

IS EVERY TRUTH KNOWABLE? REPLY TO WILLIAMSON

Neil Tennant

Abstract

This paper addresses an objection raised by Timothy Williamson to the ‘restriction strategy’ that I proposed, in *The Taming of The True*, in order to deal with the Fitch paradox. Williamson provides a new version of a Fitch-style argument that purports to show that even the restricted principle of knowability suffers the same fate as the unrestricted one. I show here that the new argument is fallacious. The source of the fallacy is a misunderstanding of the condition used in stating the restricted knowability principle. I also rebut Williamson’s criticism of my argument for the claim that any proposition of the form ‘it is known that φ ’ is decidable if φ is decidable.

1. Introduction

In Chapter 8 of *The Taming of The True*¹ (henceforth: *TToTT*) I analyzed the famous argument, due to Fitch, that purports to derive absurdity by combining some general normative principles concerning the epistemic operator ‘knows that’ with two main claims, to wit

1. All truths are knowable
2. Not all truths are known (or the intuitionistically stronger, but classically equivalent: Some truths are unknown)

The problem, in a nutshell, is that Fitch showed us how to negotiate the logical passage

$$\underbrace{\psi, \neg K\psi, \forall \varphi(\varphi \rightarrow \Diamond K\varphi)}_{\vdots} \perp$$

by using very innocent-looking inferential steps. This threatens a slide from anti-realism to actualism. The anti-realist maintains

¹ Oxford: Clarendon Press, 1997.

that every truth is knowable; while the actualist maintains that every truth is known.

The main proposal in *TTtoTT* in response to what I shall call this ‘Fitch result’ was to restrict the application of the anti-realist’s knowability principle

$$(\diamond K) \forall \varphi (\varphi \rightarrow \diamond K\varphi).$$

The new, restricted version is

$$(\diamond KC) \forall \varphi ([\varphi \wedge \neg (K\varphi \vdash \perp)] \rightarrow \diamond K\varphi).$$

The restricted principle thus holds not that every truth is knowable, but only that every *Cartesian* truth is knowable. A Cartesian proposition is a proposition φ such that $K\varphi$ is consistent – that is, one cannot derive absurdity (\perp) from the claim that φ is known. We have expressed the Cartesian character of φ by the formal condition

$$\neg (K\varphi \vdash \perp)$$

in the foregoing. 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 premisses can be suppressed. Thus if I can derive absurdity from the assumptions $K\varphi$ and ‘4 is even’, taken together, then I shall have succeeded in showing that φ is not Cartesian. In general, if ψ_1, \dots, ψ_n are all theorems, and $K\varphi, \psi_1, \dots, \psi_n \vdash \perp$, then φ is not Cartesian.

Two critical responses to this proposal raise the objection that the restriction of the knowability principle to Cartesian propositions is unprincipled and *ad hoc*. One such response is from Michael Hand and Jonathan Kvanvig.² The other is from Timothy Williamson.³

I have rebutted Hand and Kvanvig’s charge of *ad hoc*-ness elsewhere;⁴ and have nothing to add by way of response to Williamson’s version of the same charge. Here I intend, instead, to focus on other aspects of Williamson’s critique.

2. Anti-realism and anti-actualism

In *TTtoTT* a distinction had been drawn between soft and hard anti-realism. The soft anti-realist refuses to accept the purported

² ‘Tennant on Knowability’, *Australasian Journal of Philosophy*, vol. 77, no. 4 (December 1999), pp. 422–28.

³ ‘Tennant on Knowable Truth’, *Ratio*, XIII No. 2 (June 2000).

⁴ ‘Is every truth knowable? Reply to Hand and Kvanvig’, *Australasian Journal of Philosophy*, forthcoming.

Fitch result that every sentence φ is inconsistent with $\neg K\varphi$. What the *reductio* schema above shows is that if the soft anti-realist asserts (*KP*) without my proposed restriction, then he must assert both $\neg\exists\varphi(\varphi\wedge\neg K\varphi)$ and $\forall\varphi\neg(\varphi\wedge\neg K\varphi)$. But if the soft anti-realist accepts my proposed restriction on (*KP*), then the *reductio* above cannot be carried out. Consequently the soft anti-realist will be in a position to regard the schema $\neg(\varphi\wedge\neg K\varphi)$ as invalid. (If one wishes to use quantified forms, maintaining the validity of a schema ($\neg\varphi$) amounts to asserting its universal quantification $\forall\varphi(\neg\varphi)$, and maintaining the invalidity of the schema amount to denying its universal quantification, i.e. asserting $\neg\forall\varphi(\neg\varphi)$.)

