

(to appear in J. Salerno, ed., *New Essays on the Knowability Paradox*,
Oxford: Oxford University Press)

Tennant's Troubles

Timothy Williamson

First, some reminiscences. In the years 1973-80, when I was an undergraduate and then graduate student at Oxford, Michael Dummett's formidable and creative philosophical presence made his arguments impossible to ignore. In consequence, one pole of discussion was always a form of anti-realism. It endorsed something like the replacement of truth-conditional semantics by verification-conditional semantics and of classical logic by intuitionistic logic, and the principle that all truths are knowable. It did not endorse the principle that all truths are known. Nor did it mention the now celebrated argument, first published by Frederic Fitch (1963), that if all truths are knowable then all truths are known.

Even in 1970s Oxford, intuitionistic anti-realism was a strictly minority view, but many others regarded it as a live theoretical option in a way that now seems very distant. As the extreme verificationist commitments of the view have combined with accumulating decades of failure to reply convincingly to criticisms of the arguments in its favour or to carry out the programme of generalizing intuitionistic semantics for

mathematics to empirical discourse, even in toy examples, the impression has been confirmed of one more clever, implausible philosophical idea that did not work out, although here and there old believers still keep the flame alight.

A diffuse philosophical tendency cannot be refuted once and for all by a single rigorous argument. Nevertheless, such an argument can severely constrain the forms in which the tendency is expressed. The tendency labelled ‘anti-realism’ and Fitch’s argument together constitute a case in point. My first publication (1982) was a response to the Fitch argument. I argued that it was intuitionistically invalid, and therefore did not show intuitionistic anti-realism to be committed to the absurd claim that all truths are known. Naturally, my aim was not to endorse intuitionistic anti-realism; I found it as deeply implausible then as I do now. But that does not distinguish it from other forms of anti-realism, and such dispositions are not invariant across persons. My aim was rather to assess what forms of anti-realism must be argued against in some way beyond Fitch’s. The advantage in final plausibility of those other forms of anti-realism over the brazen assertion that all truths are known is tenuous at best: but it is still worth getting clear about the logical situation. Some of my later work on the Fitch argument (1988; 1992; 1994) refined the envisaged response to the Fitch argument and established its formal stability.

In *The Taming of the True* (1997), Neil Tennant objects to the specific intuitionistic anti-realist response to Fitch that I had envisaged, and proposes his own alternative responses, still of a broadly intuitionistic anti-realist kind. In response (2000b), I argued that both Tennant’s objections and his alternatives fail, and that the result illustrates a more general point: that moderate forms of anti-realism tend to be the

least stable. Tennant replied at length (2001a). For some time I thought that the problems with his 2001 reply were sufficiently evident to make any further response from me unnecessary. Later experience has taught me otherwise. The purpose of this paper is to show that Tennant's reply fails completely to meet the difficulties that I raised in 2000. Since his reply engages with many details of the 2000 paper, while missing the relevance of some of the most crucial ones, the most efficient course is to rehearse the arguments of that paper, interspersing them with discussion of Tennant's objections as they arise. Thus the present paper constitutes a self-standing critique of Tennant's treatment of the Fitch paradox that properly includes its predecessor.¹

I

The first task is to expound the Fitch argument in a form suitable for the subsequent discussion.

Anti-realists argue that truth is epistemically constrained. Their arguments are too complex, elusive and at least vaguely familiar to formulate here. We can gesture at them thus: a sentence *s* as uttered in some context expresses the content that *P* only if the link between *s* and the condition that *P* is made by the way speakers of the language use *s*; their use must be sensitive to whether the condition that *P* obtains; that requires of them the capacity in principle to recognize that it obtains, when it does so; thus *P* only if speakers of the language can in principle recognize that *P*. In brief: all truths are knowable. We can formalize the anti-realist conclusion in a schema:

$$(1) \quad \varphi \rightarrow \diamond \mathbf{K}\varphi$$

Here \diamond and \mathbf{K} abbreviate ‘it is possible that’ and ‘someone sometime knows that’ respectively; φ is to be replaced by declarative sentences.² Presumably, the relevant sense of ‘possible’ is not merely epistemic, because the anti-realist takes sensitivity to whether a condition obtains to require a genuine recognitional capacity (a metaphysical possibility of knowing), not a mere incapacity to recognize one’s ignorance (an epistemic possibility of knowing). According to (1), the truth really could have been known.

Much in the anti-realist arguments deserves to be questioned. Fitch (1963) introduced a direct objection to their conclusion with an apparent reduction ad absurdum of (1). It requires two highly plausible principles about knowledge: only truths are known, and known conjunctions have known conjuncts. More formally:

$$(2) \quad \mathbf{K}\varphi \rightarrow \varphi$$

$$(3) \quad \mathbf{K}(\varphi \wedge \psi) \rightarrow (\mathbf{K}\varphi \wedge \mathbf{K}\psi)^3$$

Principles (2) and (3) jointly entail that nothing is ever known to be an always unknown truth. For (2) yields $\mathbf{K}(\varphi \wedge \neg \mathbf{K}\varphi) \rightarrow (\varphi \wedge \neg \mathbf{K}\varphi)$ and therefore $\mathbf{K}(\varphi \wedge \neg \mathbf{K}\varphi) \rightarrow \neg \mathbf{K}\varphi$, while (3) yields $\mathbf{K}(\varphi \wedge \neg \mathbf{K}\varphi) \rightarrow (\mathbf{K}\varphi \wedge \mathbf{K}\neg \mathbf{K}\varphi)$ and therefore $\mathbf{K}(\varphi \wedge \neg \mathbf{K}\varphi) \rightarrow \mathbf{K}\varphi$, so together they give:

$$(4) \quad \neg \mathbf{K}(\varphi \wedge \neg \mathbf{K}\varphi)$$

Principles (2) and (3) are intended as necessary constraints on knowledge, and the propositional logic used to derive (4) from (2) and (3) is necessarily truth-preserving, so by a variant of the rule of necessitation in modal logic we can conclude that what (4) says is not could not have been:

$$(5) \quad \neg \diamond \mathbf{K}(\varphi \wedge \neg \mathbf{K}\varphi)$$

Now consider the special case of (1) with $\varphi \wedge \neg \mathbf{K}\varphi$ in place of φ :

$$(6) \quad (\varphi \wedge \neg \mathbf{K}\varphi) \rightarrow \diamond \mathbf{K}(\varphi \wedge \neg \mathbf{K}\varphi)$$

By (5) and (6):

$$(7) \quad \neg(\varphi \wedge \neg \mathbf{K}\varphi)$$

In classical logic, (7) is equivalent to:

$$(8) \quad \varphi \rightarrow \mathbf{K}\varphi$$

Of course, (8) is deeply implausible. It says in effect that any truth is known. As one instance, it says that there is a fragment of Roman pottery at a certain spot only if

someone sometime knows that there is a fragment of Roman pottery there. The corresponding instance of (1) for the same value of ϕ is quite plausible: there is a fragment of Roman pottery there only if it *could have been* known by someone sometime that there was a fragment of Roman pottery there. But according to (8), there is a fragment of Roman pottery there only if the possibility of knowing is actualized; that claim is quite unwarranted. Although, by (4), the attempt knowingly to identify a particular example of an unknown truth would be self-defeating, we surely have ample evidence of a less direct sort that not every truth is ever known.

Although some believe that an omniscient god makes (8) true, that issue is not very pertinent here. For if we restrict the substitutions for ϕ to sentences of our language, the anti-realist motivation for (1) allows us to read ‘someone’ in the definition of **K** as ‘some member of our speech community’; the links between sentences of English and their contents are made by human speakers of English without divine intervention. ‘There is a fragment of Roman pottery at that spot’ is a sentence of our language; surely not every truth expressible in our language will ever be known by some member of our speech community. If you think it matters, give **K** that unrestricted reading.

We should thus reject schema (8). On the assumption that (8) was derived from (1) using uncontentious principles, we should therefore reject schema (1) too.

