Knowability and the Capacity to Know

1. Introduction

Are there any unknowable truths? There are plenty of truths that are not in fact *known*; but are there any truths that are not in principle *knowable*, even by beings with greater minds than ours who have more time and energy than we have to devote to research and investigation? Some philosophers, the verificationists or anti-realists, think that the answer is "no", that there are not any unknowable truths (nor could there be any). Others, the realists, think that the answer is "yes". Since neither answer is obviously correct, and since the question is an interesting one, it is no surprise that the debate has received the attention that it has over the years. It *is* a surprise, however, especially in light of the attention that the debate has received, that the verificationist view can apparently be shown, by a swift and simple argument, to be absurd. In fact the argument, first published by F. B. Fitch (1963), is so swift that when viewed against the heady background of the realist/anti-realist debate it seems that it must contain a flaw; but equally, the argument is so simple that it seems impossible for it to contain a flaw. This is the "paradox" of knowability.

Fitch's argument, in one of its versions,[1] begins by supposing (for reductio) that the verificationist is correct: Any truth could in principle be known. Because not every truth is in fact known, the argument continues, there are truths of the form '*p* and it's not known that *p*'. Since the verificationist thinks that *any* truth could be known, she is committed to thinking that truths of this conjunctive form could be known. But such truths *couldn't* be known, since knowing one would require knowing both of its conjuncts, which is impossible. For one could know the second conjunct only if it is true; the second conjunct is true only if the first conjunct is unknown; and so one could know the second conjunct only if one didn't know the first. So the verificationist claim that any truth could in principle be known has apparently been refuted, given the existence of some truths that are in fact unknown.

All manner of responses to Fitch's argument have been offered, although there is no

[1] Fitch in fact only published one version of his argument. But others have presented arguments so similar to Fitch's that they are plausibly thought of as versions of it.

consensus on which response, or even which kind of response, is correct. Some philosophers—the revisionists—blame the reasoning, and take Fitch's argument to demonstrate either that the verificationist should reject one or more of the rules of inference used in the argument, or that the principles about knowledge appealed to in the argument are flawed. Some philosophers—the restrictionists—think the verificationist can succeed with a retreat, and should hold that their claim that any truth could be known is, was, or should be restricted, perhaps tacitly, to apply only to truths of a certain sort. Still others—the rejectionists—take Fitch's argument at face value, and see it as a surprisingly simple refutation of a false philosophical thesis. And others—the reinterpretivists—think the argument doesn't get off the ground because the verificationist claim that any truth could be known doesn't mean what Fitch's argument takes it to mean.[2]

The ultimate aim of this paper is to contribute a new response of the reinterpretivist kind. But it will help, to begin, if we see why Fitch's argument at bottom has nothing in particular to do with verificationism, anti-realism, or even knowledge. For the argument can be generalized, and as a result any solution to the "paradox of knowability" must also be capable of being generalized. As we will see, this severely constrains the range of potential solutions.

2. The paradox formalized and generalized

The version of Fitch's argument presented above was informal, and some of its important assumptions and inferential steps were correspondingly buried or missing. A slightly more careful version of the argument begins with a representation of the verificationist's principle as the schema

(V) $P \rightarrow \Diamond KP.$

Here '$\Diamond$' is the standard modal operator 'it is metaphysically possible that…', 'K' is the

[2] For a useful survey of the literature that has been generated by Fitch's argument, see Brogaard & Salerno (2004), while noting that their categorization of responses to the argument is different from that given here. For a version of the first kind of revisionist strategy, see Williamson (1982); for a refutation of the second kind of revisionist strategy, see Williamson (1993). Tennant (1997, chptr. 8) offers a much-discussed version of a restrictionist strategy. Williamson (2000, chptr. 12) is one of the few rejectionists in the literature. Edgington (1985) is a classic reinterpretivist; Restall (2004) is a very different kind of reinterpretivist.

epistemic operator 'it is known by someone at some time that. . . ',[3] and 'P' is a schematic letter, replaceable by any declarative English sentence.[4] The argument continues with the claim that there are some truths that are in fact unknown. We would be hard pressed to produce a specific example, of course, since any reason we had for thinking that a proposition was unknown would itself be reason for questioning whether it was true. But we shouldn't need a specific example in order to be convinced that there are unknown truths. For some number *n* it is true that there are exactly *n* french fries on my plate. But since no-one has counted, and since no-one will count before I finish my meal, it is also true that no-one at any time knows that there are exactly *n* french fries on my plate. So there is a sentence q, involving the numeral for the appropriate number *n*, for which the following is true:

(1) $q \mathbin{\&} \neg Kq$.