The hard anti-realist, unlike the soft anti-realist, does accept the Fitch result, and seeks to accommodate it. Williamson divides the category of hard anti-realism into two sub-categories, which he calls *moderately hard* and *very hard* anti-realism. The defining claims of the three positions are as follows:

Brand of anti-realism:	$\neg(\varphi\wedge\neg K\varphi)$	$\varphi\rightarrow K\varphi$
Soft	Invalid	Invalid
Moderately hard	Valid	Invalid
Very Hard	Valid	Valid

Williamson visits upon me the view that moderately hard anti-realism collapses into very hard anti-realism. For I had pointed out that the validity of $\neg(\varphi\wedge\neg K\varphi)$ guarantees the validity of $\varphi\rightarrow K\varphi$ if, but only if, $K\varphi$ is decidable. And I offered an argument to establish that $K\varphi$ is decidable if and only if φ is decidable. The direction that matters is

If φ is decidable then $K\varphi$ is decidable.

That would then license, for decidable φ , the application of classical *reductio ad absurdum* to $K\varphi$. Thus one would be able to deduce the actualist's schema $\varphi\rightarrow K\varphi$ from $\neg(\varphi\wedge\neg K\varphi)$, by means of the proof

$$\begin{array}{l}
 (2)\text{---} \quad \text{---}(1) \\
 \frac{\varphi \quad \neg K\varphi}{\varphi\wedge\neg K\varphi} \\
 \frac{\varphi\wedge\neg K\varphi \quad \neg(\varphi\wedge\neg K\varphi)}{\perp} \\
 (CR)\text{---} \frac{\perp}{K\varphi} (1) \\
 \frac{\quad}{\varphi\rightarrow K\varphi} (2)
 \end{array}$$

Nothing in these considerations, however, speaks to the case where φ is undecidable. In that case, the distinction between moderately and very hard anti-realism can still be sustained.⁵ So, *pace* Williamson, I do not believe that moderately hard anti-realism is an unstable position. It is just that, unlike Williamson, I do not find flagrantly counterintuitive the thought that every *decidable* true proposition is known. The theorist such as Williamson, however, who wishes to distinguish between moderately and very hard anti-realism, *and* wishes to be strenuously anti-actualist by denying (even for only decidable propositions) that all truths are known, will be forced (as Williamson indeed realized) to seek some way to prevent the validity of $\varphi \rightarrow K\varphi$ from following upon the validity of $\neg(\varphi \wedge \neg K\varphi)$, even when only decidable propositions are under consideration. Williamson raises an objection to my proof that $K\varphi$ is decidable if φ is; but I do not find that objection compelling. I shall return to a defence of my proof against his objection in due course.

In judging the stability of various positions, one does better by considering not just brands of anti-realism, but combinations of those with various brands of *anti-actualism*. The actualist theses that I shall use in order to effect a classification of brands of anti-actualism are ‘All true propositions are known’ and ‘All true decidable propositions are known’. The actualist holds the former and hence also the latter. I shall suggest the following labels for the *anti-actualist* positions characterized:

Brand of anti-actualism:	$\forall \varphi(\varphi \rightarrow K\varphi)$	
	$\forall \varphi([\varphi \wedge (\varphi \text{ is decidable})] \rightarrow K\varphi)$	
Soft	Invalid	Valid
Hard	Invalid	Invalid

Williamson seeks to establish the coherence of a moderately hard anti-realism with hard anti-actualism – even if (that brand of) the anti-realism might not be a view that he would ultimately find acceptable. The position in question is characterized by the following claims:

Valid:	$\neg(\varphi \wedge \neg K\varphi)$
Invalid:	$\varphi \rightarrow K\varphi$
Invalid:	$[\varphi \wedge (\varphi \text{ is decidable})] \rightarrow K\varphi$

⁵ Note that by ‘undecidable’ here we mean ‘not within the range of an effective decision procedure’. We do not mean that the truth-value of the claim concerned transcends all our methods for discovering truth-values.

To maintain the stability of this combination of views, Williamson must find a way to resist the passage from ‘ φ is decidable’ to ‘ $K\varphi$ is decidable’. Otherwise, by means of the previous proof, the validity of $\neg(\varphi \wedge \neg K\varphi)$ will secure the validity of $\varphi \rightarrow K\varphi$ for decidable φ , contrary to the position just characterized. Williamson’s criticism of the way in which I negotiated this passage I shall for convenience call the ‘decidability critique’. Its rebuttal comes later.