The seminal presentation of the case for (1) is Michael Dummett’s (1959, 1975 and elsewhere). Notoriously, he integrates it with a case for a comprehensive anti-realist reconception of meaning in terms of verification-conditions rather than truth-conditions, which, he argues, will invalidate classical logic, in particular the law of excluded middle, and justify its replacement by something like intuitionistic logic. The latter was originally

proposed as the logic of intuitionistic mathematics, and its intended semantics reflects that role, being formulated in terms of the notion of proof. Since the mathematical notion of proof is inappropriate to empirical statements, Dummett envisages a generalized intuitionistic semantics in which a broader notion of verification plays the key role.⁴ Even granted that his arguments are not compelling, it is pertinent to ask how the attempted reduction ad absurdum of (1) fares under intuitionistic logic, which we may provisionally suppose the anti-realist defender of (1) to have adopted.⁵

The argument from (1) to (7) is intuitionistically acceptable, but the step from (7) to (8) is not. Intuitionistically, (7) is equivalent to:

$$(9) \quad \neg \mathbf{K}\phi \rightarrow \neg\phi$$

Intuitionistically, we cannot reach (8) from (9), deleting the negations.⁶ Intuitionists can consistently accept (9) while denying that all truths are known. Since (7) generalizes, they must deny that there is an unknown truth: on their constructivist understanding of the existential quantifier, one could in any case verify that one could never verify that existential claim, because (by (4)) one could never verify a particular instance of it. Intuitionistically, to verify that one could never verify ψ is to verify $\neg\psi$. Denying that there is an unknown truth does not commit one intuitionistically to asserting that all truths are known. Given (9), the intuitionist can consistently deny the universal generalization of (8) but cannot consistently deny any particular instance of (8), for (9) is intuitionistically equivalent to the double negation of (8). In this respect, the intuitionistic status of $\phi \rightarrow \mathbf{K}\phi$ given (9) is exactly like that of the law of excluded middle. For the

intuitionist can consistently deny the universal generalization of $\phi \vee \neg\phi$ but cannot consistently deny any particular instance of it, because $\neg\neg(\phi \vee \neg\phi)$ is intuitionistically valid.⁷ Call an anti-realist position on which (7) is valid and (8) invalid *moderately hard* anti-realism. It is opposed to *very hard* anti-realism, on which both (7) and (8) are valid, and *soft* anti-realism, on which both (7) and (8) are invalid.⁸ A consistent moderately hard anti-realist view can be worked out in some detail within the framework of intuitionistic logic (Williamson 1982, 1988, 1992). Moderately hard anti-realists may regard the Fitch argument as a further reason for anti-realists in general to adopt intuitionistic rather than classical logic, since their response to the challenge depends on a distinction available within intuitionistic but not classical logic — although of course this was not Dummett’s reason for proposing intuitionistic logic as an appropriate logic for anti-realism.

Obviously, (9) itself is a deeply problematic consequence of (1). According to (9), what is never known is not true: thus if no one ever knows that there is a fragment of Roman pottery at a certain spot, there is no fragment of Roman pottery there. That sounds as bad as (8). But there are differences. Schema (8) eliminates the logical distinction between ϕ and $\mathbf{K}\phi$, for (8) and its uncontentious converse (2) jointly yield $\phi \leftrightarrow \mathbf{K}\phi$, and therefore $\Sigma(\phi) \leftrightarrow \Sigma(\mathbf{K}\phi)$ for any sentential context $\Sigma(\)$ defined with the standard intuitionistic connectives (in particular, of course, $\neg\phi \leftrightarrow \neg\mathbf{K}\phi$). Although (9) eliminates the logical distinction between $\neg\phi$ and $\neg\mathbf{K}\phi$, since (9) and (2) jointly yield $\neg\phi \leftrightarrow \neg\mathbf{K}\phi$, and therefore $\Sigma(\neg\phi) \leftrightarrow \Sigma(\neg\mathbf{K}\phi)$, it does not eliminate the logical distinction between ϕ and $\mathbf{K}\phi$, as the underivability of (8) from (9) shows. The moderately hard anti-realist loses fewer distinctions than does the very hard anti-realist. Indeed, the former can consistently deny the universal generalization of (8), while the latter must assert it. Thus

the moderately hard anti-realist, unlike the very hard anti-realist, can assert a gap between what is true and what is ever known. For reasons not peculiar to anti-realism, the acknowledgement of the gap is essentially general; it cannot be made at the level of an individual sentence, because one cannot knowingly present a specific instance of a never known truth.

But how can the moderately hard anti-realist mitigate the implausibility of particular instances of (9)? If one knew that human life was about to be eliminated by a huge meteorite, might one not be entitled to assert that no one will ever know that there is a fragment of Roman pottery at that spot without being entitled to assert that there is no fragment of Roman pottery there?⁹ Such examples suggest that the intuitionistic operator \neg does not correctly formalize the negation that we apply to empirical sentences, such as those of the form $\mathbf{K}\phi$. Dummett himself distinguishes the negation of intuitionistic mathematics from empirical negation (1977: 337). That is not to say that intuitionistic negation cannot meaningfully be applied to empirical sentences; rather, another negation operator may be needed as well to interpret ‘not’ in empirical discourse. Empirical negation must itself behave non-classically if moderately hard anti-realism is not to collapse into very hard anti-realism, for otherwise empirical negation could be substituted for \neg throughout the derivation of (8), but it must not behave non-classically in exactly the same ways as \neg , otherwise it would offer no advantage. Major difficulties face the attempt to add such an operator (Williamson 1994). In what follows, I continue to assume that the anti-realist employs intuitionistic negation; I do not seek to minimize the associated problems. Moderately hard anti-realism remains a deeply problematic

position. Nevertheless, it avoids the *most* drastic consequences of very hard anti-realism, and a full critique of it will not simply cite Fitch.

Of those engaged in the refinement of Dummett's programme and the attempted generalization of intuitionistic semantics to empirical discourse, one of the most active has been Neil Tennant (1987, 1997). In his 1997 analysis of the Fitch problem, he argues that the envisaged moderately hard anti-realist line is unstable in the crucial test cases for the Fitch argument, and suggests an alternative soft anti-realist strategy based on restricting the knowability principle (1) (1997: 245-79).¹⁰ Section II below shows that his objection to the envisaged moderately hard anti-realist line is fallacious, and that his later attempt to defend his objection merely changes the subject. Section III shows that under rather general conditions his restricted version of (1) still implies (7) and (9), so that his would-be soft anti-realism collapses into the moderately or very hard view; his later attempt to defend his restricted version of (1) results in its trivialization.

II

Tennant objects to the moderately hard anti-realist treatment of the Fitch argument that I described that it is unstable, because in what I presented as crucial test cases the distinction between (7) and (8) collapses (1997: 268). He correctly points out that in the examples I use to illustrate the argument, I instantiate ' ϕ ' in (7) and (8) with a decidable sentence: we have a decision procedure whose application would result in a verification or falsification of ϕ ('There is a fragment of Roman pottery at that spot'). This is the

simplest and most vivid form of the problem that Fitch raises, and a crucial test for any adequate treatment. Tennant also correctly points out that (7) is intuitionistically equivalent to (8) when $\mathbf{K}\phi$ is decidable.¹¹ He then attempts to argue that $\mathbf{K}\phi$ is decidable if and only if ϕ is decidable. If he is right, the envisaged moderately hard anti-realist line fails the crucial test.

Before proceeding, I note an obvious condition of adequacy on Tennant's critique. On pain of irrelevance, it must not depend on a reading of the operator ' \mathbf{K} ' other than that intended by the object of his criticism. It is to be read as 'someone sometime knows that', where 'knows' itself is understood in its predominant ordinary sense. In this sense, I may fail to know whether a given very large natural number (as presented by a corresponding numeral) is prime, even though I have a decision procedure for finding out. Of course, I know many things that I am not currently thinking about; my knowledge is stored. But I do not know something merely in virtue of its being routine for me to find out, if I do not in fact find out, or merely in virtue of its being a logical consequence of other things that I know. This is the understanding of 'know' that is evidently in play throughout my presentations of the moderately hard anti-realist line (1982, 1988, 2000) and in most other discussion of Fitch's argument. It is a very natural understanding in the context of that argument, which concerns the relation between potential and actual knowledge, a contrast obscured if ' \mathbf{K} ' is itself understood as meaning something potential. In order to preserve the relevance of Tennant's critique, I will therefore take it for the time being in terms of the usual reading of ' \mathbf{K} '. Another reading will be considered later.

Tennant's objection is sustained if he can demonstrate that $\mathbf{K}\phi$ is decidable whenever ϕ is; the converse does not concern us here. I will argue that his supposed demonstration that the decidability of ϕ implies that of $\mathbf{K}\phi$ is fallacious. Here it is:

Suppose that ϕ is decidable. Then here is a decision method for $\mathbf{K}\phi$: apply the given decision method for ϕ . If you thereby determine that ϕ is true, then you know that ϕ . So you have determined that $\mathbf{K}\phi$ is true. If, on the other hand, you determine that ϕ is false, then you have determined that $\mathbf{K}\phi$ is false, because no one could ever know a falsehood. So if ϕ is decidable, then so is $\mathbf{K}\phi$. (1997: 262)

Our possession of a decision procedure for $\mathbf{K}\phi$ entitles us to assert:¹²

(10) $\mathbf{K}\phi \vee \neg\mathbf{K}\phi$

Intuitionistically, (7) and (10) jointly entail (8). If Tennant is right, the difference between moderately and very hard anti-realism disappears in paradigms of the cases where it was supposed to help.

