignment is an admissible assignment, but it is indeterminate which one.

To see that this makes supervaluationist sense we again need to subject this part of our language to formal scrutiny by putting it into our object-meta-language. Our meta-domain D_M already contains assignment functions. To our object-meta-language we need to add singular terms that are allowed to refer to assignment functions. In particular we add the context sensitive definite description ‘the truth-tracking assignment’, interpreted by:

$$I_g(\text{‘the truth-tracking assignment’}) = g^o.$$

In order to make way for the equivalence between unqualified truth and truth relative to the truth-tracking assignment we also add an assignment relative truth-predicate ‘ α is true relative to β ’ to the language syntactically restricted so that α is a term allowed to denote an object language sentence and β is a term denoting assignments (in this fragment we only have one such term: ‘the truth-tracking assignment’).

Assignment-relative truth is *not* context sensitive, so its extension [a set of pairs (sentence, assignment)] can be given by the fixed interpretation function I :

$$I(\text{'is true relative to'}) = \{(\phi, g) : \phi \text{ is an object language sentence and } g \text{ is an object language assignment such that } \phi \text{ is true at } g\}.$$

It follows that (where α is a term denoting the name of a sentence and β is a term denoting the name of an assignment):

$$\text{'}\alpha \text{ is true relative to } \beta\text{' is true relative to } g \text{ iff } g(\alpha) \text{ is true relative to } g(\beta).$$

[Note that $g(\beta)$ will denote an assignment that only assigns values to object language variables, it won't be an assignment that assigns a value to β ; as $g(\alpha)$ denotes an object language sentence containing only object language variables, this restriction is harmless.] We are now in a position to derive the equivalence stated at the outset of this section.

Recall that in our object-meta language we allowed one free variable that was allowed to take sentences as values: the free variable 'A'. So take some assignment g . From the truth conditions given to unqualified truth we have:

$$\text{'A is true' is true relative to } g \text{ if and only if } g(\text{'A'}) \text{ is true relative to } g.$$

As $g(\text{'A'})$ denotes an object language sentence the truth value of this sentence will depend only on the assignments to object-language variables, and so it follows that:

$$\text{'A is true' is true relative to } g \text{ if and only if } g(\text{'A'}) \text{ is true relative to } g^o. \quad (4)$$

Likewise, from the truth conditions we have:

$$\begin{aligned} \text{'A is true relative to the truth-tracking assignment' is true relative to } g \text{ if and only if} \\ g(\text{'A'}) \text{ is true relative to } g(\text{'the truth-tracking assignment'}). \end{aligned}$$

As $g(\text{'the truth-tracking assignment'}) = g^o$ we have:

$$\begin{aligned} \text{'A is true relative to the truth-tracking assignment' is true relative to } g \text{ if and only if} \\ g(\text{'A'}) \text{ is true relative to } g^o. \end{aligned}$$

Combining this last step with (4) we have:

$$\begin{aligned} \text{'A is true' is true relative to } g \text{ if and only if} \\ \text{'A is true relative to the truth-tracking assignment' is true relative to } g. \end{aligned}$$

So:

$$\text{'A is true if and only if A is true relative to the truth-tracking assignment' is true relative to } g.$$

This holds for arbitrary g , so:

$$\text{'A is true if and only if A is true relative to the truth-tracking assignment' is determinately true.}$$

Throughout it has been assumed that A is an object language sentence. So, assuming only that A is an object language sentence:

A is true if and only if A is true relative to the truth-tracking assignment.

So we have an explicit characterisation of the extension of unqualified truth in terms of assignment-relative truth. But we do not have a reductive analysis of truth *simpliciter* in terms of truth relative to an assignment; the reference of the singular term ‘the truth-tracking assignment’ is, after all, fixed by appeal to (and waxes and wanes with the indeterminacy of) what is true. We cannot get further. If it were possible to define truth *simpliciter* by means of properties with a determinate extension and terms with a determinate reference, the result would be a property with a determinate extension. What we can conclude, however, is that there is some admissible assignment g such that the predicate ‘true’ has the same extension as ‘true relative to the assignment g ’, we just can’t say which assignment that has this property as this is indeterminate. (The epistemicist is in no better position; for the epistemicist will have to acknowledge that it is in principle impossible to state—in a non-question begging manner—which of the admissible assignments that is the truth-tracking assignment.)

