Beyond the Limits of Knowledge

Graham Priest

May 22, 2006

1 Introduction

Are there limits to knowledge? Well, there are certainly many things that we do not, as a matter of fact, know. We do not know (at the moment) whether Iraq will continue its downward spiral into anarchy. We will know in due course. We do not know how to make the Theory of Relativity and Quantum mechanics consistent with each other. Maybe we will in due course. More interesting is the question of whether there are things that it is not *possible* to know. Perhaps there are things that are so difficult, remote, or recondite, that they transcend anything we could find out. If this is the case, there are even limits to what it is possible to know. Whether or not this is so is the main topic of this paper.

'Possible' is a highly ambiguous word in philosophy. It can mean 'logically possible', 'physically possible', 'epistemically possible', and doubtless many other things. It may therefore reasonably be asked what sense of possibility is at issue here. The answer is that it doesn't really matter. For most of the purposes of this paper, it can mean any sense of possibility one likes.

One group of people who assert that all truths are knowable (in some appropriate sense) comprises verificationists, including mathematical intuitionists. For them, this is a constraint on truth itself (or maybe on meaning): everything that is true is such that it is possible (at least in principle) to know it. What sense of 'possible' verificationists have in mind here, I leave them to explain. But at least in their honour, I call the principle that all truths are possibly known the *Verification Principle.* This Principle settles the matter at issue in one way.

On the hand, there is a well known argument, usually attributed to Fitch, to the effect that Verification Principle is false. If this is the case, then there are some truths that it is impossible to know. This resolves the issue in the opposite direction. We may therefore approach the matter by considering the tenability of the Verification Principle in the light of the Fitch argument.

2 Setting up the Issue

Let us start by getting the geography of the issue straight. Some notation: I will use lower case Greek letters for sentences of whatever language is at issue. $\Diamond$ and $\Box$ are the usual modal operators of possibility and necessity. Kx will be the predicate 'is known', and $\langle . \rangle$ is a name-forming device. Now, let $\mathcal{T}$ be the set of truths. The question is how what we know relates to this. There are two relevant subsets. The first comprises the truths that are (actually) known, $\mathcal{K} = \{x : Kx\}$. The second comprises the truths that it is possible to know, $\mathcal{P} = \{x : x \in \mathcal{T} \wedge \Diamond Kx\}$. (Note that $K \langle \alpha \rangle$ entails that α is true; but $\Diamond K \langle \alpha \rangle$ does not—only that α is possibly true.) Since what is known is possibly known (in any normal sense of possibility), the general relationship between the three sets is as follows:

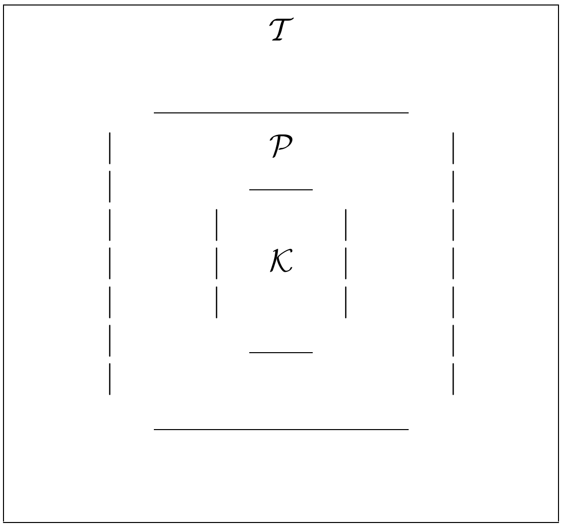

$\mathcal{K}$ is certainly non-empty. Melbourne, for example, is known to be in Australia. $\mathcal{P} - \mathcal{K}$ is also non-empty. As I have already observed, there are things about the future that we do not know, but will; so *that* knowledge is certainly possible. Similarly, the Ancient Greeks did not know that there was a planet beyond Uranus; but it is possible to know this: we do. The status of $\mathcal{T} - \mathcal{P}$ is less clear. The Verification Principle says that $\alpha \rightarrow \Diamond K \langle \alpha \rangle$. If this is true, $\mathcal{T} - \mathcal{P}$ is empty; if there is a counter-example to the Principle then there are truths that it is not possible to know.