Why does Tennant suppose that our possession of a decision procedure for $\mathbf{K}\phi$ entitles us to assert (10)? He is relying on a principle about the assertibility of disjunctions:

(DIS) Our possession of a method whose application will either verify ϕ or verify ψ entitles us to assert $\phi \vee \psi$.

(DIS) allows us to assert the disjunction in advance of actually applying the method.¹³ Since the application of Tennant's decision procedure will supposedly either verify $\mathbf{K}\phi$ or verify $\neg\mathbf{K}\phi$, by (DIS) our mere possession of the decision procedure, in advance of actually applying it, entitles us to assert $\mathbf{K}\phi \vee \neg\mathbf{K}\phi$. At first sight, (DIS) looks very plausible. It is surely correct when ϕ and ψ are mathematical.

Nevertheless, we can easily see that Tennant's argument must be unsound on the present reading of 'K'. For if it were sound, there would also be a sound argument for a much stronger conclusion, namely, that our possession of a decision procedure for ϕ entitles us to assert this:

$$(11) \quad \mathbf{K}\phi \vee \mathbf{K}\neg\phi$$

We can argue for this conclusion in Tennant's style:

Suppose that ϕ is decidable. Then apply the given decision method for ϕ . If you thereby determine that ϕ is true, then you know that ϕ . So you have determined that $\mathbf{K}\phi$ is true. If, on the other hand, you determine that ϕ is false, then you know that $\neg\phi$. So you have determined that $\mathbf{K}\neg\phi$ is true.

The reasoning for the case where ϕ is true is in Tennant's own words; the reasoning for the case where ϕ is false in effect merely substitutes $\neg\phi$ for ϕ throughout that reasoning. By (DIS), we can conclude that our mere possession of the decision procedure for ϕ

entitles us to assert (11). But it is utterly implausible to claim that whenever ϕ is decidable, someone sometime will know whether ϕ holds, in the usual sense of ‘know’. Mere possession of the decision procedure does not entitle us to assert that anyone will ever have that knowledge. For in advance of applying the procedure, we may have no reason to think that it will ever be applied; indeed, we may have reason to think that it will never be applied. Perhaps its application is costly in time and other scarce cognitive resources and ϕ is a proposition whose truth-value is unlikely ever to be of interest to anyone. Moreover, it may be unlikely that anyone will ever come to know whether ϕ holds without applying such a decision procedure. Alternatively, we may know that the meteorite is about to strike.

Let us provisionally accept the Tennant-style argument for the proposition that if ϕ is decidable then we have a method whose application will either verify $\mathbf{K}\phi$ or verify $\mathbf{K}\neg\phi$. Nevertheless, our mere possession of that method does not entitle us to assert $\mathbf{K}\phi \vee \mathbf{K}\neg\phi$. Thus the problem lies with (DIS). What has gone wrong is that the application of the decision procedure for ϕ *brings about* the state of affairs expressed by (11). That is why our mere possession of the method is not enough. To assert (11), we need some reason to think that someone sometime will apply the method. Exactly the same problem affects Tennant’s assumption that we are entitled to assert (10) whenever ϕ is decidable. For, if ϕ is true, the application of the decision procedure for ϕ brings about the state of affairs expressed by $\mathbf{K}\phi$. Our mere possession of the method, in advance of actually applying it, does not entitle us to assert (10). If, as Tennant assumes, our possession of a decision method for ψ always entitles us to assert $\psi \vee \neg\psi$, then it has not been shown that our possession of a decision method for ϕ puts us in possession of a

decision method for $\mathbf{K}\phi$. Alternatively, if we did count the Tennant-style procedure as a decision method for $\mathbf{K}\phi$, then it has not been shown that mere possession of a decision method in that weak sense for $\mathbf{K}\phi$ would entitle us to assert $\mathbf{K}\phi \vee \neg\mathbf{K}\phi$, and Tennant's argument would still fail because we could not bridge the gap from (7) to (8).¹⁴

For a more dramatic example of the fallacy, consider a paradigm of a potentially undecidable sentence, 'A city was, is or will be built on this spot' (compare Dummett 1959). Assume that no city has ever been built on the spot, and there is no present plan to build one, but equally no special reason why one should never be built there. For the Dummettian anti-realist, we are not entitled to assert 'Either a city was, is or will be built on this spot or no city was, is or will be built on this spot', for we have no procedure for determining which disjunct holds. Now imagine someone claiming:

We do have a decision method for the sentence. For you can in principle build a city on this spot. Having done so, you will have determined that a city was, is or will be built on this spot.

Although we have the capacity in principle to build a city on the spot, it does not put us in possession of a decision procedure in the relevant sense, for by intuitionistic standards it does not entitle us to assert 'Either a city was, is or will be built on this spot or no city was, is or will be built on this spot' in advance of exercising the capacity. Indeed, the problem for (DIS) is even worse, for we can take both ϕ and ψ in it to be 'A city was, is or will be built on this spot'. We possess a method (building a city) whose application will verify 'A city was, is or will be built on this spot'. Therefore, by (DIS) we are now

entitled ‘Either a city was, is or will be built on this spot or a city was, is or will be built on this spot’. Since even in intuitionistic logic $\phi \vee \phi$ is trivially equivalent to ϕ , we are now entitled to assert ‘A city was, is or will be built on this spot’, in advance of building one or even planning to do so. That is absurd. More generally, such an argument would conclude that whenever one has the power in principle to make ϕ true, one is entitled to assert ϕ in advance of exercising that power or even planning to exercise it. That conclusion involves one in contradictions, since one often has both the power to make ϕ true and the power to make $\neg\phi$ true, for example when ϕ is ‘I shall count to a thousand by midnight’.¹⁵ By (DIS), one is now both entitled to assert ϕ and entitled to assert $\neg\phi$, irrespective of one’s intentions. Evidently, (DIS) can fail for sentences whose truth-values depend on our will; more specifically, it fails when the truth-value of ϕ or of ψ depends on whether the method is actually applied. Of course, this problem does not arise when ϕ and ψ are mathematical.

If (DIS) does not govern the assertibility of disjunctions, what does? Intuitionistic semantics relies on some notion of canonical verification (Dummett 1977: 389-403). One is entitled to make an assertion for which one lacks a canonical verification, if one knows that such a verification exists. The existence of the verification does not consist in its being possessed by anyone; nevertheless, since the intuitionist conceives it as essentially *capable* of being possessed by someone, it does not import any sort of platonism inconsistent with the intuitionistic view.¹⁶ The natural suggestion is then that a canonical verification of a disjunction consists of a canonical verification of a disjunct. We may be entitled to assert a disjunction without being entitled to assert any disjunct, because we know that a canonical verification exists for some disjunct without knowing which. We

are in that position with respect to $\phi \vee \neg\phi$ when we have a genuine decision procedure for ϕ but have not yet applied it. But if we possess the purported decision procedure for $\mathbf{K}\phi$ without having applied it, we do not thereby know that a canonical verification for $\mathbf{K}\phi$ or a canonical verification for $\neg\mathbf{K}\phi$ exists, even though we know how to bring such a canonical verification into existence. Consequently, we do not know that a canonical verification for $\mathbf{K}\phi \vee \neg\mathbf{K}\phi$ exists; so by intuitionistic standards we are not entitled to assert the disjunction.

Tennant's response to this critique is to insist on a different reading of 'K' (2001a: 273-7). In his terminology, he understands 'know' to mean *virtually or implicitly know* rather than *occurrently know*, with a corresponding difference in the interpretation of 'K' (2001b: 109). In this sense, in merely possessing a decision method for primeness we already know whether any given natural number is prime. Unfortunately, he never squarely faces the obvious problem that this makes his discussion *prima facie* quite irrelevant to the point at issue, namely, the stability of the moderately hard anti-realist response to Fitch as I proposed it, which involves a logical distinction between (7) and (8) for decidable ϕ on the usual reading of 'know', not Tennant's. We have just seen that his arguments do not work on the usual reading.

To make his arguments relevant, Tennant would have to show that the usual reading of 'know' was somehow unavailable to the moderately hard anti-realist. The boldest, least plausible strategy would be to argue that such a reading makes no sense. But it is hopeless for the anti-realist to pretend that actually applying a decision procedure makes no cognitive difference at all. Indeed, Tennant describes in his own terms what cognitive difference it makes:

The whole point of having a decision procedure is to discover the canonical form of expression of a proposition that, at the outset, can be identified only by description: as *the result of applying the decision procedure*. (2001a: 276)

We are supposed to discover the canonical form by applying the decision procedure. We thereby discover the canonical form only if we did not already know what it was. But in Tennant's special sense we *did* already know what it was, since we had a procedure for finding out. Thus the point of applying the procedure is to gain knowledge in the *usual* sense of the canonical form. Consequently, Tennant himself relies on the coherence of something like the usual reading.¹⁷ Indeed, he speaks of it as of a genuine sense of 'know' (2001a: 276, 2001b: 109). Although he writes favourably of a conceptual reform that would impose his special reading on 'know', not all moderately hard anti-realists need feel bound by such a reform; moreover, eliminating all unreformed epistemological terminology would leave us unable to articulate the point of applying a decision procedure.

