

Fitch's Paradox

James R. Shaw

To know that p requires no less than that p obtain—a phenomenon referred to as the “factivity” of knowledge. Fitch's paradox capitalizes on this characteristic factivity of knowledge to argue that there are structural limits on possible knowledge about the world: things *in principle unknowable*.

The paradox consists in the derivation of the existence unknowable truths from the existence of unknown truths.¹ The argument proceeds roughly as follows. Suppose there is a fact unknown* in a special sense (signalled by the asterix)—unknown to anyone at any time. Let this fact be, say, that I snored last night. Then what is expressed by the following conjunction, which we may call a Fitch sentence, is unknowable: “James snored last night and that is unknown*”. To see this suppose, for *reductio*, that someone in hypothetical circumstances came to know it. Then they would know that I snored last night. But in such circumstances it would clearly be known* that I snored last night. So the second conjunct of our Fitch sentence would no longer be true and hence neither would the conjunction itself. Thus, by the factivity of knowledge, it would not be the sort of thing which *could* be known, and we have a contradiction.

The paradox is alternatively formulated as part of an argument against the doctrine of *weak verificationism*, comprising the claim that everything true is in principle knowable. If every truth is knowable then, since any unknown* truth would be unknowable by the above argument, we may conclude that every truth is known*. This result seems untrue, and at least is not true *a priori*. Thus, it is argued, weak-verificationism collapses into an untenable position. In some sense, the reduction of weak verificationism to absurdity in this way depends on the plausibility of the first formulation of the paradox and it is this that I would like to make my present focus.

Several avenues of resolution to the paradox have already been explored. One has been in restricting inferential principles used in the paradoxical deduc-

¹The *locus classicus* for the paradox is Fitch (1963).

tion, such as that on which knowledge of a conjunction entails knowledge of its conjuncts. These have not met with much success. Another more promising avenue takes there to be an implicit restriction of the range of propositions quantified over in certain relevant claims of knowability—for example when it is claimed that “everything is knowable”. Views of this form have been advanced by both Dummett (2001) and Tennant (2002). I would like to propose a version of such a view which I will bolster by appeal to intuitions about when quantified statements involving “knowability” would be truthfully asserted.

Fitch’s paradox is often framed in modal logic with a knowledge operator ‘K’ on propositions (attributing knowledge* of them) and quantifiers ranging over propositions. In the context of this logic one can prove by seemingly reasonable steps the formula

$$(1) (\exists p)(p \wedge \neg \diamond Kp)$$

from

$$(2) (\exists p)(p \wedge \neg Kp).$$

The deduction can disguise the sophisticated air of the paradox hinted at with any particular example, such as the snoring case above. The paradox relies on the idea that the more one comes to know, the more various claims about ignorance, among others, become false: thus whatever places one in a condition to know the first conjunct in a Fitch sentence will also falsify the sentence as a whole and render it unsuited to be knowledge. The paradox in effect asserts the unknowability of given facts by noting that in hypothetical circumstances favorable to knowledge generally, they cease to *be* facts. But being unable to know various falsehoods in such circumstances hardly seems a genuine strike against one’s possible epistemic state.

Here is a way of pressing these worries: it is clearly logically and conceptually compatible with the formalized statement (1) (barring a more convincing case for structural or other unknowability) that there is a world with an omniscient being—someone knowing all the facts about their world, including the modal facts. There is an intuitive sense in which we might maintain that such a possibility would be enough on its own to secure the truth of the claim that everything is possible to be known.

We need to be slightly more careful, though. This hypothesized world with an omniscient being could easily be highly different from our own. It might be claimed: “Surely it is not enough to point to this scenario to show that it is

possible to know everything there is to know. We are eager to discover whether it is possible to know what is going on *in our world* or at least a world very much like it”.

The motivation for this reply is well-founded, but does not get to the heart of the objection. To flesh out our example, suppose there were a world w_a with a ‘Tree of Knowledge’ that bore special fruit: eating it gives one complete knowledge of the world one is in. Suppose in w_a one person said to another, who would eventually be the only person to partake of the fruit, “it is possible to know everything there is to know. After all, just eat the fruit from the Tree of Knowledge”. Suppose we were to render the logic of their initial statement as

$$(3) (\forall p)(p \rightarrow \diamond Kp).$$

Then the speaker in such circumstances would have spoken the truth. Indeed, everything will not only be possibly, but actually known. But now imagine a second world w_b *exactly* like the first save that the speaker’s interlocutor is never properly moved to eat the fruit in question. In these circumstances the utterance, if interpreted in the same way, would likely come out *false*. So much the worse, one would think, for that interpretation.

The utterance in w_b *on the interpretation in question* is potentially false, despite the *nearly identical* neighboring world w_a where the fruit is eaten, creating an ostensibly omniscient being, because of potentially unknown* truths in that world like the following: “no one eats the fruit of knowledge, a fact itself unknown*”. Of course, were the fruit eaten that would cease to *be* a fact to be known. But the modal statement above makes no such distinctions since all and only *actual* facts satisfy the antecedent of the conditional statement.

The point of this example is not, of course, that we might ever be in such a fantastical world but that the statement “every truth is knowable” cannot be captured by the logical rendition proposed. For there is an intuitive sense in which the utterance in w_b would be true precisely (or at least in part) because of the near actualizing of w_a . This suggests, it seems, that a proper rendering of the conditions for the truth of that statement would take into account how this occurs. One very plausible story seems to be that certain ‘facts’ are being excluded from consideration in evaluating the claim. Many true propositions which would become false as knowledge was obtained would fall into this category.