Now to the Fitch argument.[1] This is to the effect that if it is possible to know whatever is true, then everything true is not just *possibly* known, but *actually* known. That would appear to be a *reductio ad absurdum* of the view. It is clear that not everything is actually known—even if one is a verificationist. *A priori*, there is something highly suspect about the argument, however. Surely one cannot get from the mere fact that it is *possible* to know something to the fact that it is known?

Informally, Fitch's argument goes as follows. Suppose that everything true is knowable, and suppose for *reductio* that there is something, α, which is true but not known, $\alpha \wedge \neg K \langle\alpha\rangle$. Then it must be possible to know this, $\Diamond K \langle\alpha \wedge \neg K \langle\alpha\rangle\rangle$. By a few straightforward inferences concerning knowledge, it follows that it is possible to both know α and not know it, $\Diamond(K \langle\alpha\rangle \wedge \neg K \langle\alpha\rangle)$, which it isn't.

3 The Fitch Argument

3.1 Stage 1: Knowledge

Let me spell out the argument in detail (in natural deduction form), so that we may look at the moves in it more carefully. For the purpose of discussing the argument, and in the cause of simplicity, I will write $K \langle\alpha\rangle$ as $K\alpha$, effectively turning the predicate K into the more usual operator. (As long as we are not quantifying-in, there is no real difference.)

The part of the proof concerning knowledge goes as follows. Call it Π_1.

$$\dfrac{\dfrac{K(\alpha \wedge \neg K\alpha) \qquad \dfrac{[\alpha \wedge \neg K\alpha]}{\alpha}}{K\alpha} \qquad \dfrac{\dfrac{K(\alpha \wedge \neg K\alpha)}{\alpha \wedge \neg K\alpha}}{\neg K\alpha}}{K\alpha \wedge \neg K\alpha}$$

Π_1 uses four inferences:

$$\dfrac{\beta \quad \gamma}{\beta \wedge \gamma} \qquad \dfrac{\beta \wedge \gamma}{\gamma} \qquad \dfrac{K\beta}{\beta} \qquad \dfrac{K\beta \quad \begin{array}{c}[\beta]\\ \vdots \\ \gamma\end{array}}{K\gamma}$$

In the fourth of these, the column from β to γ represents an argument with premise β (and only β), and conclusion γ . The square brackets represent

[1] Fitch (1963). Fitch himself attributes the argument to an anonymous source. It lay dormant for some time, but was published again by Hart and McGinn (1976), whose attention was drawn to it, again, anonymously.

the fact that the inference discharges β, so that the final argument no longer depends on it.

The first two inferences require no comment; nor does the third: what is known is true. The fourth says that knowledge is closed under entailment. This is certainly not correct. What is actually known is not closed under entailment. For example, medieval monks knew that Aristotle was Greek. They did not know that (Aristotle was Greek or the formalism of quantum mechanics deploys Hilbert spaces), even though this entails it. Or consider the Peano postulates. I know all these. But I do not know all their consequences (amongst which are probably the solutions to some famous unsolved problems in number theory).

But the Fitch argument cannot be dismantled by simply rejecting this principle of inference. This is because the only use made of the principle in the argument is to infer a special case: that the knowledge of a conjunct follows from the knowledge of a conjunction. Hence, the rule could be replaced by the much simpler:

$$\frac{K(\beta \wedge \gamma)}{K\gamma}$$

This seems much harder to contest.[2] In particular, the sorts of counter-example just mentioned relevant to the failure of the closure of knowledge under entailment (in general) seem to get very little grip on it. The knowledge of a conjunct seems *implicit* in the knowledge of a conjunction.

There is therefore little scope for faulting this part of the argument.

3.2 Stage 2: Possibility

The second part of the argument embeds Π_1 in an argument concerning possibility. This is as follows, where the right-hand column represents Π_1. Call this part Π_2.

$$\dfrac{\dfrac{\alpha \wedge \neg K\alpha}{\Diamond K(\alpha \wedge \neg K\alpha)} \qquad \begin{matrix}[K(\alpha \wedge \neg K\alpha)] \\ \nabla \\ K\alpha \wedge \neg K\alpha\end{matrix}}{\Diamond(K\alpha \wedge \neg K\alpha)}$$

[2]Harder, but not impossible. Connexivist logicians (including some medievals) held that $\beta \wedge \gamma$ does not entail γ—for example, if β is $\neg\gamma$, this simply cancels out the γ. Such a logician could know $\beta \wedge \gamma$, but not believe, and *a fortiori* know, γ. To avoid this kind of problem we can just restrict the class of knowers in question to those who have the normal beliefs about the validity of inferences concerning conjunction—which includes us.