Tennant points out that his special reading of 'know' is better suited to epistemic logics that assume logical omniscience. He claims that 'if this idealizing assumption were disallowed, the original Fitch argument would not go through' (2001a: 274). But that is to ignore crucial differences. The only aspect of logical omniscience used in the argument is principle (3), that knowledge of a conjunction implies knowledge of its conjuncts. But one can accept that very weak closure principle even for knowledge in the usual sense without accepting logical omniscience in general. The modest idea that in knowing a

conjunction one knows its conjuncts does not commit one to the extravagant idea that in knowing anything one knows anything that it entails. Moreover, there are revisions of the Fitch argument that do not even rely on (3).¹⁸ Nothing that Tennant says compels the moderately hard anti-realist to formulate their knowability principle (1) in terms of a reading of ‘know’ that satisfies logical omniscience.

Thus the effect of Tennant’s insistence on his special reading of ‘know’ is that his arguments completely fail to engage with the version of moderately hard anti-realism that he was supposed to be attacking. Indeed, they fail to engage with just about any form of anti-realism that endorses the knowability principle (1) on the usual reading of ‘know’, since the most pertinent version of the Fitch argument will then involve that reading throughout. To refrain from endorsing the principle that every truth is capable of being known in the usual sense is to take a significant step back from full-blooded anti-realism.

Do Tennant’s arguments establish something of interest on his preferred reading of ‘know’, even though they do not establish what they were supposed to? On such a reading, (11) is no longer obviously absurd when ϕ is decidable but not occurrently decided. His purported decision procedure for $\mathbf{K}\phi$ is at much less risk of violating the constraint, which he accepts, that ‘[p]roper decision procedures do not interfere with the states of affairs’ that they are supposed to determine (2001a: 276-7). Nevertheless, crucial unclarities remain.

First, Tennant’s argument that the decidability of ϕ implies the decidability of $\mathbf{K}\phi$ appears to commit him to a version of the highly controversial principle that when one knows, one knows that one knows. For the key passage is this:

If you [...] determine that φ is true, then you know that φ . So you have determined that $\mathbf{K}\varphi$ is true. (1997: 262)

The principles invoked in the two sentences can be formalized by schemas (D1) and (D2) respectively, with $\mathbf{D}y$ for determination by you as true and $\mathbf{K}y$ for your knowledge:

$$(D1) \quad \mathbf{D}y\varphi \rightarrow \mathbf{K}y\varphi$$

$$(D2) \quad \mathbf{K}y\varphi \rightarrow \mathbf{D}y\mathbf{K}\varphi$$

By substitution of ' $\mathbf{K}\varphi$ ' for ' φ ' in (D1):

$$(D3) \quad \mathbf{D}y\mathbf{K}\varphi \rightarrow \mathbf{K}y\mathbf{K}\varphi$$

By transitivity from (D2) and (D3):

$$(D4) \quad \mathbf{K}y\varphi \rightarrow \mathbf{K}y\mathbf{K}\varphi$$

In other words, if you know something, you know that someone sometime knows it. But that principle is extremely doubtful, even on a special reading of 'know' that satisfies logical omniscience (Williamson 2000a: 114-34). In terms of epistemic logic: a knowledge operator can satisfy deductive closure without satisfying the S4 principle. An

argument that relies on assumptions that are jointly as strong as (D4) is some way from establishing its conclusion.

Second, it is not altogether clear what Tennant means by ‘decidable’. He is attracted by a version of moderately hard anti-realism different from the one I envisaged. On his version, one accepts schema (7) and rejects schema (8), but accepts this restricted version of (8) (2001a: 266-7):

$$(8a) \quad [\varphi \wedge (\varphi \text{ is decidable})] \rightarrow \mathbf{K}\varphi$$

If one accepts the Tennant-style argument from the decidability of φ to (11) on his special reading, (8a) is a simple consequence given principle (2) (the factivity of knowledge), since $\mathbf{K}\neg\varphi$ implies $\neg\varphi$. Clearly, if (8a) is not to collapse into (8), one must reject:

$$(8b) \quad \varphi \rightarrow (\varphi \text{ is decidable})$$

For of course (8a) and (8b) jointly entail (8). Thus ‘ φ is decidable’ cannot be regarded as a notational variant of $\varphi \vee \neg\varphi$, as it often is in intuitionistic writings, for that would immediately validate (8b). Moreover, that reading disqualifies undecidability as a genuine option, as Tennant takes it to be (2001a: 266), since $\neg(\varphi \vee \neg\varphi)$ is intuitionistically inconsistent. At first sight, it is unclear how an intuitionist can reject (8b). For suppose that we have a proof of φ . Then that proof decides φ , and thereby constitutes a proof of its decidability. Thus (8b) seems to be intuitionistically provable,

because we can transform any proof of its antecedent into a proof of its consequent. Presumably, the way to block (8b) is to insist that its consequent concerns our actual possession of a decision method, not the mere possibility in principle of possessing one. Attempts to validate (8b) by means of the intuitionistic semantics for the conditional can then be blocked like similar attempts to validate (8).¹⁹ But then a more than virtual notion of possessing a decision procedure is doing the crucial work. To clarify his envisaged version of moderately hard anti-realism, Tennant needs to explain how it is supposed to reconcile virtual and non-virtual aspects of cognition.

Tennant's alternative version of moderately hard anti-realism remains undeveloped. On an understanding of the language relevant to the version that I mooted, the decidability of ϕ does not imply the decidability of $\mathbf{K}\phi$ in any sense conducive to (10). Tennant fails in his attempt to collapse (7) into (8) for decidable ϕ . He has located no instability in the envisaged version of moderately hard anti-realism.

III

Tennant also proposes another response to Fitch's argument, this time a soft anti-realist one, by constructing and defending a modified knowability principle. Having defined a sentence ϕ to be *Cartesian* if and only if the contradiction \perp does not follow from $\mathbf{K}\phi$, he endorses this restricted variant of (1), formulated as a rule of inference:²⁰

(\diamond KC) ϕ ; ergo $\diamond\mathbf{K}\phi$, where ϕ is Cartesian

Informally: (\Diamond KC) says that truth entails knowability except when Fitch's problem occurs. Tennant admits that it may be undecidable whether a given step is an instance of a rule like (\Diamond KC), because it is undecidable whether \perp follows from $\mathbf{K}\phi$. He claims that this does not matter, on the grounds that we will apply it only when we do know that the condition is met. Since (\Diamond KC) is intended for use only in a few philosophical arguments, not for systematic application in mathematics or science, the undecidability is supposed not to defeat its purpose.

Tennant's rule looks desperately *ad hoc*. He replied in detail (2001b) to a similar charge from Michael Hand and Jonathan Kvanvig (1999). However, his reply depends on a misunderstanding of the nature of the charge. He writes as though what is wrong with *ad hoc* principles is that they are restricted to the point of total or partial triviality. For instance, he argues that the restriction in (\Diamond KC) is analogous to restrictions in other principles that are nevertheless 'substantive, informative and important' (2001b: 110, 111, 113).²¹ He seems to assume that if a principle P* is more restricted than a principle P, then P is *ad hoc* only if P* is also *ad hoc*.²² He assumes that '[a]d hoc emendations to general laws in natural science' detract from their applicability (2001b: 112). However, imagine a scientist who has always maintained that emeralds are always green, but at time t is confronted with a blue emerald, and adopts the revised theory that emeralds are always grue, where something is grue at a time if and only if either it is green and the time is before t or it is blue and the time is not before t . The gruesome, gerrymandered revised theory is clearly *ad hoc*, even though it is just as general, substantive, informative and important (if correct) as the old theory. The problem is not triviality but ill-motivated

complexity. Although the new theory predicts the same data before t as the old theory, and improves on the old theory with respect to the datum at t , the previous evidential support for the old theory does not transfer to the new theory, even before counterexamples to the new theory emerge. Analogous problems face $(\diamond KC)$. If Fitch's argument forces anti-realists to restrict the original knowability principle (1), then something is wrong with their original meaning-theoretic arguments for (1). Until we have an adequate diagnosis of the fallacies in those arguments, we cannot assume that such considerations confer any support whatsoever on $(\diamond KC)$ or any other attempted approximation to (1) that does not immediately succumb to the Fitch argument. A subtle fallacy in an argument can easily mean that it establishes nothing of interest whatsoever. $(\diamond KC)$ is a gruesome principle.