A similar story (I will spare the reader the details) establishes the following identity:

(The value of the variable ‘ x ’) = (the value of the variable ‘ x ’ relative to the truth tracking assignment).

This does not amount to a reductive analysis of unqualified talk of reference to talk of reference relative to an assignment—again: no such analysis is in the offing—but it does show how the two are related.

So, turn to the concepts of truth and reference *in a context of use u* (where, I am assuming, u denotes some possible situation that determines some non-empty set of admissible assignments). Just as with their cousins the non-modal, non-relational concepts of truth and reference *simpliciter*, we can’t hope for a reductive analysis. However, as opposed to their non-modal, non-relational counterparts we can’t even turn to the T-schema or R-schema in order to characterise their extension. Still, I claim (again sparing the reader the now onerous details), we can make sense of the definite description ‘the truth tracking assignment in u ’, that is, the description denoting the assignment g such that:

A is true in u iff A is true relative to g .

So placing the predicates ‘true in the context of use u ’ and ‘true relative to the truth tracking assignment in u ’, and the functionals “The value of ‘ x ’ in u ” and “The value of ‘ x ’ relative to the truth tracking assignment in u ” in the meta-object language, one can show that it is possible to make supervenient sense of these expressions in such a way as to guarantee that the following holds:

1. A is true in u iff A is true relative to the truth tracking assignment in u .
2. (The value of ‘ x ’ in u) = (the value of ‘ x ’ relative to the truth tracking assignment in u).

What this shows is that there is some assignment g (the truth tracking assignment in u) such that A is true in u iff A is true relative to g . Such an existence claim cannot be cashed out in terms of a description that determinately refers to an assignment, as it is indeterminate which one of the in u admissible assignments that tracks truth in u . Again, this does not amount to a supervaluational reductive analysis of truth and reference in a context of use, but it does show that the predicate ‘true in u ’ is conceptually distinct from (and need not have the same extension as) the predicate ‘determinately true in u ’.

4.4 ‘Determinately’ in the object language?

Having studied what happens when the semantic vocabulary is put in the object language, a natural question becomes: what happens when we add the ‘Determinately’ operator to the object language, giving us sentences of the form ‘Determinately A ’, or ‘*Det* A ’ for short. Here one should proceed with considerable caution.

On the supervaluational analysis, a sentence is determinately true iff it is true relative to all assignments that make all the posits true. The problem is that in the course of an argument the sets of posits change. New posits can be added and old posits can be retracted. On letting x be such that $x^2 = 4$, ‘*Det* ($x = 2$ or $x = -2$)’ is true. On dropping this posit ‘*Det* ($x = 2$ or $x = -2$)’ is false. This has nothing to do with the fact that determinately is a modality, we get the same phenomenon if we replace determinately true by its *analysans*. On letting x be such that $x^2 = 4$, ‘ $x = 2$ or $x = -2$ ’ is true relative to all assignments that make all the posits currently in force true. Now drop the assumption that $x^2 = 4$. It is then no longer the case that ‘ $x = 2$ or $x = -2$ ’ is true relative to all assignments that make all the posits currently in force true. I have added the phrase *currently in force* to make explicit what hitherto has been implicit: the posits that can be said to be currently in force changes over time, in the course of an argument what counts as a posit currently in force changes, accordingly what counts as determinately true changes.

The standard principles of logic, to the extent that they focus on syntactic form, are not designed to capture logical relations between expressions that can change their semantic value in ongoing discourse or reasoning. An assertion of ‘You [pointing to Bill] are guilty and you [pointing to Anne] are not guilty’ does not express an inconsistent proposition even though it has the syntactic form of a contradiction (‘You are guilty and you are not guilty’). In a similar way, logic is affected when the *Det*-operator is introduced. While *Det* ($x = 2$ or $x = -2$) follows on the assumption that $x^2 = 4$, we cannot on the basis of this drop the assumption, apply the principle of conditional proof, and infer the material conditional:

$$(x^2 = 4) \supset \text{Det}(x = 2 \text{ or } x = -2).$$

For when the posit is dropped the context has changed. As a result the consequent of the conditional becomes false, which thus entails the falsity of the antecedent, i.e. we would wrongly infer that $x^2 \neq 4$. Conditional proof is no longer a valid rule of inference.