Π_2 applies two new rules, which are as follows:

$$\frac{\beta}{\Diamond K\beta} \qquad \frac{\Diamond\beta \quad \begin{matrix}[\beta] \\ \vdots \\ \gamma\end{matrix}}{\Diamond\gamma}$$

The first of these is simply the Verification Principle, which is what the argument assumes (for the sake of *reductio*). The second says that possibility is closed under entailment. This seems to hold for any notion of possibility. If α is true in a possible world (of any appropriate kind), and α entails β, then β is true in that world, and so possible (in the same sense).

There is little in this stage of the inference that one can balk at, then.

3.3 Stage 3: Contraposition

The third part of the argument embeds Π_2 in an argument deploying negation. This is as follows, where the left-hand column represents Π_2. Call this Π_3.

$$\frac{\begin{matrix}[\alpha \wedge \neg K\alpha] \\ \nabla \\ \Diamond(K\alpha \wedge \neg K\alpha)\end{matrix} \qquad \neg\Diamond(K\alpha \wedge \neg K\alpha)}{\neg(\alpha \wedge \neg K\alpha)}$$

Π_3 employs one premise and one further rule of inference. The premise is $\neg\Diamond(\beta \wedge \neg\beta)$, or equivalently, given the usual connections between $\Box$ and $\Diamond$:

$\Box\neg(\beta \wedge \neg\beta)$

The inference is contraposition:

$$\frac{\begin{matrix}[\beta] \\ \vdots \\ \gamma\end{matrix} \qquad \neg\gamma}{\neg\beta}$$

The only plausible way to contest these steps is to suppose that contradictions may be true. The rationale for contraposition is that if β delivers something that is not true, γ, it must be false. This rationale collapses if γ can be true despite the truth of $\neg\gamma$. Unsurprisingly, then, the inference fails in many paraconsistent logics (including the one whose semantics I will describe below). Suppose, for example, that the logic contains the Law of Excluded Middle (LEM), $\beta \vee \neg\beta$. Then we have $\gamma \vdash \beta \vee \neg\beta$. Contraposing, $\neg(\beta \vee \neg\beta) \vdash \neg\gamma$, that is (assuming De Morgan Laws), $\beta \wedge \neg\beta \vdash \neg\gamma$—which fails, since γ was arbitrary. This stage of the argument may therefore be broken by appealing to dialetheism.

It might be thought that dialetheism would invalidate the new premise of the argument as well: if contradictions may be true, one might expect $\neg(\beta \wedge \neg\beta)$, and so its necessitation, to fail. Surprising as it might be to those meeting paraconsistency for the first time, it does not. There are many paraconsistent logics where the law holds (including the one whose semantics I will describe below). Of course, any contradiction, $\beta \wedge \neg\beta$, will then generate a secondary contradiction, $(\beta \wedge \neg\beta) \wedge \neg(\beta \wedge \neg\beta)$, but there is nothing in a paraconsistent logic to rule this out.

Actually, the simplest way of avoiding $\neg(\beta \wedge \neg\beta)$ (and so its necessitation) is to appeal, not to truth-value gluts, but to truth-value gaps. If β is neither true nor false, so (given the natural semantics for the connectives) is $\neg(\beta \wedge \neg\beta)$. Appealing to truth-value gaps also invalidates contraposition unless the logic is paraconsistent. If the logic is not paraconsistent, we have $\beta \wedge \neg\beta \vdash \gamma$, and so $\neg\gamma \vdash \neg(\beta \wedge \neg\beta)$, i.e., $\neg\gamma \vdash \beta \vee \neg\beta$, which is not the case if we do not have the LEM.