We need not dwell on the charge that $(\diamond KC)$ is *ad hoc*, for that is not the worst of its problems. The point of restricting $(\diamond KC)$ to Cartesian cases is to enable the soft anti-realist to avoid asserting (8), that something is true only if known, even for decidable sentences. Since Tennant holds that (7) collapses into (8) when φ is decidable, his soft anti-realist rejects (7) too. The restriction on $(\diamond KC)$ blocks the original derivation of (7) from (1). But Tennant overlooked a more complex derivation of (7) from $(\diamond KC)$ and some plausible assumptions.

The argument will be conducted in Tennant's preferred background logic, which is not the standard intuitionistic one but his weaker system IR of intuitionistic relevance logic (Tennant 1997: 343-4). The differences do not affect the arguments below. For definiteness, let φ be the decidable sentence 'There is a fragment of Roman pottery at that spot' (we assume a suitable context). Introduce a proper name ' \mathbf{n} ' by the stipulation that

it is to designate (rigidly) the number of books actually now on my table. Thus ‘**n**’ is not a numeral such as ‘9’ but rather a name whose reference is fixed by an empirical description. Let ‘**E**’ be the predicate ‘is even’.

We first argue that the conjunction $\phi \wedge (\mathbf{K}\phi \rightarrow \mathbf{E}\mathbf{n})$ is Cartesian. For suppose that $\mathbf{K}(\phi \wedge (\mathbf{K}\phi \rightarrow \mathbf{E}\mathbf{n}))$ is inconsistent, in the sense that \perp follows from $\mathbf{K}(\phi \wedge (\mathbf{K}\phi \rightarrow \mathbf{E}\mathbf{n}))$ according to the logic adverted to in the definition of ‘Cartesian’. Then this story contains an inconsistency:

STORY I find a fragment of Roman pottery at this spot and identify it correctly; I thereby come to know that there is a fragment of Roman pottery there. I also count the books actually now on my table and discover that the number is even; I deduce that if someone sometime knows that there is a fragment of Roman pottery at that spot then **n** is even.²³ By putting the two pieces of knowledge together, I acquire the knowledge expressed by the conjunction $\phi \wedge (\mathbf{K}\phi \rightarrow \mathbf{E}\mathbf{n})$. Thus $\mathbf{K}(\phi \wedge (\mathbf{K}\phi \rightarrow \mathbf{E}\mathbf{n}))$ holds.

But STORY is obviously consistent; we cannot exclude its truth on purely logical grounds (just try!). Of course, **n** may *in fact* be odd, in which case, since that number could not have been even, STORY expresses an impossible state of affairs. Nevertheless, STORY itself is still *consistent*; we cannot discern by reason alone that the description which fixes the reference of ‘**n**’ picks out an odd number. Someone who asserts **E****n** because he failed to see one of the books is not guilty of an inconsistency. Although we might produce an inconsistent story by substituting for ‘**n**’ throughout STORY a numeral

with the same reference as ‘**n**’, it does not follow that STORY itself is inconsistent.

More precisely, the result of substituting a coreferential numeral for ‘**n**’ in the sentence $\mathbf{K}(\varphi \wedge (\mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n}))$ is a different sentence; the inconsistency of the latter does not imply the inconsistency of the former. Thus \perp does not follow from $\mathbf{K}(\varphi \wedge (\mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n}))$, so $\varphi \wedge (\mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n})$ is Cartesian.

For the rest of the argument, let \vdash be the consequence relation of a system of modal epistemic logic based on IR with the additional rule ($\diamond\mathbf{K}\mathbf{C}$) and the axiom schemas (2) and (3).²⁴ Since its condition is met in this case, ($\diamond\mathbf{K}\mathbf{C}$) gives:

$$(12) \quad \varphi \wedge (\mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n}) \vdash \diamond\mathbf{K}(\varphi \wedge (\mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n}))$$

Moreover:

$$(13) \quad \varphi \wedge \neg\mathbf{K}\varphi \vdash \varphi \wedge (\mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n})$$

For $\neg\alpha \vdash \alpha \rightarrow \beta$ holds even in IR (Tennant 1997: 344); from $\neg\mathbf{K}\varphi \vdash \mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n}$ we can derive (13) by the rules for \wedge . Now (12) and (13) yield:

$$(14) \quad \varphi \wedge \neg\mathbf{K}\varphi \vdash \diamond\mathbf{K}(\varphi \wedge (\mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n}))$$

The next step is a Fitch-like argument for:

$$(15) \quad \mathbf{K}(\varphi \wedge (\mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n})) \vdash \mathbf{E}\mathbf{n}$$

For (3) and \wedge -elimination yield $\mathbf{K}(\varphi \wedge (\mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n})) \vdash \mathbf{K}\varphi$, while (2) and \wedge -elimination yield $\mathbf{K}(\varphi \wedge (\mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n})) \vdash \mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n}$, and even in IR we can then move to (15).²⁵

Since the rules used to derive (15) are truth-preserving in all possible situations, not just the actual one, if the premise of (15) expresses a possibility, so does its conclusion (if $\alpha \rightarrow \beta$ is a theorem of a normal modal logic, so is $\diamond\alpha \rightarrow \diamond\beta$):

$$(16) \quad \diamond\mathbf{K}(\varphi \wedge (\mathbf{K}\varphi \rightarrow \mathbf{E}\mathbf{n})) \vdash \diamond\mathbf{E}\mathbf{n}$$

Now (14) and (16) yield:

$$(17) \quad \varphi \wedge \neg\mathbf{K}\varphi \vdash \diamond\mathbf{E}\mathbf{n}$$

Uncontentiously, it is not contingent whether \mathbf{n} is even. Since I can count the books on my table, it is decidable whether \mathbf{n} is even; hence \mathbf{n} is either odd or even. But if \mathbf{n} is odd, it could not have been even, for the mathematical properties of numbers are not contingent. Thus \mathbf{n} could have been even only if it is even. We can symbolize that as the argument $\mathbf{E}\mathbf{n} \vee \neg\mathbf{E}\mathbf{n}, \neg\mathbf{E}\mathbf{n} \rightarrow \neg\diamond\mathbf{E}\mathbf{n}, \diamond\mathbf{E}\mathbf{n} \vdash \mathbf{E}\mathbf{n}$, since ‘ \mathbf{n} ’ is a rigid designator.²⁶ Thus, treating the uncontentious auxiliary assumptions concerning $\mathbf{E}\mathbf{n}$ as part of the background logic, we can strengthen (17) to:

$$(18) \quad \varphi \wedge \neg\mathbf{K}\varphi \vdash \mathbf{E}\mathbf{n}$$

Since the two cases are symmetric, we can now repeat the argument for (18) with ‘odd’ in place of ‘even’ to derive:

$$(19) \quad \varphi \wedge \neg \mathbf{K}\varphi \vdash \neg \mathbf{E}\mathbf{n}$$

But (18) and (19) together yield:

$$(20) \quad \vdash \neg(\varphi \wedge \neg \mathbf{K}\varphi)$$

That is to make (7) a theorem. But the point of Tennant’s restricted knowability principle ($\diamond\mathbf{KC}$) was precisely to enable the soft anti-realist not to assert (7) for decidable φ (as in the present case), since on Tennant’s view (7) collapses into (8), the principle definitive of hard anti-realism, for decidable φ . Thus Tennant’s restriction is futile.

Evidently, the foregoing critique of Tennant’s principle ($\diamond\mathbf{KC}$) depends on careful respect for the distinction between the logical notion of *inconsistency* and the metaphysically modal notion of *impossibility*. One or other of $\mathbf{E}\mathbf{n}$ and $\neg\mathbf{E}\mathbf{n}$ expresses a metaphysical impossibility, but each of them is logically consistent, since the reference of ‘ \mathbf{n} ’ is fixed empirically. Similarly, the sentence ‘George W. Bush = Tony Blair’ is logically consistent, even though it expresses an impossibility. Unfortunately, Tennant’s reply to the critique displays a startling insensitivity to the distinction. In expounding ($\diamond\mathbf{KC}$), he writes:

It should be clear to anyone with a sympathetic understanding of the spirit of the proposed restriction that for a proposition to be Cartesian one ought to be unable to derive absurdity from it *modulo* any necessarily true propositions. It is a logical convention of long standing that mention of theorems as premises can be suppressed. (2001b: 264)

This passage conflates necessary truth and theoremhood. Of course, given the cut rule, the use of theorems of a given logic as premises of derivations in that logic does not enable one to reach any conclusions that could not be reached without those premises. But it is certainly not ‘a logical convention of long standing’ that mention of *necessary truths* as premises can be suppressed. The undecidable Gödel sentence for first-order arithmetic is a necessary truth, but that does not mean that mention of it as a premise can be suppressed in the proof theory of first-order arithmetic. Similarly, one cannot declare ‘George W. Bush = Tony Blair’ logically inconsistent just on the grounds that its negation expresses a necessary truth.