For my own part, although I did not argue explicitly for as much in *TToTT*, I find attractive the combination of moderately hard anti-realism with soft anti-actualism – a combination deriving from the conviction that ‘ φ is decidable’ entails ‘ $K\varphi$ is decidable’. Pursuit of the restriction strategy in *TToTT*, however, was presented as worthwhile precisely because of the unwillingness that I anticipated, among anti-realists, to be merely soft anti-actualists. Hard anti-actualism is very beckoning, since it involves no reform of pre-philosophical intuitions about the nature of knowledge. So I was concerned to rescue some still powerful and contentious form of the knowability principle that would satisfy the hard anti-actualists among fellow anti-realists.

Williamson now charges that my attempt so to restrict the knowability principle founders on a modified Fitch-style result, within an extended logical system that includes the operators \diamond and K . To this charge we now turn.

3. On Williamson’s critique of the restricted knowability principle

3.1. Logical exegesis of Williamson’s reasoning

Williamson sets out the following argument:

... let φ be the decidable sentence ‘There is a fragment of Roman pottery at that spot’ (we assume a suitable context). Introduce a name ‘ n ’ by the stipulation that it is to designate (rigidly) the number of books actually now on my table. Let E be the predicate ‘is even’. We first argue that the conjunction $\varphi \wedge (K\varphi \rightarrow En)$ is Cartesian.

Let \vdash be the consequence relation of a system of epistemic logic based on *IR* with the additional rule $(\diamond KC)$. . . [I omit the argument that Williamson gives at this point for the Cartesian status of $\varphi \wedge (K\varphi \rightarrow En)$; we shall return to it below. – *NT*] . . . Since its condition is met in this case, $(\diamond KC)$ gives:

$$(12) \quad \varphi \wedge (K\varphi \rightarrow En) \vdash \diamond K \varphi \wedge (K\varphi \rightarrow En))$$

Moreover:

$$(13) \quad \varphi \wedge \neg K\varphi \vdash \varphi \wedge (K\varphi \rightarrow En)$$

(12) and (13) yield

$$(14) \quad \varphi \wedge \neg K\varphi \vdash \diamond K(\varphi \wedge (K\varphi \rightarrow En))$$

The next step is an argument following Fitch for:

$$(15) \quad K(\varphi \wedge (K\varphi \rightarrow En)) \vdash En$$

. . . Since the rules used to derive [this last result] are truth-preserving in all possible situations, not just the actual one, if the premise of [the last result] expresses a possibility, so does its conclusion . . . :

$$(16) \quad \diamond K(\varphi \wedge (K\varphi \rightarrow En)) \vdash \diamond En$$

(14) and (16) yield:

$$(17) \quad \varphi \wedge \neg K\varphi \vdash \diamond En$$

We can strengthen (17) to:

$$(18) \quad \varphi \wedge \neg K\varphi \vdash En$$

For it is not contingent whether n is even We can now repeat the argument for (18) with 'odd' in place of 'even' to derive:

$$(19) \quad \varphi \wedge \neg K\varphi \vdash \neg En$$

But (18) and (19) yield:

$$(20) \quad \vdash \neg(\varphi \wedge \neg K\varphi)$$

. . . But the point of Tennant's restricted knowability principle ($\diamond KC$) was precisely to enable the soft anti-realist not to assert $[\neg(\varphi \wedge \neg K\varphi)]$. . . Thus Tennant's restriction is futile.

The following regimentation of Williamson's argument employs the placeholder θ where Williamson chooses to have either En or $\neg En$. Thus the reader should bear in mind that the following proof-schema with conclusion θ has to be instantiated *twice* in order to produce Williamson's intended reductio of the set $\{\varphi, \neg K\varphi, (\diamond KC)\}$. (The first instantiation will use En for θ , the second will use $\neg En$; finally, one applies $\neg E$ to obtain \perp .) We also indicate by Π_θ whatever will count as the proof, for the particular

proposition θ concerned, of the claim that the proposition $\varphi \wedge (K\varphi \rightarrow \theta)$ is Cartesian:

$$\begin{array}{c}
 \frac{\frac{\neg K\varphi \quad \text{---} (1)}{K\varphi}}{\perp} (1) \qquad (2) \frac{\text{---}}{K[\varphi \wedge (K\varphi \rightarrow \theta)]} \\
 \\
 \frac{\frac{\varphi \quad K\varphi \rightarrow \theta}{\varphi \wedge (K\varphi \rightarrow \theta)} \quad \frac{\Pi_\theta \quad K[\varphi \wedge (K\varphi \rightarrow \theta)] \quad \varphi \wedge (K\varphi \rightarrow \theta)}{\neg(K[\varphi \wedge (K\varphi \rightarrow \theta)] \vdash \perp)}}{(\diamond KC) \frac{\varphi \wedge (K\varphi \rightarrow \theta) \quad \neg(K[\varphi \wedge (K\varphi \rightarrow \theta)] \vdash \perp)}{\diamond K[\varphi \wedge (K\varphi \rightarrow \theta)]}} \quad \frac{K\varphi \quad K\varphi \rightarrow \theta}{\theta} (2)}{\frac{\diamond K[\varphi \wedge (K\varphi \rightarrow \theta)] \quad \theta}{\diamond \theta}} \\
 \frac{\diamond \theta}{\theta}
 \end{array}$$

Let us call this proof-schema $\Xi[\Pi_\theta, \theta, \varphi]$. The subproof of θ on the right regiments the ‘argument following Fitch’ for Williamson’s (15) in the quote above.

The final step of the proof-schema just given is to be justified, according to Williamson, by appeal to the special character of the proposition θ : ‘it is not contingent whether $[\theta]$ ’. To repeat: the proof-schema will be instantiated once with ‘ n is even’ in place of θ , and a second time with ‘ n is not even’ in place of θ .

3.2. Locating the fallacy in Williamson’s reasoning

Let us call any proposition θ which, if possibly true, is necessarily true, and, if possibly false, is necessarily false, a polar proposition. (We are dealing primarily with *logico-mathematical* possibility and necessity here.) Examples are ‘ n is even’, or ‘ n is not even’, where ‘ n ’ rigidly designates a particular natural number. And of course there are many other examples besides these. Any provable or refutable mathematical proposition is polar. So too – if we countenance *metaphysical* necessity and impossibility – are propositions such as ‘Water is H_2O ’ and ‘Water is XYZ ’. Not all polar propositions are effectively decidable. Moreover, if we countenance metaphysical necessity and impossibility, then not all polar propositions’ truth-values will be able to be determined *a priori*.

Which propositions are Cartesian in the actual world will depend, in general, on which polar propositions are true. Recall that φ is Cartesian just in case $\neg(K\varphi \vdash \perp)$. To say that absurdity is not derivable from $K\varphi$ is equivalent to saying that absurdity is not derivable from $K\varphi$ in conjunction with any set X of necessarily

true propositions. Whether this definition calls for the consideration only of sets X all of whose members are knowable *a priori*, or calls for the consideration also of sets 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.⁶

Williamson's argument is ingenious – for invoking such polar propositions – but flawed – for the same reason. We have to take a closer look at how the subproof Π_θ is put together. Here is Williamson's reasoning – to be regimented by Π_θ – in the case where θ is the (polar) proposition ' n is even' (this is the reasoning that we omitted from the earlier quote):

... the conjunction $\varphi \wedge (K\varphi \rightarrow En)$ is Cartesian. For suppose that $K(\varphi \wedge (K\varphi \rightarrow En))$ is inconsistent. Then this story contains an inconsistency:

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. [fn] By putting the two pieces of knowledge together, I acquire the knowledge expressed by the conjunction $\varphi \wedge (K\varphi \rightarrow En)$. Thus $K(\varphi \wedge (K\varphi \rightarrow En))$ is true.

But the story is obviously consistent. Thus \perp does not follow from $K(\varphi \wedge (K\varphi \rightarrow En))$, so $\varphi \wedge (K\varphi \rightarrow En)$ is Cartesian.