If we are to do logic on the basis of syntactic form when we are dealing with these kinds of expressions we have two choices. Either we keep close to the syntactic structure of natural language at the risk of abandoning the standard principles of logic as applied to syntactic form, or we make sure to de-contextualise the expressions in question. For instance, adopting the latter strategy we can rewrite ‘You are guilty and you are not guilty’ as ‘You₁ are guilty and you₂ are not guilty’ where the indices make sure that repeated uses of ‘you’ are not necessarily to be treated as referring to the same individual. In a similar way one can de-contextualise the determinately operator. For instance, one can relativise it to a context of use, as in ‘Det[in the context of use u] B ’, or relativise it to a specific collection of posits, as in ‘Det[given only posits A_1, \dots, A_n] B ’. So, for instance, Det[given only posit $x^2 = 4$] ($x = 2$ or $x = -2$) would be true regardless of what is currently being assumed while Det[given only posit $x^2 = 4$] ($x = 2$) and Det[given no posit] ($x = 2$ or $x = -2$) would be false regardless of what is currently being assumed. Now all the standard principles of classical logic are safe.

In the literature on the *Det*-operator, the first strategy is by far the most popular. Some [e.g. Keefe (2000)] embrace the resulting clash with principles of classical logic. Williamson (1994, p. 151–152) considers such an approach to be ‘revisionary’ with respect to logic. This is certainly true in one sense, but is perhaps a bit unfair. If one wants to do logic keeping the surface syntactic structure of ‘You are guilty and you are not guilty’ (making sure it doesn’t logically behave like a contradiction), clearly some principles of classical logic need to be revised, but this does not necessarily amount to a deep break with the semantic commitments that we take to underlie classical logic (e.g. the fact that we allow that a singular term like ‘you’ can denote different individuals when it occurs in different parts of a sentence does not mean that we allow that one and the same person can, at a given moment, be both guilty and not guilty). The same can be said for ‘revisionary’ approaches to the *Det*-operator: the classical principles of reasoning were not developed for cases where the truth value of a sentence changes as we reason (and *because* we reason). In any case, one can avoid such charges of revisionism by de-contextualising the determinately modality in the way suggested above.

5 The epistemology of indeterminacy

5.1 Indeterminacy is a non-epistemic modality

An assignment of values to variables is admissible in a context if and only if all posits of the context become true relative to the assignment. This is the core modal notion through which determinate truth, determinate reference, etc., have been analysed here. On this analysis admissibility is clearly not an epistemic notion. The analysis proceeds entirely in terms of truth relative to an assignment and the posits that govern the context: two non-epistemic concepts. The resulting notion of ‘determinate’ truth, which is analysed through purely non-epistemic concepts, is clearly not an epistemic modality.

There is still the possibility that *determinately true* and *knowably true* are co-extensional concepts. This would be a substantial thesis; for it would, among other things, imply that every true closed sentence is knowably true (closed sentences are true iff they are determinately true). If knowably true in the context of mathematics is equivalent to provably true, it would imply that every true closed sentence of mathematics is provably true. Since Gödel we have reason to be suspicious of such claims. Barring a revision of classical logic, we have no reason to think that determinate truth and knowable truth are co-extensional concepts.

5.2 Knowledge requires determinate truth

Knowledge implies truth, on this all agree. But we should be prepared to make a stronger claim: knowledge implies *determinate* truth. This is encoded in the principle (DK): if a sentence has indeterminate truth value, it cannot be known true. If it is either knowable that P or knowable that not- P , then it is either determinately true that P or it is determinately false that P . Is this to be considered a bold revisionary analysis of knowledge? No. With only a minor change in vocabulary, the non-revisionary epistemicist will agree: if a sentence A is not true in all admissible assignments, it cannot be known true. Barring a fundamental change in logic, it just won't happen that mathematicians one day will learn that on assuming no more than that $x^2 = 4$, it turns out, say, that $x = -2$. Even if it is true that $x = -2$, we know that this can't be known. So on any account, knowledge requires truth in all admissible assignments. For the radical supervaluationist, this just amounts to (DK). Rather than being a covert epistemic modality, (in)determinate truth provides a boundary for what is knowable, a boundary drawn by purely non-epistemic means.