It might therefore be thought that appealing to truth-value gaps is a way of avoiding the argument without an appeal to gluts. Unfortunately (for the friends of consistency) it is not. As Π_2 shows, $\alpha \wedge \neg K\alpha$ already leads to $\Diamond(K\alpha \wedge \neg K\alpha)$, and thus to the possibility of true contradictions. Moreover, if the logic is not paraconsistent, we have, for an arbitrary β, $K\alpha \wedge \neg K\alpha \vdash \beta$. By the closure of possibility under entailment, we have $\Diamond(K\alpha \wedge \neg K\alpha) \vdash \Diamond\beta$. Given that $\Diamond(K\alpha \wedge \neg K\alpha)$, everything is possible—not an enticing conclusion. One way or another, then, true contradictions are required to break this step of the argument.

3.4 Stage 4: Double Negation

There is one final part of the argument. This embeds Π_3 in the argument which actually takes us from α to $K\alpha$. This goes as follows, where the right-hand column represents Π_3.

$$\dfrac{\dfrac{\dfrac{\alpha \quad [\neg K\alpha]}{\alpha \wedge \neg K\alpha} \qquad \begin{matrix}\neg\Diamond(K\alpha \wedge \neg K\alpha)\\ \nabla \\ \neg(\alpha \wedge \neg K\alpha)\end{matrix}}{\neg\neg K\alpha}}{K\alpha}$$

This stage of the argument uses contraposition again, discharging $\neg K\alpha$. (And in this application, there is also another assumption in the sub-proof. As is to be expected, this does nothing to restore validity in a paraconsistent logic. It just makes matters worse.) It uses one further rule, double negation:

$$\frac{\neg\neg\beta}{\beta}$$

Double negation fails in intuitionist logic, which is intimately connected with verificationism. Hence, breaking the argument by denying this step is a very plausible move. If we do, we can get from α only to $\neg\neg K\alpha$, which is not so bad. Well, not really. Given $\alpha \rightarrow \neg\neg K\alpha$, we obtain $\neg\neg\neg K\alpha \rightarrow \neg\alpha$ by a form of contraposition that is intuitionistically valid. And in intuitionist logic, $\neg\beta \leftrightarrow \neg\neg\neg\beta$. So by transitivity, $\neg K\alpha \rightarrow \neg\alpha$. Even intuitionists cannot accept this in general. Let α be 'Alpha Centauri has a planetary system'. I do not know that α; I do not know that $\neg\alpha$. (Nor does anybody else—maybe for ever.) It cannot follow that $\neg\alpha$ and $\neg\neg\alpha$.[3]

4 A Simple Model

We have seen that appealing to dialetheism breaks the Fitch argument against verificationism. We can do more than this, however. It can be shown that once contraposition (and only contraposition) is removed from the principles employed, the inference from α to $K\alpha$ is not forthcoming. I demonstrate this with a semantics for a simple paraconsistent modal/epistemic logic.[4]

Interpretations are of the form $\langle W, \infty, R, S, \nu \rangle$. W is a set of worlds. ∞ is a distinguished member of W. R is the modal binary accessibility relation, and we require that for every $w \in W$, $wR\infty$. S is the epistemic binary accessibility relation, which is at least reflexive. ν maps every world and propositional parameter to $\{1\}$, $\{0\}$ or $\{1, 0\}$ (true, false, both). I write the value of α at w as $\nu_w(\alpha)$. Truth conditions at worlds, w, other than ∞ are as follows:

$1 \in \nu_w(\alpha \wedge \beta)$ iff $1 \in \nu_w(\alpha)$ and $1 \in \nu_w(\beta)$

$0 \in \nu_w(\alpha \wedge \beta)$ iff $0 \in \nu_w(\alpha)$ or $0 \in \nu_w(\beta)$

[3] On this and related objections, see Percival (1990).

[4] This is an extension of the propositional paraconsistent logic LP. (See Priest (1987), ch. 5.) The existence of the trivial world, ∞, does not affect the logic of the extensional connectives.

$1 \in \nu_w(\neg\alpha)$ iff $0 \in \nu_w(\alpha)$

$0 \in \nu_w(\neg\alpha)$ iff $1 \in \nu_w(\alpha)$

$1 \in \nu_w(\Diamond\alpha)$ iff for some w' such that wRw', $1 \in \nu_{w'}(\alpha)$

$0 \in \nu_w(\Diamond\alpha)$ iff for all w' such that wRw', $0 \in \nu_w(\alpha)$

$1 \in \nu_w(K\alpha)$ iff for all w' such that wSw', $1 \in \nu_{w'}(\alpha)$

$0 \in \nu_w(K\alpha)$ iff for some w' such that wSw', $0 \in \nu_{w'}(\alpha)$

∞ is the trivial world. That is, for every α:

$\nu_\infty(\alpha) = \{1, 0\}$

Validity is defined in terms of truth-preservation at all worlds.