Let us abbreviate by φ the decidable proposition 'There is a fragment of Roman pottery at that spot', and by θ the decidable polar proposition 'the number that is actually the number of books now on my table is even'. In Williamson's story, it is being assumed (correctly, but implicitly) that $K\varphi$ and $K\theta$ can be true together. Any logical consequence of their combination, therefore, will be consistent, assuming the soundness of the underlying proof system. By means of the following proof, the proposition $K(\varphi \wedge (K\varphi \rightarrow \theta))$ is shown to be just such a consequence:

⁶ Nowhere in *TTtoTT* did I claim that the Cartesian character of a proposition would always be an *a priori* matter. But as it happens, Williamson invokes only *a priori* polar propositions, such as ' n is even'. (Bear in mind that this is a *mathematical* proposition, since n is a rigid designator.)

$$\frac{\frac{\frac{\frac{\quad}{\theta} (1)}{\varphi} (1)}{K\varphi \rightarrow \theta}}{K\varphi \quad K\theta \quad \varphi \wedge (K\varphi \rightarrow \theta)} (1)}{K(\varphi \wedge (K\varphi \rightarrow \theta))} (1)$$

Hence the proposition $\varphi \wedge (K\varphi \rightarrow \theta)$ is Cartesian. *But* it is hereby shown to be Cartesian only conditionally on the assumption that it is possible for both $K\varphi$ and $K\theta$ to hold.

Suppose that θ had instead been the (necessarily) false proposition ‘3 is even’. Let φ be as before. Then the proposition $\varphi \wedge (K\varphi \rightarrow \theta)$ would *not* be Cartesian. For, in this case, we would have $K(\varphi \wedge (K\varphi \rightarrow \theta)) \vdash \perp$, by the following proof:

$$\begin{array}{c} (2) \frac{\frac{\frac{\quad}{K[\varphi \wedge (K\varphi \rightarrow \theta)]} (2)}{\varphi \wedge (K\varphi \rightarrow \theta)} (2)}{K\varphi \quad K\varphi \rightarrow \theta} \\ \hline \theta \\ \Sigma \\ \perp \end{array}$$

where Σ is the *reductio* of the (necessarily false) mathematical claim θ (that 3 is even).

The logical moral is that the Cartesian character of $\varphi \wedge (K\varphi \rightarrow \theta)$ depends on the truth-value of the proposition θ when θ is polar. For polar θ , it is only when θ is true that $\varphi \wedge (K\varphi \rightarrow \theta)$ is Cartesian. Thus no matter how we fill in the details for the subproof Π_θ within the proof-schema $\Xi[\Pi_\theta, \theta, \varphi]$, it will have to involve at least the hidden assumption θ , even if not the stronger assumption $K\theta$. All that $\Xi[\Pi_\theta, \theta, \varphi]$ succeeds in establishing, therefore, is that θ follows logically from the set of assumptions $\{\varphi, \neg K\varphi (\diamond KC), \theta\}$ – not a very informative result, since the conclusion is one of the premisses.

Let us now reassess Williamson’s original argument. We can say that if the number n that happens to be the number of books on his table is even, then he can construct the uninformative, because circular, proof

$$\Xi[\Pi_{En}, En, \varphi]$$

of the conclusion En from the set of assumptions

$$\{\varphi, \neg K\varphi, (\diamond KC), En\},$$

but will be unable to construct

$$\Xi[\Pi_{\neg En}, \neg En, \varphi],$$

since when En holds the proposition

$$\varphi \wedge (K\varphi \rightarrow \neg En)$$

is not Cartesian. On the other hand, if the number n that happens to be the number of books on Williamson's table is odd, then he can construct the uninformative, because circular, proof

$$\Xi[\Pi_{\neg En}, \neg En, \varphi]$$

of the conclusion $\neg En$ from the set of assumptions

$$\{\varphi, \neg K\varphi, (\diamond KC), \neg En\},$$

but will be unable to construct

$$\Xi[\Pi_{En}, En, \varphi],$$

since when $\neg En$ holds the proposition

$$\varphi \wedge (K\varphi \rightarrow En)$$

is not Cartesian. However many books there happen to be on Williamson's table, he will not be able to produce both a proof of En from just the assumptions $\varphi, \neg K\varphi$ and $(\diamond KC)$, and a proof of $\neg En$ from those same assumptions. So he will be unable to construct his intended *reductio ad absurdum* of the restricted knowability principle proposed in *TTtoTT*.

4. On Williamson's decidability critique

Williamson concedes that my objection to moderately hard anti-realism (as it concerns *decidable* propositions) is sustained if I can demonstrate that $K\varphi$ is decidable whenever φ is. He quotes the argument I gave in *TTtoTT*, at p. 262:

Suppose that φ is decidable. Then here is a decision method for $K\varphi$: 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 $K\varphi$ is true. If, on the other hand, you determine that φ is false, then you have determined that $K\varphi$ is false,

because no one could ever know a falsehood. So if φ is decidable, then so is $K\varphi$.