Dorr (2003) has identified an apparent tension in the epistemology of indeterminacy, a tension between the principle of unknowability (DK) and the idea that in cases of genuine indeterminacy there is 'no ignorance' involved as there is 'nothing to know'. Dorr gives the following analysis of ignorance: S is ignorant of the fact that P iff P and S doesn't know that P , and shows, in just a few steps, that the 'no ignorance' intuition is in conflict with (DK).¹³

I share Dorr's 'no ignorance' intuition. In letting x be such that $x^2 = 4$ the fact that we are neither in a position to judge that $x = 2$ nor in a position to judge that $x \neq 2$ has nothing to do with ignorance in any normal sense of the word: there is no ignorance, *for there is nothing—no determinate fact—to know*. Still, surely it is correct to say that it is not the case that one knows that $x = 2$. There is a lack of knowledge, but no ignorance. The problem, as I see it, is not that the 'no ignorance' intuition conflicts with (DK), but rather the analysis of ignorance offered by Dorr. If the 'no ignorance' intuition is driven by the idea that there is no determinate fact to

¹³ E.g. assume, in accordance with the 'no ignorance' intuition, that it is not the case that S is ignorant about the fact that $x = 2$ and it is not the case that S is ignorant about the fact that $x \neq 2$. So it is not the case that [$x = 2$ and S does not know that $x = 2$], and it is not the case that [$x \neq 2$ and S does not know that $x \neq 2$]. By (DK), S does not know that $x = 2$ and S does not know that $x \neq 2$. Given some classical logic it thus follows that $x = 2$ and $x \neq 2$, a contradiction.

know, this suggests the following alternative analysis: S is ignorant of the fact that P iff it is *determinately* the case that P and S doesn't know that P . Given this analysis it is perfectly ok to say that there is no ignorance involved even though we do not know whether $x = 2$ or $x \neq 2$.

5.3 Valid arguments preserve determinate truth

Valid arguments preserve determinate truth. A valid rule of inference is a rule that guarantees that, once the conditions for its application are met, the conclusion drawn is determinately true if the assumptions are determinately true. To the extent that a valid rule of inference invokes one or more premises,¹⁴ the conditions for its application are met only if the premises are determinately true given that the assumptions in force are determinately true. Determinate truth, not truth *simpliciter*, is the property preserved in a valid argument.

Given that valid arguments are an important source of knowledge, and that knowledge requires determinate truth, it should come as no surprise that it is preservation of determinate truth, not truth *simpliciter*, that is the semantically defining property of a valid argument. But it is instructive to see why truth *simpliciter* cannot fill this role. In particular, consider the rule of Universal Generalisation:

$$(UG) \quad \frac{A}{\forall xA}, \quad \text{provided } x \text{ does not occur free in any surviving premises.}$$

This is a classically valid rule, yet there is a distinct sense in which it is not truth-preserving: the truth of the premise does not suffice for the truth of the conclusion. For say that it is truth-preserving in this sense, then the following holds for arbitrary A :

(UG-TP) If A is true and x does not occur free in any surviving premises (x nfp), then ' $\forall xA$ ' is true.

So the following instances hold:

(UG-TPa) If ' $P(x)$ ' is true and x nfp, then ' $\forall xP(x)$ ' is true.

(UG-TPb) If ' $\neg P(x)$ ' is true and x nfp, then ' $\forall x\neg P(x)$ ' is true.

Combining these using classical logic we get:

(X) If either ' $P(x)$ ' or ' $\neg P(x)$ ' is true, and x nfp, then either ' $\forall xP(x)$ ' or ' $\forall x\neg P(x)$ ' is true.