Leaving aside the Verification Principle for the moment, it is not difficult to check that the semantics verify all the inferences involved in the Fitch argument (including the closure of knowledge under entailment, and the premise $\neg\Diamond(\beta \wedge \neg\beta)$) except contraposition.

For the verificationist inference: for any α, $1 \in \nu_\infty(K\alpha)$, and so for every w (including ∞), $1 \in \nu_w(\Diamond K\alpha)$. The inference $\alpha \vdash \Diamond K\alpha$ is therefore (vacuously) valid.

To finish the job, we just need an interpretation where there are worlds, w_0 and w_1, such that $1 \in \nu_{w_0}(p)$, $w_0 S w_1$, but $1 \notin \nu_{w_1}(p)$. Then $1 \notin \nu_{w_0}(Kp)$. We can depict the simplest interpretation of this kind as follows (+ indicates that a formula holds; − indicates that it fails; square brackets indicate things that hold at worlds, other than what is part of the specification):

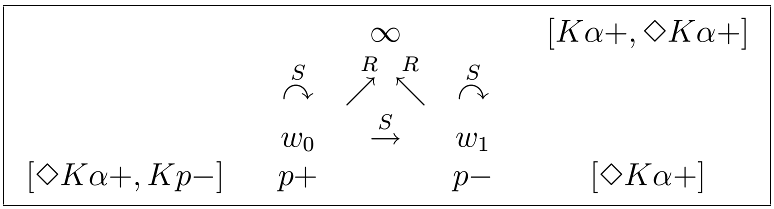

Notice that R and S can be made as strong as one likes without ruining the argument. In other words, the modal logic of K and $\Diamond$ ($\Box$) can be beefed up to $S5$ without affecting the result.

Note, also, that we may take the language to contain a conditional operator, $\rightarrow$, with strict truth/falsity conditions as follows.[5] At any world, w, other than ∞:

[5]See Priest (1987), ch. 6.

$$1 \in \nu_w(\alpha \to \beta) \text{ iff for all } w' \text{ such that } wRw', \text{ if } 1 \in \nu_{w'}(\alpha) \text{ then } 1 \in \nu_{w'}(\beta)$$

$$0 \in \nu_w(\alpha \to \beta) \text{ iff for some } w' \text{ such that } wRw', 1 \in \nu_{w'}(\alpha) \text{ and } 0 \in \nu_{w'}(\beta)$$

Assuming that R is reflexive, these semantics verify (at least) the inferences:

$$\frac{\alpha \quad \alpha \to \beta}{\beta} \qquad \frac{\begin{matrix}[\alpha] \\ \vdots \\ \beta\end{matrix}}{\alpha \to \beta}$$

(where α is the only undischarged assumption in the second inference). In the above model (with the additional proviso that R is reflexive), $\alpha \to \Diamond K\alpha$ holds for all α at w_0, but $p \to Kp$ fails.

5 Enter the Knower

We have seen that the Fitch argument may be blocked by an appeal to dialetheism. Moreover, it is the *only* way that we have found in which the argument may be blocked.[6] But—it might well be argued—an appeal to dialetheism in this context is extreme and unmotivated. Better to take the argument to be a simple *reductio* of the Verification Principle.

Matters are not that simple, though. First, there are situations in which the Verification Principle appears to hold (at least in some sense of possibility) and where the agent in question does not know everything true. It is coherent, I take it, to suppose the existence of an omniscient (and omnipotent) being. Let us call them 'God'. Everything true it is possible for God to know; indeed, everything true God actually does know. But God has a friend; call him 'Gabriel'. Gabriel is not omniscient. There are many things that Gabriel doesn't know, and doesn't care about—such as who won the 4.30 at Flemington. But Gabriel knows at least that God is omniscient. Moreover, he knows that he can always ask God if he wants to know something; God, being a decent and trustworthy fellow, will tell him. Hence, anything that it is true, it is possible for Gabriel to know—just by asking. Yet Gabriel does not know everything true. The Fitch argument must therefore fail.