Williamson correctly points out that in the case where φ is false, we would (by applying the decision procedure) determine not only that $K\varphi$ is false, but also that $K\neg\varphi$ is true. Thus one has a decision procedure that will justify not only the disjunction $K\varphi \vee \neg K\varphi$ but also the disjunction $K\varphi \vee K\neg\varphi$.

Nothing in the quoted passage, however, implies an interpretation on which the knowing in question is attained only upon completing the application of the decision procedure. In a perfectly admissible sense of 'knows that', one already knows that φ (if φ is true) or knows that $\neg\varphi$ (if φ is false) simply by being minded of a decision procedure for φ .

Suppose a teacher in a mathematics class asks the students 'Who knows the answer to the following question: At what value of x does the function $3x^2+12x+48$ attain its minimum?' An intelligent child can legitimately and immediately answer 'I do' simply by recognizing the problem as belonging to a general class for which he has a decision procedure. He need not wait to *apply* the decision procedure in question before being entitled to say 'I do' – that is (in this context), 'I know at what value of x that function attains its minimum.' The child knows that all he has to do is find the derivative of the quadratic, set it equal to 0, and solve for x . Here, the derivative is $6x+12$, so the answer is -2 . Now of course the child might be able to work this answer out so quickly after saying 'I do' that he is ready to supply it as soon as the teacher says 'Well, Johnny, what is the answer then?' But suppose he is not so quick. Suppose, instead, that Johnny does not have a ready answer to the teacher's follow-up question, but talks aloud as he slowly determines the value of x according to the decision procedure. Does that mean that his first claim 'I do' (i.e., 'I know at what value of x that function attains its minimum') is *false*? I maintain that it need not. Knowledge of the correct answer is attributable on the basis of knowledge of the relevant decision procedure, whether or not the procedure is actually applied. So when one reads $K\varphi$ as $\exists t \exists x (xK_t\varphi)$, the temporal quantification is satisfied by the time at which one is minded of a correct decision procedure for φ . One does not have to wait for the time at which the application of the procedure terminates, and one first becomes occurrently aware of the result.

All this is part of a standard intuitionistic understanding of

direct and indirect proofs, and a standard understanding, on the part of epistemic logicians, that the notion of knowledge being regimented in an epistemic logic is that of so-called *virtual* knowledge. In epistemic logic, the knower is logically committed to knowing the logical consequences of what he knows. In the absence of such an idealizing assumption, there is, arguably, no such thing as an epistemic logic. Indeed, if this idealizing assumption were disallowed, the original Fitch argument would not go through.

Most knowledge-claims in mathematics are backed only by indirect proofs, which encode the effective procedures that are required in order (among other things) to produce, on demand, the various objects satisfying the existential claims involved. Whenever an intuitionist proves a result ψ by using a constructive dilemma on a decidable proposition ϕ :

$$\frac{\frac{\text{---} \quad (i) \quad \text{---} \quad (i)}{\begin{array}{c} \phi \\ \vdots \\ \psi \end{array}} \quad \frac{\text{---} \quad (i)}{\begin{array}{c} \neg\phi \\ \vdots \\ \psi \end{array}} \quad (i)}{\psi} \quad (i)$$

he is thereby entitled to claim that he knows that ψ , even though he has not actually applied the decision procedure that would decide ϕ . Like the schoolboy in the mathematics class, our intuitionist knows that *if called upon to produce a canonical proof* of ψ , he will be able to do so, by (among other things) applying the decision procedure for ϕ in order to determine which horn of the constructive dilemma to extract for inclusion in the canonical proof that he will have to produce. Actually obtaining a canonical proof is analogous to actually obtaining the value $x = -2$ in the schoolboy example. One does not need the terminus in order to be able to make the knowledge-claim. One needs only a method that is guaranteed to get one there.

I therefore disagree with Williamson when he writes

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 have no reason to think that it will ever be applied . . .

This displays a vestige of realist thinking, according to which knowing is a matter of being in a certain occurrent representational state, rather than possessing an effective means for attaining it.