Nevertheless, Tennant is quite explicit in his notion of a Cartesian proposition that:

To say that absurdity is not derivable from $\mathbf{K}\phi$ is equivalent to saying that absurdity is not derivable from $\mathbf{K}\phi$ in conjunction with any set \mathbf{X} of necessarily true propositions. (2001a: 269-70)

We had best understand Tennant as defining ‘Cartesian’ in terms of a special consequence relation for which, by stipulation, all necessities are theorems. For the time being let us read his term ‘proposition’ as equivalent to ‘sentence’, since elsewhere he treats the constituents of arguments in the logically standard way as linguistic (1997: 313-15); in his reply (2001a) he does not object to my treatment of premises and conclusions as sentences. We will reconsider the talk of propositions later.

According to Tennant, ‘We are dealing primarily with *logico-mathematical* possibility and necessity here’ (2001a: 269). On Tennant’s proposal, if \mathbf{n} is odd then $\neg\mathbf{En}$ is a logico-mathematical necessity because ‘ \mathbf{n} ’ is a rigid designator; since \mathbf{En} follows from $\mathbf{K}(\phi \wedge (\mathbf{K}\phi \rightarrow \mathbf{En}))$ by (15), absurdity is derivable from $\mathbf{K}(\phi \wedge (\mathbf{K}\phi \rightarrow \mathbf{En}))$ in conjunction with the necessity $\neg\mathbf{En}$; thus $\phi \wedge (\mathbf{K}\phi \rightarrow \mathbf{En})$ is not Cartesian after all. By parallel reasoning, if \mathbf{n} is even then $\phi \wedge (\mathbf{K}\phi \rightarrow \neg\mathbf{En})$ is not Cartesian. Either way, one half of my argument is supposed to break down. Tennant seems to assume that logico-mathematically necessary truths are knowable *a priori*, for he glosses the account quoted above of what it is for absurdity not to be derivable from $\mathbf{K}\phi$ thus:

Whether this definition calls for the consideration only of sets \mathbf{X} all of whose members are knowable *a priori*, or calls for the consideration also of sets \mathbf{X} some of whose members might be knowable only *a posteriori*, is an issue of principle on which we are not at present forced to take a stand. (2001a: 270)

Tennant is not forced to take a stand on the issue of principle only if he is entitled to assume that if \mathbf{n} is odd then $\neg\mathbf{En}$ is knowable *a priori* and if \mathbf{n} is even then \mathbf{En} is

knowable *a priori*. But recall that the reference of ‘**n**’ was fixed by the description ‘the number of books actually now on my table’. Thus we cannot know *a priori* whether **n**, so presented, is even! That was the crux of my argument. **En** and $\neg\mathbf{En}$ are not sentences of the language of mathematics, because ‘**n**’ is not a term of that language: although it refers to a number, it does so in a non-mathematical way. Tennant describes ‘**n** is even’ as ‘a *mathematical proposition*’ on the grounds that ‘**n**’ is a rigid designator (2001a: 270), but it is unclear what he means by ‘mathematical proposition’. At any rate, his discussion ignores the significance of the empirical way in which the reference of ‘**n**’ was fixed.

In order to make Tennant’s discussion relevant to the argument that I presented, we should understand him as defining ‘Cartesian’ in terms of a special consequence relation for which, by stipulation, all necessary truths may occur as premises in the derivation of absurdity, whether or not they are knowable *a priori*. The result does not constitute a formal system, but never mind. Consider the following variant of ($\diamond\mathbf{KC}$):

($\diamond\mathbf{KC}^*$) φ ; *ergo* $\diamond\mathbf{K}\varphi$, where $\neg\Box\neg\mathbf{K}\varphi$ holds

We can derive ($\diamond\mathbf{KC}$) from ($\diamond\mathbf{KC}^*$) by showing that if φ is Cartesian, $\neg\Box\neg\mathbf{K}\varphi$ holds. Suppose that $\Box\neg\mathbf{K}\varphi$ holds. Then $\neg\mathbf{K}\varphi$ is permitted to occur as a premise in the derivation of absurdity, in the sense used in the definition of ‘Cartesian’. So absurdity is derivable from $\mathbf{K}\varphi$ in that sense. Therefore φ is not Cartesian. Thus if $\Box\neg\mathbf{K}\varphi$ holds, φ is not Cartesian. By an intuitionistically valid form of contraposition, if φ is Cartesian, $\neg\Box\neg\mathbf{K}\varphi$ holds. Thus ($\diamond\mathbf{KC}$) is a simple consequence of ($\diamond\mathbf{KC}^*$); one can also argue for the converse, although that is not our present concern. But ($\diamond\mathbf{KC}^*$) is not a distinctively anti-

realist principle. For a realist who accepts classical modal logic, $(\Diamond KC^*)$ is trivially truth-preserving, since \Diamond is equivalent to $\neg\Box\neg$ simply by the duality of the two modal operators. Of course, it is odd to present $\Diamond K\phi$ as derived from the ostensible premise ϕ rather than from the condition for the applicability of $(\Diamond KC^*)$, that $\neg\Box\neg K\phi$ holds, which is what really guarantees the truth of where $\Diamond K\phi$. But that oddity is harmless once one appreciates that the connection signalled by ‘*ergo*’ outside the scope of the ‘where’ clause need not be knowable *a priori*. For exactly the same reason, Tennant’s principle $(\Diamond KC)$ as now interpreted, a corollary of $(\Diamond KC^*)$, is also not distinctively anti-realist; for a realist who accepts classical modal logic, it is trivially truth-preserving. The two principles may not be quite so innocent from an intuitionistic perspective, since the inference $\neg\Box\neg\alpha$ to $\Diamond\alpha$ is structurally analogous to the intuitionistically invalid inference from $\neg\forall x\neg\alpha$ to $\exists x\alpha$. However, that extra piece of logical content from an intuitionistic perspective is, if anything, a slight concession to classicism, not the articulation of a distinctively anti-realist claim about the possibility of knowledge. Thus Tennant’s stipulations about the notion of derivability in the definition of ‘Cartesian’ are relevant to my original objection to $(\Diamond KC)$ only by voiding $(\Diamond KC)$ of all interest as a formulation of an anti-realist principle of knowability.

I had in mind considerations of the kind above when I wrote in my original critique ‘A more liberal interpretation of inconsistency might trivialize $\Diamond KC$; it is not what Tennant intends’ (2000b: 110). Perhaps it was what Tennant intended, and he has indeed fallen into the trap that I warned against.

Does it make a difference if we suppose that by ‘proposition’ Tennant means something significantly more coarse-grained than a sentence, even though he does not

mention the distinction between sentences and propositions in this connection? Let ‘**q**’ be a numeral with the same reference as ‘**n**’. If proper names are directly referential, then the two sentences **Eq** and **En** express the same proposition, even though the sentences do not have the same cognitive significance for us. If knowledge is a relation to propositions, it follows that knowing **Eq** *a priori* is knowing **En** *a priori*, although one can be in a position to express one’s knowledge by one sentence without being in a position to express it by the other. Such a view might fit what Tennant says in his discussion of decidability about different forms of expression of a given proposition (2001a: 276-7). Would this approach enable Tennant to restrict the use of necessarily true propositions in the derivation of absurdity to those knowable *a priori* under some mode of presentation or other (such as **Eq**)?

The appeal to a directly referential semantics for proper names deals with only one class of examples. An analogue of my argument can be developed for any decidable sentence ϕ whatsoever, by substituting for **En** the sentence **A ϕ** , where **A** is the rigidifying ‘actually’ operator, so that **A ϕ** is *a priori* equivalent to ϕ , but $\diamond\mathbf{A}\phi$ entails **A ϕ** and $\diamond\neg\mathbf{A}\phi$ entails $\neg\mathbf{A}\phi$. To extend the approach just sketched to such examples, Tennant would need to argue that **A ϕ** expresses the same proposition as some sentence that one can use to express *a priori* knowledge. It is hard to see what principled justification there could be for such a claim, short of the identification of all necessarily equivalent propositions. But on that approach there is just one necessary truth, which is known *a priori* under the mode of presentation ‘ $0 = 0$ ’. We can then reproduce the derivation of (\diamond KC) from (\diamond KC*), and regain just the trivialization of Tennant’s principle that the appeal to coarse-grained propositions was supposed to avoid.

Thus Tennant's response to the critique of (\diamond KC) serves only to emphasize his difficulties; his attempt to locate a fallacy in it merely trivializes his own view.