The Fitch argument itself suggests an objection to this. Let us suppose that Red King Hit won the 4.30 at Flemington—call this κ—and that, as

[6]Human ingenuity being what it is, there may, of course, be other suggestions. A number of these are discussed (and rejected) in Williamson (2000), ch. 12. The chapter also contains references to other discussions of the argument in the literature.

a matter of fact, Gabriel does not know this, since he never bothers to ask. Then:

(*) κ and Gabriel does not know (at any time) that κ

is true. God knows it. It might be argued that it is, none the less, not possible for Gabriel to know it. To do so, he would have to know κ and know that he does not know κ (at any time), which is impossible.

But could he not ask God whether (*) is true, and get an answer? Of course he could. If, as we suppose, (*) is true, God will tell him so. Hence, Gabriel will know κ, and (*) is false. Suppose, on the other hand, that (*) is false. Then God will tell him so. At this point Gabriel still does not know whether κ is true or false. Suppose we then shoot him; he never will. So (*) is true.

None of this shows that Gabriel cannot know (*); all it shows is that, if he *does* ask the question, the situation is a paradoxical one. In fact, the paradox is a version of a well known one—the Bridge. A person has to cross a bridge; on the other side there is a bridge-keeper who asks a question. If the person answers truly, they are allowed to pass; if not, the bridge-keeper hangs them. The bridge-keeper asks 'what will you do when you get to the other side of the bridge?'. The person answers 'I will be hanged by you'.[7] Again, the question forces a paradoxical situation.

A much simpler version of the paradox is forthcoming by just letting κ be the sentence 'Gabriel (or even God) does not know κ'. Let us make this more precise. By applying techniques of self-reference, we can construct a sentence, κ, that says of itself that it is not known. That is, κ is of the form $\neg K\langle\kappa\rangle$. (I now revert to writing K as a predicate. Self-referential constructions require this.[8]) Suppose that $K\langle\kappa\rangle$; then κ is true, so $\neg K\langle\kappa\rangle$. Hence, $\neg K\langle\kappa\rangle$. That is, κ, but we have just demonstrated this, so it is known to be true, $K\langle\kappa\rangle$. (This is the Knower paradox.) We have demonstrated $K\langle\kappa\rangle \wedge \neg K\langle\kappa\rangle$. This is therefore necessarily true (in whatever sense of necessity one cares for); *a fortiori*, $\Diamond(K\langle\kappa\rangle \wedge \neg K\langle\kappa\rangle)$. And the Verification Principle figures nowhere in the argument for this. We see, in particular, that quite independently of the Fitch argument there are sentences of the form required to invalidate the contraposition in Π_3 . Appealing to di-

[7]The paradox is one of Buridan's sophismata but, according to Sorensen, it probably goes back to Chrysippus. A version of it is told by Cervantes in *Don Quixote*. See Sorensen (2003), pp. 207-9.

[8]In fact, we can maintain K as an operator provided that we have a truth predicate, T, in the language. We can then *define* an appropriate predicate, $K'x$, as KTx. (Thanks to JC Beall for this observation.)

aletheism to break the Fitch argument is, therefore, not at all *ad hoc* or unmotivated. In the context, it is very natural.[9]

6 Contradiction and the Limits of Knowledge

We can bring this to bear explicitly on the question of the limits of knowledge as follows. Let $X \subseteq \mathcal{K}$. Provided that X has a name, and given appropriate techniques of self-reference, we can form a sentence that says of itself that it is not in X; that is, a sentence, α_X, of the form $\langle\alpha_X\rangle \notin X$. We can show that $\langle\alpha_X\rangle \notin X$ but that $\langle\alpha_X\rangle \in \mathcal{K}$ as follows:

$$\begin{aligned} \langle\alpha_X\rangle \in X \quad &\Rightarrow \quad \langle\alpha_X\rangle \in \mathcal{K} \\ &\Rightarrow \quad \alpha_X \\ &\Rightarrow \quad \langle\alpha_X\rangle \notin X \end{aligned}$$

Hence, $\langle\alpha_X\rangle \notin X$. But this is α_X, and we have just established this, so it is known to be true; that is, $\langle\alpha_X\rangle \in \mathcal{K}$. The situation may be depicted thus:

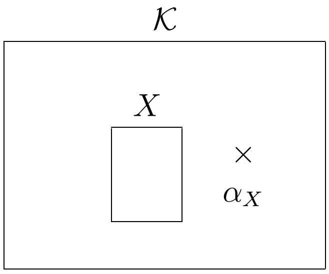

When X is the empty set, α_X can be located anywhere in $\mathcal{K} - X$ $(= \mathcal{K})$. As X gets bigger and bigger,[10] there is less and less space in which α_X can be consistently located; until, at the limit, when X coincides with $\mathcal{K}$ there is nowhere consistent for α_X to go. $\langle\alpha_{\mathcal{K}}\rangle \in \mathcal{K} \wedge \langle\alpha_{\mathcal{K}}\rangle \notin \mathcal{K}$. (This is the Knower paradox. κ is just $\alpha_{\mathcal{K}}$.) The limit of what is known is dialetheic. That is, there are certain truths that are both within the known and without it.

Exactly the same is true of $\mathcal{P}$. Let $X \subseteq \mathcal{P}$. As before, we can construct a sentence, α_X, of the form $\langle\alpha_X\rangle \notin X$.

[9] The first person to moot the possibility of a connection between the Fitch argument and the Knower was Routley (1981). (See esp. p, 112, n. 26.) The connection was made more robustly by Beall (2000).

[10] Of course, X does not literally grow. In particular, we are not considering the case where more and more is known. (That would be a case of K growing.) This is just a picturesque way of saying that for larger and larger X...

$$\begin{aligned}\langle\alpha_X\rangle \in X &\Rightarrow \langle\alpha_X\rangle \in \mathcal{P}\\ &\Rightarrow \langle\alpha_X\rangle \in \mathcal{T} \wedge \Diamond K\langle\alpha_X\rangle\\ &\Rightarrow \alpha_X\\ &\Rightarrow \langle\alpha_X\rangle \notin X\end{aligned}$$

Hence, $\langle\alpha_X\rangle \notin X$. But this is α_X, and we have just established this, so it is true and known to be so, $K\langle\alpha_X\rangle$. *A fortiori*, it is possible to know it, $\Diamond K\langle\alpha_X\rangle$. Thus, $\langle\alpha_X\rangle \in \mathcal{T} \wedge \Diamond K\langle\alpha_X\rangle$. That is, $\langle\alpha_X\rangle \in \mathcal{P}$. Just as with $\mathcal{K}$, when X is small, there is plenty of room for α_X to reside, consistently, outside it but inside $\mathcal{P}$. As X gets bigger and bigger, there is less and less room, until when X is $\mathcal{P}$, a contradiction arises: $\langle\alpha_{\mathcal{P}}\rangle \in \mathcal{P} \wedge \langle\alpha_{\mathcal{P}}\rangle \notin \mathcal{P}$. The boundary of possible knowledge is inconsistent too.

An *Inclosure* involving a set, Ω, a predicate, ψ, and a function, δ, is a structure satisfying the following conditions:

1. $\psi(\Omega)$
2. if $X \subseteq \Omega$ and $\psi(X)$
   (a) $\delta(X) \notin X$ (Transcendence)
   (b) $\delta(X) \in \Omega$ (Closure)

Whenever we have an Inclosure, a contradiction arises at the limit, when $X = \Omega$. For we then have $\delta(\Omega) \notin \Omega \wedge \delta(\Omega) \in \Omega$. All the standard paradoxes of self-reference are limit-paradoxes of this kind.[11]

The two contradictions we have just looked at are of this form. In the first, Ω is $\mathcal{K}$; in the second, Ω is $\mathcal{P}$. In both, $\psi(X)$ is 'X is definable (has a name)', and $\delta(X)$ is α_X. Hence, both are inclosure contradictions.

[11]See Priest (1995), Part 3. For the Knower paradox, see 10.2. There, Ω is defined as $\{x : \varphi(x)\}$, where φ is the appropriate predicate.