Substituting (1) for 'P' in (V) yields

(2) $(q \mathbin{\&} \neg Kq) \rightarrow \Diamond K(q \mathbin{\&} \neg q)$,

which together with (1) implies

(3) $\Diamond K(q \mathbin{\&} \neg Kq)$.

Necessarily, a conjunction is known only if each conjunct is known. So if it's possible that a conjunction is known, as (3) says it is, then it's possible that each conjunct is known:

(4) $\Diamond(Kq \mathbin{\&} K\neg Kq)$.

Necessarily, whatever is known is true. Applied to the second conjunct in (4) this fact gives us

(5) $\Diamond(Kq \mathbin{\&} \neg Kq)$,

[3] Since Fitch's argument requires the existence of unknown truths, believers in divine omniscience might prefer to read 'K' as 'it is known by some non-omniscient being at some time that. . . '. See Kvanvig (1995) for a discussion of this point.

[4] I am taking English to include the expressions '$\Diamond$' and 'K', the standard Boolean connectives, and punctuation. Since there are more truths than there are sentences of English to express them, (V) is a particularly weak representation of the verificationist's principle. Fitch's apparent refutation of it is therefore correspondingly strong, which is one reason to prefer this schematic presentation of the principle to the more usual presentations which involve propositional quantification.

which expresses the possibility of a contradiction. Since contradictions are impossible, the argument concludes, schema (V) has false instances.

This version of Fitch's argument is little more than the first, informal version dressed up in abbreviatory symbolic form. But enough clutter has been cleared away for the argument's assumptions to be highlighted. First, there is the assumption that the verificationist's principle that every truth could be known may be represented as the schema (V). This assumption will be my target in what follows. Second, there is the implicit assumption that $\Diamond$ is the dual of a necessity operator which obeys a classical normal modal logic. I accept this assumption and will not discuss it further here. Third, there is the assumption that K is factive, and that it distributes over conjunctions. I accept this assumption too, although we will briefly address the question of its significance below. Finally, there is the assumption that there are true sentences of the form '$q \,\&\, \neg Kq$'. I take this to be less of an assumption than a fact demonstrated by the discussion of french fries above.

Since we are accepting a background of classical normal modal logic, Fitch's argument may usefully be viewed as an uncontrovertible *proof* of a relatively simple theorem:

FITCH'S THEOREM. There is no sentential operator O satisfying all four of the principles

1. $P \rightarrow \Diamond OP$;
2. $\Box(OP \rightarrow P)$;
3. $\Box(O(P \,\&\, Q) \rightarrow (OP \,\&\, OQ))$; and
4. There are true sentences of the form '$(q \,\&\, \neg Oq)$'.

The paradox of knowability can be thought of as an application of Fitch's Theorem. Taking K as an instance of O in the theorem, it seems obvious that principles 2, 3 and 4 are true; the proof of Fitch's theorem is sufficiently simple that it should also be obvious—now that this proof is before us—that principle 1 is false; yet the verificationist claim that every truth could be known, which seems to be what is expressed by principle 1, is *not* obviously false.

This way of looking at the paradox of knowability highlights two important points. First, the issue is not directly that of whether or not the verificationist claim is false. Rather, the issue is that of whether or not the verificationist claim is *obviously* false. The well-worn joke in which a mathematics lecturer, when asked by a student whether a certain claim is obvious, pauses to think for forty minutes before finally answering, "Yes, it is obvious," exploits the fact that obviousness is hard to measure. But the fact that philosophers have for *decades* argued on both sides of the verificationist claim should incline us to think that the claim is not obviously false. The rejectionist solution to the paradox of knowability, according to which Fitch's Theorem constitutes a simple refutation of verificationism, doesn't sit well with this point.

The second point brought out by our way of looking at Fitch's argument is the one advertised earlier, namely that it has nothing in particular to do with verificationism. For similar paradoxes arise with operators other than K taken as instances of O in Fitch's Theorem. Consider the operator 'it will eventually be revealed by the Oracle that…', where the Oracle is a possible being that reveals only truths. Assuming that the Oracle doesn't in fact exist, principles 2, 3 and 4 of Fitch's Theorem are clearly satisfied by this operator. Yet it doesn't seem *obvious* that there are truths which could never be revealed by the Oracle, as Fitch's argument would have us conclude. Or consider Williamson's (1993) operator 'it is T-conceived that…', where a proposition is T-conceived iff it is both true and conceived of by someone. Again, principles 2, 3 and 4 of Fitch's Theorem seem obviously satisfied, yet it seems equally *un*obvious that there are truths that could not be T-conceived.