Should Tennant's soft anti-realist seek, instead of (\diamond KC), a knowability principle with a stronger restriction to avoid (7)? The natural suspicion is that such a restriction would have to be very draconian indeed, and thereby risk trivialization again. But the search is in any case ill-motivated, for reasons already indicated in the discussion of Tennant's unsuccessful response to the charge that (\diamond KC) is *ad hoc*. The original knowability principle (1) was the outcome of an anti-realist argument (albeit a very dubious argument). If (1) has false consequences, then something must be wrong with the argument. If the argument is irreparably fallacious, one has lost one's reason for postulating any knowability principle at all, restricted or unrestricted. If the argument can be repaired, the nature of the repairs should dictate the nature of the restrictions on the resultant knowability principle.²⁸ Tennant does not indicate any form of argument for anti-realism that would motivate a principle restricted in the manner of (\diamond KC) on a non-trivializing interpretation. To say that ϕ is non-Cartesian is not to explain on anti-realist terms how ϕ could be unknowably true, how speakers' use of ϕ could be sensitive to a condition they could not in principle recognize to obtain or how ϕ could express its content without such sensitivity; it is merely to say that broadly logical considerations do (not not) rule out knowledge of its truth. Without such an explanation, from the perspective of principled anti-realism it is quite premature to endorse anything like (\diamond KC).

If Fitch's argument does not by itself refute all forms of anti-realism, it certainly shows how much would have to be done before there was a working anti-realist

semantics for empirical language, even in the toy examples that we have been considering. The attempts on behalf of anti-realism to deal with the Fitch problem give every sign of a degenerating research programme.²⁹

Notes

- 1 Where appropriate, sentences or paragraphs from Williamson 2000 have been absorbed into the present text.
- 2 Tennant briefly considers other possible readings of **K**. He complains (1997: 270) that ‘Williamson [...] appears not to have anticipated the possibility’ of interpreting **K** in Fitch’s arguments as ‘it is known at t that’ for a particular time t . He has overlooked the discussions of such readings of the argument at Williamson 1982: 204, 1988: 425-8 and 1994: 141, 144. It is unnecessary to add to them here.
- 3 For difficulties facing the attempt to evade Fitch’s argument by rejecting (3) see Williamson 1993 and 2000a: 275-85.
- 4 It is sometimes claimed that one can meet Dummett’s demand simply by treating the notion of truth in a truth-conditional compositional semantics for empirical discourse as verifiability. That is a mistake. The key notion in the intuitionistic compositional semantics for mathematical language is ‘ Π is a proof of φ ’, *not* ‘ φ is provable’ (‘Something is a proof of φ ’); for example, the semantic clause for \rightarrow concerns the transformability of proofs of the antecedent into proofs of the consequent, which makes no sense in terms of an undifferentiated notion of provability. Thus the key notion in an analogous verification-conditional compositional semantics for empirical discourse is ‘ Π is a verification of φ ’, *not* ‘ φ is verifiable’. The truth-conditional clause for negation, ‘ $\neg\varphi$

is true if and only if ϕ is not true' (or something similar), cannot be interpreted as a constraint on verifications, for if something is not a verification of ϕ , it does not follow that it is a verification of $\neg\phi$. The idea of Dummett's original argument is that the key notion in a compositional semantics should be one to which speakers' use is sensitive, and that their use is sensitive to a given condition in a given context only if they can decide in that context whether it obtains. On such a view, they can decide in a given context whether they have a proof or verification of ϕ in that context (for an argument against this decidability claim, see Williamson 2000a: 110-13), but no notion of provability or verifiability that would constitute a not wildly subjectivist notion of truth is decidable within the limitations of every given speech context. Thus no notion of provability or verifiability that would constitute a not wildly subjectivist notion of truth meets Dummett's constraints on the key notion in a compositional semantics, even if it is the existential generalization of a notion of proof or verification that does meet Dummett's constraints.

5 Dummett's own response to Fitch (2001) does not appeal to intuitionistic logic; rather, it restricts the knowability principle (1) to atomic sentences. This restriction is hard to reconcile with Dummett's original motivation for the knowability principle, a motivation that applies to complex sentences just as much as to atomic ones. It will not do to say that the use of complex sentences is indirectly epistemically grounded because their atomic constituents are. For connectives such as conjunction and negation are used as constituents of complex sentences, not by themselves. Thus any epistemic grounding of the use of connectives must derive from an epistemic grounding of complex sentences

in which they occur, not vice versa: yet Dummett's strategy against Fitch is just to avoid any such direct epistemic grounding of the use of complex sentences. Thus his anti-realism unravels. Note also that his original (1959) examples of sentences that the realist contentiously treated as verification-transcendent involved complex constructions such as universal quantification and the counterfactual conditional: 'A city will never be built on this spot' and 'If Jones had encountered danger, he would have acted bravely' are not atomic sentences. See Brogaard and Salerno 2002 (some points of which are anticipated at Williamson 1990: 300) and Tennant 2002 for criticism of Dummett's response to Fitch and Rosenkranz 2004 for more discussion.

6 Williamson 1992 proves model-theoretically that schema (8) is not derivable from schemas (1)-(7) and (9) even in a very strong system of intuitionistic modal epistemic logic.

7 See Williamson 1982: 206. Tennant (1997: 267-8) objected that $\neg(\phi \rightarrow \mathbf{K}\phi)$ is not intuitionistically inconsistent given (1) (and the other principles used to derive (9)) unless it intuitionistically implies both ϕ and $\neg\mathbf{K}\phi$, and that, intuitionistically, although it implies $\neg\mathbf{K}\phi$ and $\neg\neg\phi$ it does not in general imply ϕ (it does in the special case when ϕ is decidable, but then $\phi \vee \neg\phi$ holds and the analogy is not useful). The objection rests on an error. Intuitionistically, $\neg\mathbf{K}\phi$ and $\neg\neg\phi$ imply $\neg(\neg\mathbf{K}\phi \rightarrow \neg\phi)$, the negation of (9); thus $\neg(\phi \rightarrow \mathbf{K}\phi)$ is intuitionistically inconsistent given (1) (and the other principles used to derive (9)), although (without those principles) it does not intuitionistically imply ϕ . Tennant (2001a: 277-9) concedes and amplifies this criticism of his objection.

8 This adapts Tennant's terminology (1997: 261).

9 For related points see Percival 1990. Other relevant discussions of Fitch's argument in the context of intuitionistic logic include Wright 1993: 427-30, Cozzo 1994, Pagin 1994, Usberti 1995: 65-6, 121-8 and DeVidi and Solomon 2001.

10 Tennant (1997: 276-8) also rejects the attempt in Edgington 1985 to reconstrue the knowability principle by means of something like an 'actually' operator, for reasons given in Williamson 1987 and 2000a: 290-301 and Wright 1993: 426-32. See also Percival 1991. Edgington's idea is developed rigorously by Rabinowicz and Segerberg 1994, Lindström 1997 and Rückert 2003, but none of these papers fully answers the philosophical objections.

11 Tennant claims that 'the validity of $\neg(\phi \wedge \neg\mathbf{K}\phi)$ guarantees the validity of $\phi \rightarrow \mathbf{K}\phi$ if, but only if, $\mathbf{K}\phi$ is decidable' (2001a: 265). The 'only if' direction, which is inessential to his argument, is an error. The validity of $\neg\neg\mathbf{K}\phi \rightarrow \mathbf{K}\phi$ is easily seen to be sufficient for the validity of $\neg(\phi \wedge \neg\mathbf{K}\phi)$ to guarantee the validity of $\phi \rightarrow \mathbf{K}\phi$ but is intuitionistically a weaker condition than the decidability of $\mathbf{K}\phi$.

12 Tennant (1997: 268) is explicit that if ϕ is decidable then $\phi \vee \neg\phi$ is intuitionistically acceptable.

13 The phrase ‘possession of a method whose application will ...’ is to be read as implying ‘recognition that application of the method will ...’.

14 As observed in n. 11, we do not need (10) to get from (7) to (8); the weaker $\neg\mathbf{K}\phi \rightarrow \mathbf{K}\phi$ would suffice. But the same considerations apply; the purported decision procedure would entitle us to assert $\neg\mathbf{K}\phi \rightarrow \mathbf{K}\phi$ only by entitling us to assert (10).

15 Any incompatibility between free will and determinism is irrelevant here. That I am causally determined not to apply a method is compatible with my possession of it in the relevant sense.

16 For more detail see Williamson 1982: 206-7 and 1988: 429-32, where the idea is used to show the invalidity of an argument for $\phi \rightarrow \mathbf{K}\phi$ from the intuitionistic semantics of \rightarrow and the premise that $\mathbf{K}\phi$ is verifiable whenever ϕ is verifiable (see Hart 1979: 165, Wright 1993: 430); the existence of a canonical verification of ϕ does not imply the existence of a canonical verification of $\mathbf{K}\phi$. Tennant describes this objection as ‘compelling’ (1997: 264).