Meanwhile, by classical logic the law of excluded middle ' $P(x)$ or $\neg P(x)$ ' is true. Given that 'or' is truth-functional we get bivalence: either ' $P(x)$ ' is true or ' $\neg P(x)$ ' is true. In a context where x does not occur free in any surviving premises, both conjuncts in the antecedent of (X) thus hold and so it follows that either ' $\forall xP(x)$ ' or ' $\forall x\neg P(x)$ ' is true. But it is not hard to find a predicate P such that both are false (e.g.

¹⁴ For instance, the rule of modus ponens $A, A \supset B/B$ contains two premises, A and $A \supset B$. By contrast in the rule of *reductio ad absurdum* (if a contradiction can be derived from the assumption A , then $\neg A$) there is no premise, only a condition to be fulfilled in order to draw the conclusion $\neg A$.

let $P(x)$ be $x = 2$). On the basis of classical principles and the assumption that Universal Generalisation is truth-preserving (UG-TP) we have reached a contradiction, and the culprit is the assumption that Universal Generalisation is truth-preserving: it is not.

The rule of Universal Generalisation does not preserve truth *simpliciter*. Still, there is no doubt that it is valid. Why? Because it preserves the property of determinate truth:¹⁵

If A is *determinately* true and x does not occur free in any surviving premises, then ' $\forall xA$ ' is determinately true.

Valid rules of inference preserve determinate truth, not (not necessarily anyway) truth *simpliciter*. There are various ways of misunderstanding this claim.

First of all, there is no denying that if one is in a position to assert that A is true (if, that is, one has warrant for this claim), and there is a valid argument from A to B , then one is in a position to assert that B is true. In particular, if one is in a position to assert that A is true, and x does not occur free in any remaining assumptions, then one is in a position to assert that ' $\forall xA$ ' is true. In this sense UG is truth preserving. But this is only because whenever one is in a position to assert that A is true *simpliciter* one is also in a position to assert that A is determinately true;¹⁶ any proof or evidence that A is true is also proof or evidence that A is determinately true.

Likewise, there is no denying that if A logically entails B , then B is true whenever A is true: the consequence relation is truth-preserving (given the T-schema and classical logic this is trivial). A potential source of confusion, however, is a certain tendency to view rules of inference as equivalent to or implying corresponding claims about entailment. But the question whether a rule of inference is valid does not in general reduce to the question whether the premises of the inference entail its conclusion. Universal Generalisation does not amount to or imply the claim that A logically entails ' $\forall xA$ ' whenever there are no surviving premises with x free. For first, the proviso that there be no surviving premises makes no sense in this context: the relation of logical consequence is a relation between sentences or propositions, talk of 'premises' has no place in this context. Second, it is clear that a sentence A can be true relative to some assignment and some model even though ' $\forall xA$ ' is false relative to that assignment in that model; so it is not in general the case that A logically entails ' $\forall xA$ '.

What does hold, of course, is the following: if a set of sentences Γ entails A and x does not occur free in Γ , then Γ entails ' $\forall xA$ '. So there is no denying that there is a

¹⁵ To see this assume that A is determinately true and that x does not occur free in any surviving premise. From the latter we know that for every element in the domain there is some admissible assignment that takes ' x ' to refer to that element. As A is true in all admissible assignments it is thus true for all x .

¹⁶ So: one is in the position to assert ' A is true' iff one is in the position to assert ' A is determinately true'. This does not imply that the two sentences are logically equivalent, or that the two predicates are substitutable *salva veritate*. For instance, 'Either A is true or $\neg A$ is true' is classically a logical truth, but 'Either A is determinately true or $\neg A$ is determinately true' may well be false. What it does imply is that the predicate " x is in the position to assert ' A is true'" can be substituted *salva veritate* for " x is in the position to assert ' A is determinately true'".

close knit relationship between the validity of the inference rule UG and the consequence relation. But this is just another way of bringing out the role of determinate truth in valid inferences. For if Γ entails A and constitutes a set of proper posits, then A is determinately true in the context in which Γ has been posited.