As I mentioned already, proposals to explicitly restrict the domain of quantification in the modal statement have already been proposed by both Tennant and Dummett. In fact these proposals try to specify, though syntactic or other means, *exactly which* propositions are quantified over. To that extent they have been subject to strong criticisms that the restriction is unprincipled.²

But from the present perspective Dummett and Tennant's responses can seem to go too far. The point of the above example is to shift the burden of proof back onto the purveyor of the paradox. We now know that we can come to accept (1): it is true, but does not present us with an intuitively recognizable and controversial claim about possible knowledge. Our ordinary claims about the general achievement of possible knowledge factor out certain propositions, including those that would tend to become false as we acquire more knowledge. How we restrict the domain in those general claims is obviously a complicated matter that is witnessed in actual cases, such as the fantastical one given by w_a and w_b . But we are not beholden to be able to *systematize* that restriction in order to answer the paradox. Indeed, there may even be no obvious rule-governed way to represent which propositions we take to be relevant to assessing a claim like "every truth is knowable".

Note that I do not merely mean, in invoking the "intuitive" understanding of certain possible knowledge claims, to fall back on ordinary linguistic practice as an unassailable source of truths about knowledge and modality. The point is that w_a and w_b present us with a case which we would have thought (surprisingly) to have been ruled out when presented with the conclusion of Fitch's paradox as originally stated. Seeing that this is not so robs the paradox of its force. This is the sense in which the example above is a call for a supplementation to Fitch's paradox: any proof of the claim that the existence of unknown truths entails the existence of unknowable truths *should show us that finding a pair of worlds like w_a and w_b is logically or conceptually impossible*. Anything short of that will, I think, fail to have demonstrated that claim on its most significant interpretation.

It might be urged that although the reply given in the preceding paragraphs is available to some, it is not available to the anti-realist who endorses weak verificationism. If one wishes to *positively* endorse a claim about possible knowledge, one ought to specify exactly what propositions one is quantifying over in the interest of full disclosure. Otherwise how are we to know what would

²For example, in Hand & Kvanvig (1999).

count for or against that claim? It is not my purpose to defend the anti-realist here, but I would like to make some quick remarks about why I suspect even she might be unaffected by Fitch’s paradox.

Note first that *if* there is an intuitive sense in which “all truths are knowable” is stably assessable across speakers (say that they recognize in given cases what would and would not count as a counterexample to that claim), there need be no harm in falling back on that interpretation—whatever it is and however the domain is (perhaps contextually) restricted. Additionally, however, the example from w_a and w_b highlights a *feature* of Fitch’s paradox on account of which it would be surprising for it to constitute an objection to *any* reasonable philosophical position. This can perhaps be brought out by a more realistic case.

Suppose one amateur philosopher-scientist were speaking to another, explaining a simplified version of the verificationist position. The first scientist claims that it is in principle possible for them to know everything about the state of their lab at a given time—the positions of the instruments, the behavior of their bacteria samples, and so on—by observation. The second scientist retorts: “Well, given that we aren’t observing the bacteria in that petri dish right now, that can’t be right. Here’s why: Suppose that the microbes were just preparing to undergo fission. Then there is clearly something in principle unknowable by observation about the state of the lab, namely that *that the microbes were just preparing to undergo fission and no microscopes were near the bacteria*. We couldn’t know *that* by observation.” The first scientist is confused, so the second continues “well suppose one of us knew that the microbes were about to undergo fission by observation. To know that, one of us must have looked under a microscope at that appropriate time. But then we couldn’t know the fact I mentioned by observation—for then we would know there were no microscopes near the bacteria, when of course there were!” The first scientist, no doubt, would have a hard time seeing this as an objection to her claim, *however* she meant it.

The sense in which the scientists ‘couldn’t know’ the lab-state in question is of no particular interest. The interpretation required of the first scientist’s claim required to make that ‘unknowability’ pertinent to it is *extremely* irregular. Fitch’s paradox, which is structurally analogous to the second scientists’ claim, is likely just as uninteresting to the verificationist, and for the same reasons: no *reasonable* talk of possible knowledge or knowledge by observation will come into conflict with it. The anti-realist *might* have more work to do in spelling

out her position, but reflection on Fitch’s paradox needn’t specially prompt that endeavor and doubtless won’t significantly constrain its outcome.

Quine had a name for paradoxes such as Fitch’s. In “Ways of Paradox” he raises the puzzle of Frederick from the *Pirates of Penzance* who has reached the age of 21 before his sixth birthday. The resolution of the paradox comes in recognizing that age is chalked up in days lived while birthdays are tallied by the advent of certain calendar days, with leap-years slowing the count for Frederick. This is an example of what Quine calls a *veridical* paradox—a seemingly outlandish statement which, on due reflection, we have reason to accept. In fortunate circumstances the initially implausible paradox is, however, only that: *initially* implausible. Understanding the basis for the ‘paradoxical’ claim’s truth not only reveals grounds for accepting the claim but also for rejecting it as an affront to common sense. I suggest we lump Fitch in with Frederick in that regard.

References

- M. Dummett (2001). ‘Victor’s error’. *Analysis* **61**(269):1–2.
- F. Fitch (1963). ‘A Logical Analysis of Some Value Concepts’. *The Journal of Symbolic Logic* **28**(2):135–142.
- M. Hand & J. Kvanvig (1999). ‘Tennant on knowability’. *Australasian Journal of Philosophy* **77**(4):422–428.
- W. Quine (1979). *The Ways of Paradox*. Harvard University Press.
- N. Tennant (2002). *The Taming of the True*. Oxford University Press.
- T. Williamson (2002). *Knowledge and Its Limits*. Oxford University Press.