Moreover, as Andrew Arlig has observed, the intuition being invoked by Williamson here does not sit well with the overall position whose coherence Williamson seeks to defend. For Williamson's moderately hard anti-realist accepts, as a result of the Fitch argument, the conclusion

$$\neg K\phi \rightarrow \neg\phi.$$

In the spirit of the last quote above from Williamson, we can ask him to imagine a case of some *true* and decidable proposition ϕ for which indeed the decision procedure is never applied. Then, according to the intuition on which Williamson sets so much store, we would have $\neg K\phi$. Thus, courtesy of Fitch, it would follow that $\neg\phi$ – that is, that ϕ was *false*. It would therefore appear that Williamson can countenance only *false* decidable propositions as those whose decision procedures, fortuitously, are never applied. Such a curious asymmetry (between truth and falsity) surely calls for explanation.

An opponent will probably object that our preferred 'potentiality' sense of 'knows that' – what we called virtual knowledge above – is for its own part highly counterintuitive in common examples, and that it should not be allowed to displace the 'occurrent' sense of 'knows that' in our theorizing about knowledge. For example, the proposition

There is a silver dollar in this box

is decidable. The decision procedure is to take off the lid and have a look. But – the opponent continues – suppose no one ever does that. Are we really to say, simply because we have this decision procedure, that

Either someone at some time knows that there is a silver dollar in this box or someone at some time knows that there is no silver dollar in this box?

Isn't this a clear case where complete ignorance (as to the contents of the box) might prevail? Might not the decision procedure remain forever unapplied?

The opponent's intuition appears very strong: for many a decidable proposition ϕ , it is possible that no one ever (occurrently) knows whether ϕ . How are we to reform this intuition, and in a theoretically sensible way, when seeking to maintain that if the proposition ϕ is decidable then so is the proposition $K\phi$ (i.e. someone at some time knows that ϕ)?

Williamson sees the knowledge in question (as to the truth or falsity of φ) as a state that is *brought about* by the application of the decision procedure for φ . He writes

What has gone wrong is that the application of the decision procedure for φ *brings about* the state of affairs expressed by $[K\varphi \vee K\neg\varphi]$.

By contrast, I want to regard the knowledge whether φ as implicit in the possession of the decision procedure. This appears to be the most sensible way, theoretically, to reform the intuition that lays undue stress on the occurrent sense of 'knows that'. Suppose Jones is minded of a decision procedure Φ for the proposition φ . Then Jones can correctly describe the correct answer to the question 'Is it the case that φ ?'. He can say 'the answer is the unique member of the pair $\{\varphi, \neg\varphi\}$ that is determined by applying the procedure Φ .' All that the actual application of the decision procedure Φ would do is supply another form of words for the expression of the proposition thus identified.

Notice that there is a strong disanalogy between the case of applying a known decision procedure for φ (hence for $K\varphi$) and Williamson's case – allegedly, a 'more dramatic example of the fallacy' he thinks is involved in the potentialist's way of viewing matters – of the person who claims

We do have a decision procedure for the sentence ['A city was, is or will be built on this spot.'] 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.

The difference here is that a person who applies this so-called 'decision procedure' can say in advance which member of the pair of propositions

A city was, is or will be built on this spot.

It is not the case that a city was, is or will be built on this spot.

will thereby be determined as the correct answer, in canonical form, to the question 'Is it the case that a city was, is or will be built on this spot?' That is not the case with decision procedures properly so-called. 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*. Proper decision

procedures do not interfere with the states of affairs in whose propositional representations we are interested.

Applying a proper decision procedure for φ , however, *unavoidably* determines an answer not only to the question whether φ holds, but also to the epistemic question whether someone, at some time, knows that φ . No decision procedure for φ can be limited only to the question whether φ holds, and be non-decisive with respect to the epistemic question. Moreover, its decisiveness in this latter regard is not at all like that of the would-be city-builder seeking to 'determine the answer' to the question whether a city was, is or will be built on this spot. For, in advance of applying the decision procedure for φ , one cannot say what answer in canonical form will thereby (unavoidably) be determined for the epistemic question whether it is known that φ . This is a serious enough point of disanalogy to render Williamson's diagnosis of an alleged fallacy (in my argument from ' φ is decidable' to ' $K\varphi$ is decidable') unacceptable to the epistemic logician.