7 Conclusion

Let us recall our original diagram, and take stock:

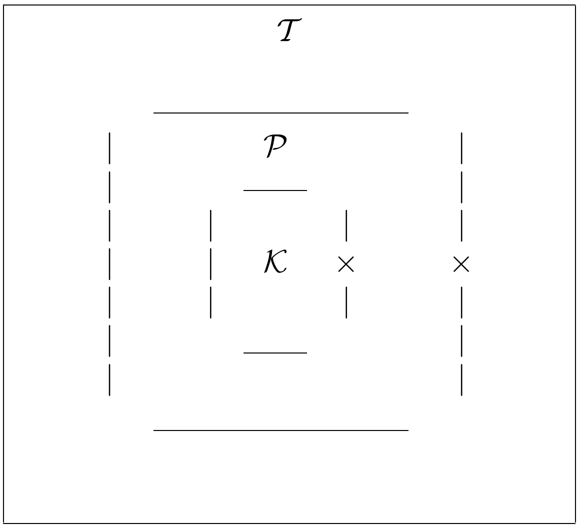

$\mathcal{K}$, we know, is non-empty, as is $\mathcal{P} - \mathcal{K}$. And this is so if $\mathcal{T} - \mathcal{P}$ is empty, and so the Verification Principle is correct, since the Fitch argument fails. We have also learned that the boundaries between $\mathcal{K}$ and $\mathcal{T} - \mathcal{K}$, and $\mathcal{P}$ and $\mathcal{T} - \mathcal{P}$ are dialetheic. That is, there is a true sentence, $\alpha_{\mathcal{K}}$, such that $\langle \alpha_{\mathcal{K}} \rangle \in \mathcal{K}$ and $\langle \alpha_{\mathcal{K}} \rangle \notin \mathcal{K}$, and a true sentence, $\alpha_{\mathcal{P}}$, such that $\langle \alpha_{\mathcal{P}} \rangle \in \mathcal{P}$ and $\langle \alpha_{\mathcal{P}} \rangle \notin \mathcal{P}$. (This is what the '$\times$'s on the new version of the diagram indicate.) And since $\alpha_{\mathcal{P}}$ is true, $\langle \alpha_{\mathcal{P}} \rangle \in \mathcal{T} \wedge \langle \alpha_{\mathcal{P}} \rangle \notin \mathcal{P}$, so $\mathcal{T} - \mathcal{P}$ is also non-empty. For all I have said, this might be its only denizen. It cannot, therefore, be ruled out that $\mathcal{T} - \mathcal{P}$ is empty as well (which it is if Verification Principle is correct). Whether or not this is so might well depend on the sense of possibility at issue. It is, at any rate, a matter for another occasion.[12]

References

[1] Beall, JC (2000) 'Fitch's Proof, Verificationism, and the Knower Paradox', *Australasian Journal of Philosophy* 78, 241-7.

[12] Ancestors of this paper were given under the title 'The Limits of Knowledge' at Mt Holyoke College, the Graduate Center, City University of New York, the University of California at San Diego, the University of Melbourne; and at the conferences *The Philosophy of Uncertainty: Epistemic Limits, Probability, and Decision*, held at the University of Tokyo, and *Logica 2005*, held in the Czech Republic. I am grateful to the audiences present for their helpful discussions, and to Masake Ichinose, Tim Childers and Vladimir Svoboda, for organising the conferences.

[2] Fitch, F. B. (1963) 'A Logical Analysis of some Value Concepts', *Journal of Symbolic Logic* 28, 135-42.

[3] Hart, W. D., and McGinn, Colin (1976) 'Knowledge and Necessity', *Journal of Philosophical Logic* 5, 205-8.

[4] Percival, Philip (1990) 'Intuitionistic Knowability', *Analysis* 60, 182-7.

[5] Priest, Graham (1987) *In Contradiction*, Dordrecht: Martinus Nijhoff. Second edition, Oxford: Oxford University Press, 2006.

[6] Priest, Graham (1995) *Beyond the Limits of Thought*, Cambridge: Cambridge University Press. Second edition, Oxford: Oxford University Press, 2002.

[7] Routley, Richard (1981), 'Necessary Limits to Knowldege: Unknowable Truths', pp. 97-118 of E. Morscher (ed.), *Essays in Scientific Philosophy*, Bad Reichenhall: Comes Verlag.

[8] Sorensen, R. (2003), *A Brief History of the Paradox*, Oxford: Oxford University Press.

[9] Williamson, Timothy (2000) *Knowledge and its Limits*, Oxford: Oxford University Press.