17 Tennant says of a passage in which I emphasize the cognitive difference that applying a decision procedure makes that it ‘displays a vestige of realist thinking’ (2001a: 274); if so, Tennant has not explained how he himself can do without that vestige of realist thinking. Tennant’s further complaint concerning the conditional $\neg\mathbf{K}\phi \rightarrow \neg\phi$

and ‘a curious asymmetry (between truth and falsity)’ (2001a: 275) raises an issue that is discussed in a more acute version in Williamson 1994, to which the reader is referred.

That issue is separate from Tennant’s main argument.

18 See n. 3.

19 See n. 16.

20 For (\diamond KC) to be a well-defined rule, the notion of consequence used to define ‘Cartesian’ should be given independently of (\diamond KC) and cannot be assumed to be closed under it (see Tennant 1997: 275). This does not affect the argument in the text.

21 The examples are unconvincing. They too look *ad hoc*. Moreover, his restricted thesis about truth treats ‘This sentence is false’ as a *premise* of the Liar paradox (2001b: 110), whereas the mere well-formedness of the sentence is what makes trouble. His restricted thesis about wondering (2001b: 113) results from a derivation that makes the radically idealizing assumption (SI) that ‘a rational thinker is one whose attitudes are *self-intimating* to the thinker himself’ (1997: 248) so that I rationally wonder whether something is the case only if I believe that I so wonder (*ibid.*: 255).

22 Some such assumption is required for the relevance of his claim that ‘The restricted thesis about truth [Tarski’s] to which almost every philosopher subscribes is in fact *even more restricted* than’ a thesis about truth that Tennant wants to show not to be

ad hoc, given that ‘Tarski can hardly be accused of making an *ad hoc* restriction to his disquotational Thesis about truth’ (2001b: 111).

23 $\beta \vdash \alpha \rightarrow \beta$ in IR (Tennant 1997: 342).

24 \vdash may differ from the consequence relation used to define ‘Cartesian’. Tennant uses the weaker set of inference rules $\mathbf{K}(\varphi \wedge \psi) \vdash \mathbf{K}\varphi$ and from $\Delta, \varphi \vdash \perp$ to $\Delta, \mathbf{K}\varphi \vdash \perp$ in place of (3) and (2) respectively (1997: 259-60). Nevertheless, his discussion makes clear that on his view we can reasonably treat instances of (2) and (3) as theorems of epistemic logic.

25 The cut rule used to chain inferences together does not hold unrestrictedly in IR, but fails only in case of a redundant premise or conclusion, which (15) does not contain.

26 For the reason sketched in section I, the argument assumes a nonepistemic reading of \diamond . The mere epistemic possibility that \mathbf{n} is even does not entail that \mathbf{n} is even. $(\diamond\mathbf{K}C)$ is in any case quite unpromising on an epistemic reading of \diamond . If \diamond is read as ‘for all we know’, the principle will be unacceptable to soft anti-realists of the sort for whom Tennant seems to intend $(\diamond\mathbf{K}C)$, since, according to them, we sometimes know empirically that a given decidable proposition will never in fact be decided. If \diamond is read as ‘for all we know *a priori*’, $(\diamond\mathbf{K}C)$ is more or less trivialized because $\diamond\mathbf{K}\varphi$ then says little more than that φ is Cartesian.

27 In his examples, Tennant replaces ‘**n**’ by a numeral (2001a: 264, 271), thereby obscuring the vital point.

28 A failure to fit the philosophical arguments for anti-realism may also affect the revisions of knowability principle proposed in Edgington 1985, Melia 1991, Rabinowicz and Segerberg 1994, Kvanvig 1995, Lindström 1997 and Rückert 2003, although the point cannot be argued here; see also the comments on Dummett 2001 in n. 5. The upshot of the treatment of Fitch’s argument in Usberti 1995 is a much more drastic restriction of ϕ in (1) to mathematical sentences, which excludes those containing **K**. The proposal is grounded in Usberti’s anti-Dummettian analysis of the arguments for an intuitionistic approach (see also Williamson 1998). For a discussion of Fitch’s argument in the context of classical logic see Williamson 2000a: 270-301, 318-19.

29 Thanks to Peter Milne and Peter Pagin for comments on a draft of Williamson 2000b, which have in turn benefited the present paper.

References

- Brogaard, B., and Salerno, J. 2002. 'Clues to the Paradoxes of Knowability: Reply to Dummett and Tennant'. *Analysis* 62: 143-50.
- Cozzo, C. 1994. 'What Can we Learn from the Paradox of Knowability?'. *Topoi* 13: 71-78.
- DeVidi, D., and Solomon, G. 2001. 'Knowability and Intuitionistic Logic'. *Philosophia* 28: 319-34.
- Dummett, M. 1959. 'Truth'. *Proceedings of the Aristotelian Society* 59: 141-62.
- Dummett, M. 1975. 'The Philosophical Basis of Intuitionistic Logic', in H. Rose and J. Shepherdson (eds.), *Logic Colloquium '73*, Amsterdam: North-Holland.
- Dummett, M. 1977. *Elements of Intuitionism*. Oxford: Clarendon Press.
- Dummett, M. 2001. 'Victor's Error'. *Analysis* 61: 1-2.
- Edgington, D. 1985. 'The Paradox of Knowability'. *Mind* 94: 557-68.
- Fitch, F. 1963. 'A Logical Analysis of Some Value Concepts'. *The Journal of Symbolic Logic* 28: 135-42.
- Hand, M., and Kvanvig, J. 1999. 'Tennant on Knowability'. *Australasian Journal of Philosophy* 77: 422-28.
- Hart, W. 1979. 'The Epistemology of Abstract Objects: Access and Inference'. *Aristotelian Society* sup. 53: 152-65.
- Kvanvig, J. 1995. 'The Knowability Paradox and the Prospects for Anti-Realism'. *Noûs* 29: 481-500.
- Lindström, S. 1997. 'Situations, Truth and Knowability: A Situation-Theoretic Analysis

- of a Paradox by Fitch', in E. Ejerthed and S. Lindström (eds.), *Logic, Action and Cognition — Essays in Philosophical Logic*. Dordrecht: Kluwer.
- Melia, J. 1991. 'Anti-Realism Untouched'. *Mind* 100: 341-42.
- Pagin, P. 1994. 'Knowledge of Proofs'. *Topoi* 13: 93-100.
- Percival, P. 1990. 'Fitch and Intuitionistic Knowability'. *Analysis* 50: 182-87.
- Percival, P. 1991. 'Knowability, Actuality and the Metaphysics of Context-Dependence'. *Australasian Journal of Philosophy* 69: 82-97.
- Rabinowicz, W., and Segerberg, K. 1994. 'Actual Truth, Possible Knowledge'. *Topoi* 13: 101-15.
- Rosenkranz, S. 2004. 'Fitch Back in Action Again?'. *Analysis* 64: 67-71.
- Rückert, H. 2003. 'A Solution to Fitch's Paradox of Knowability', in S. Rahman, J. Symons, D. Gabbay and J. van Bendegem (eds.), *Logic, Epistemology, and the Unity of Science*. Dordrecht: Kluwer.
- Tennant, N. 1987. *Anti-Realism and Logic: Truth as Eternal*. Oxford: Clarendon Press.
- Tennant, N. 1997. *The Taming of the True*. Oxford: Clarendon Press.
- Tennant, N. 2001a. 'Is Every Truth Knowable? Reply to Williamson'. *Ratio* 14: 263-80.
- Tennant, N. 2001b. 'Is Every Truth Knowable? Reply to Hand and Kvanvig'. *Australasian Journal of Philosophy* 79: 107-13.
- Tennant, N. 2002. 'Victor Vanquished'. *Analysis* 62: 135-42.
- Usberti, G. 1995. *Significato e Conoscenza: Per una Critica del Neoverificazionismo*. Milan: Guerini Scientifica.
- Williamson, T. 1982. 'Intuitionism Disproved?'. *Analysis* 42: 203-7.
- Williamson, T. 1987. 'On the Paradox of Knowability'. *Mind* 96: 256-61.

- Williamson, T. 1988. 'Knowability and Constructivism'. *The Philosophical Quarterly* 38: 422-32.
- Williamson, T. 1990. 'Two Incomplete Anti-Realist Modal Epistemic Logics'. *The Journal of Symbolic Logic* 55: 297-314.
- Williamson, T. 1992. 'On Intuitionistic Modal Epistemic Logic'. *Journal of Philosophical Logic* 21: 63-89.
- Williamson, T. 1993. 'Verificationism and Non-Distributive Knowledge'. *Australasian Journal of Philosophy* 71: 78-86.
- Williamson, T. 1994. 'Never Say Never'. *Topoi* 13: 135-45.
- Williamson, T. 1998. Review of Usberti 1995. *Dialectica* 52: 63-69.
- Williamson, T. 2000a. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, T. 2000b. 'Tennant on Knowable Truth'. *Ratio* 13: 99-114.
- Wright, C. 1993. *Realism, Meaning and Truth*, 2nd ed. Oxford: Blackwell.