5. Corrigendum

One aspect of Williamson's critique, though charitably confined to a footnote,⁷ is important enough to highlight and concede. For the purpose of describing my earlier mistake and correcting it, assume that one is equipped with the unrestricted knowability principle ($\diamond K$). Assume also an epistemic (intuitionistic) logic that includes the principles (basic or derived)

$$K(\varphi \wedge \psi) \rightarrow K\varphi$$

$$K(\varphi \wedge \psi) \rightarrow K\psi$$

along with such other principles as are needed, and would suffice, for proof of the Fitch result. Fitch, Hand and Kvanvig, and Williamson all help themselves for this purpose to the factive principle

$$K\varphi \rightarrow \varphi$$

One of the minor merits that can be claimed, in passing, for the exposition of Fitch's result in *TT₀TT* was that a strictly weaker principle than this was shown to suffice. The principle in question

⁷ *Loc. cit.*, footnote 4.

(which I called (I)) is that if ϕ is inconsistent (*modulo* certain other assumptions), then so is the claim $K\phi$ (where it is assumed that we are talking about the knowledge of an ideally rational agent).

Remember that the turnstile here is that of intuitionistic logic. In *TTtoTTI* I claimed, incorrectly, that even if

$$(1) (\diamond K), \psi, \neg K\psi \vdash \perp,$$

it need not follow that

$$(2) (\diamond K), \neg(\psi \rightarrow K\psi) \vdash \perp.$$

This was a careless claim. It was prompted by the correct, but inconclusive, observation that from $\neg(\psi \rightarrow \theta)$ the intuitionist can conclude only $\neg\neg\psi$, but cannot conclude ψ . This observation puts the kybosh on the following classical vindication of the move from (1) to (2):

$$\begin{array}{c}
 \neg(\psi \rightarrow K\psi) \\
 \vdots \quad \neg(\psi \rightarrow K\psi) \\
 \neg\neg\psi \quad \vdots \\
 \psi \quad \neg K\psi \\
 (\diamond K) \\
 \underbrace{\hspace{10em}} \\
 \vdots \\
 \perp
 \end{array}$$

But of course it is much too swift to conclude from the strictly classical character of this proof-schema (with its application of Double Negation Elimination to infer ψ from $\neg\neg\psi$) that the passage from (1) to (2) is denied to the intuitionist. For one need only observe that (2) is a statement to the effect that a certain set of sentences is inconsistent. And by the well-known Glivenko theorem, any classically inconsistent set of sentences (in propositional logic) is intuitionistically inconsistent.⁸ It would remain to check only that the presence of the modal-epistemic operator $\diamond K$ and the epistemic operator K presented no obstacles to the usual way of securing the inconsistency intuitionistically. And indeed they do not:

⁸ For a proof of this result, I blush to confess, the reader may consult N. Tennant, *Natural Logic*, (Edinburgh: University Press, 1978), p. 94.

$$\begin{array}{c}
 \frac{}{K\psi} \quad (1) \\
 \hline
 \psi \rightarrow K\psi \neg(\psi \rightarrow K\psi) \\
 \hline
 \perp \\
 (2) \frac{}{\psi, \neg K\psi, (\diamond K)} \quad (1) \\
 \hline
 \vdots \\
 \frac{}{K\psi} \\
 \hline
 \psi \rightarrow K\psi \neg(\psi \rightarrow K\psi) \quad (2) \\
 \hline
 \perp
 \end{array}$$

The foregoing is a slightly more direct intuitionistic proof of inconsistency than the one provided by Williamson.

The point just conceded in the previous section should not, however, be allowed to detract from the main dialectical thrust in that part of the discussion in *TTtoTT*. As it happens, the error committed there arose while making a tangential point. The main point at issue, which I still maintain, is that the Fitch result allows one to infer the ‘actualist’ claim $\forall \phi(\phi \rightarrow K\phi)$ only for decidable propositions ϕ .

6. Conclusion

Regardless of how semantic and epistemic intuitions eventually incline an anti-realist to one reflective equilibrium rather than another, it should be clear that it is not at all straightforward how one is to avoid the collapse of moderately hard anti-realism into hard anti-realism that I claim takes place with regard to decidable propositions. There is a plausibly reformed sense of ‘knows that’ (the potentialist one described above) that cannot be so easily dismissed as part of a coherent anti-realist approach in the event that such a collapse be granted.

But even if Williamson can convince others that such collapse should not be granted, he has still failed to clinch his case for moderately hard anti-realism as the sole sensible anti-realist response to the Fitch paradox. For there is still the strategy, which I have made available to the soft anti-realist, of restricting the

principle of knowability to Cartesian propositions. Williamson has failed to show that the restricted knowability principle suffers from the same problems as the unrestricted one.

*Department of Philosophy
The Ohio State University
Columbus, Ohio 43210
USA
tennant.9@osu.edu*