As long as ‘determinate truth’ is analysed as ‘truth in all admissible assignments’, the non-revisionary epistemicist will agree to the above. Valid arguments preserve the property of being true in all admissible assignments. A valid rule of inference is a rule that guarantees that, once the conditions for its application are met, the conclusion drawn is true in all admissible assignments given the assumptions in force. And so on. The distinctive semantic value preserved in a valid argument has philosophical significance, no matter what we call it, and the fact that this semantic role is filled by *determinate truth* (by this or any other name) rather than truth *simpliciter*, is of philosophical significance.

Summing up, I hope to have established that there is at least one domain—the semantics of free variables—where the concept of determinate truth can be given a coherent analysis in which it coincides with neither truth *simpliciter* nor knowable truth and where it plays a philosophically significant role, a role that the non-revisionary epistemicist must concede. Establishing this much has no immediate consequences as regards indeterminacy in other domains—like vagueness, the interpretation of quantum mechanics, or talk about the future—but I think it at least undermines complaints to the effect that the concept of indeterminacy is in itself philosophically empty, logically or semantically revisionary, unintelligible, or in some other way incoherent.

Acknowledgements I wish to thank Lee Walters, an anonymous referee, and the participants at the philosophy seminars at KTH and Uppsala, for valuable comments and pointers.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Barnes, E., & Williams, J. R. G. (2011). A theory of metaphysical indeterminacy. In K. Bennett & D. W. Zimmerman (Eds.), *Oxford studies in metaphysics* (Vol. 6, pp. 103–148). Oxford: Oxford University Press.
- Belnap, N. (2009). Truth values, neither-true-nor-false, and supervaluations. *Studia Logica*, 91(3), 305–334.
- Breckenridge, W., & Magidor, O. (2012). Arbitrary reference. *Philosophical Studies*, 158, 377–400.
- Dorr, C. (2003). Vagueness without ignorance. *Philosophical Perspectives*, 17, 83–113.
- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30, 265–300.
- Fine, K. (1985). *Reasoning with arbitrary objects*. Oxford: Basil Blackwell.
- Gentzen, G. (1969). Investigations into logical deduction. In M. Szabo (Ed.), *The collected works of Gerhard Gentzen*. Amsterdam: North Holland.
- Kaplan, D. (1978). On the logic of demonstratives. *Journal of Philosophical Logic*, 8(1), 81–98.

- Kaplan, D. (1989). Demonstratives and afterthoughts. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan*. Oxford: Oxford University Press.
- Keefe, R. (2000). *Theories of vagueness*. Cambridge: Cambridge University Press.
- Keefe, R. (2008). Vagueness: Supervaluationism. *Philosophy Compass*, 3(2), 315–324.
- Lewis, D. (1972). General semantics. In D. Davidson & G. Harman (Eds.), *Semantics of natural language*. Dordrecht: Reidel.
- Lewis, D. (1979). Score-keeping in a language game. *Journal of Philosophical Logic*, 8(1), 339–359.
- Lewis, D. (1986). *On the plurality of worlds*. Oxford: Blackwell.
- Lewis, D. (1993). Many, but almost one. In K. Campbell, J. Bacon, & L. Reinhardt (Eds.), *Ontology, causality and mind: Essays on the philosophy of D. M. Armstrong*. Cambridge: Cambridge University Press.
- McGee, V., & McLaughlin, B. (1994). Distinctions without a difference. *The Southern Journal of Philosophy*, 33, 203–251.
- McGee, V., & McLaughlin, B. (2004). Logical commitment and semantic indeterminacy: A reply to Williamson. *Linguistics and Philosophy*, 27, 123–136.
- Mehlberg, H. (1958). *The reach of science*. Cambridge: Cambridge University Press.
- Price, R. (1922). Can there be random individuals? *Analysis*, 22, 94–96.
- Quine, W. V. (1960). Variables explained away. *Proceedings of the American Philosophical Society*, 104(3), 343–347.
- Smith, N. J. J. (2008). *Vagueness and degrees of truth*. Oxford: Oxford University Press.
- Stalnaker, R. (2014). *Context*. Oxford: Oxford University Press.
- van Fraassen, B. C. (1966). Singular terms, truth value gaps, and free logic. *Journal of Philosophy*, 63(17), 481–495.
- Williamson, T. (1994). *Vagueness*. London: Routledge.