These two operators are so similar to K that their differences from it might seem superficial in the present context. There are plenty of other examples, however, among them operators that aren't even epistemic. Fitch himself describes one, in his "Theorem 3":

> If an agent is all-powerful in the sense that for each situation that is the case, it is logically possible that that situation was brought about by that agent, then whatever is the case was brought about (done) by that agent. (Fitch 1963, p. 138)

Let 'God' name such an agent, and consider the operator 'God brought it about that…'. Fitch's Theorem apparently shows that if God is all-powerful then every truth was in fact brought about by God, a surprisingly swift demonstration of a version of theological determinism. Or consider

Peter van Inwagen's operator N, where 'Np' is to be read 'p, and no-one has, or ever had, any choice about whether p' (van Inwagen 1983). Fitch's Theorem seems to show that if any truth *could* have been fixed then every truth is *in fact* fixed (where a truth is fixed iff no-one has, or ever had, any choice about it). Again, this is a surprisingly quick derivation of a version of fatalism from the interesting thesis that no truth is essentially unfixed.

Another example mentioned by Fitch is worth exploring:

> If there is some true proposition about proving that nobody has ever proved or ever will prove, then there is some true proposition about proving that nobody can prove. (Fitch 1963, p. 139)

This is an application of a restricted form of Fitch's Theorem, where principles 1, 2 and 3 in the theorem are restricted to sentences expressing propositions about proving, and where the domain of the quantifier in 4 is similarly restricted. Although this particular case is unilluminating (it is not surprising that there are true but unprovable propositions about proving), there are other examples involving restricted forms of the Theorem that need explaining just as much as the original paradox of knowability.

Consider, for example, the claim that every empirical truth could be discovered *a posteriori*. If we take 'empirical truth' to *mean* 'truth that is discoverable *a posteriori*', then Fitch's Theorem seems to show that there are no empirical truths of the form 'q and no-one has discovered or will discover *a posteriori* that q', which seems absurd. Or consider two applications of Fitch's Theorem working in tandem against the pair of principles

(H1) Every matter of fact is discoverable *a posteriori*.

(H2) Every relation of ideas is knowable *a priori*.

Fitch's Theorem applied to these cases would seem to show, quite trivially, that Hume was wrong when he claimed that

> All the [true] objects of human reason or enquiry may naturally be [exhaustively]

divided into two kinds, to wit, *Relations of Ideas*, and *Matters of Fact*. (Hume 1777, IV. I.)

For, taking (H1) and (H2) as definitional of the expressions 'matter of fact' and 'relation of ideas' respectively, let q be a matter of fact that is actually undiscovered (say, the matter of fact concerning the number of french fries on my plate). Since (H1) is true, an application of Fitch's Theorem in restricted form shows that 'q and it is not discovered *a posteriori* that q', although true, is not a matter of fact.[5] If Hume were right in his claim, then this conjunction must be a relation of ideas and hence knowable *a priori*. But if a conjunction is knowable *a priori* then so is each conjunct. In particular, therefore, q would be knowable *a priori*, contradicting our assumption that q was a matter of fact.

Other applications of Fitch's argument emerge once we recognize, as Timothy Williamson (2000) and others have recognized, that principle 2 of Fitch's Theorem is stronger than is required for that Theorem's proof. Instead of the factivity of the operator O, all that Fitch's argument relies on is a weaker principle (WF for "Weak Factivity"):

(WF) $\Box(O\neg OP \to \neg OP)$.

In the case of knowledge this principle clearly holds, for the simple reason that knowledge is factive. But are there any *non*-factive operators that obey (WF) in addition to apparently obeying the other three principles mentioned in Fitch's Theorem?[6]

One candidate weakly factive operator is 'it is rational to believe that...', where the idea might be that part of being rational involves being correct in one's beliefs about what it's not rational for one to believe.[7] Provided the sense of 'rational' being used here is sufficiently

[5]Fitch's Theorem does not by itself establish this, since it is not plausible to suppose that the operator 'it is discovered *a posteriori* that...' distributes over conjunctions. But it does distribute over conjunctions of the form 'p and it is not discovered *a posteriori* that p' when p expresses an empirical truth, which is all that is needed in the present context.

[6]Dorothy Edgington (1985) argues that the factivity of knowledge is not required for Fitch-like reasoning to threaten verificationism, since the same kind of reasoning can be marshalled against forms of verificationism stated in terms of what is believed or what is supported by evidence. Jonathan Kvanvig (1995) disagrees. As Kvanvig effectively points out, however, this debate is orthogonal to the point that only a weakened version of factivity is required for the proof of Fitch's Theorem.

[7]See Tennant (1997) for a developed suggestion along these lines, and Williamson (2000, pp. 274–5) for some

idealized, this is a plausible idea. At the very least it is plausible to suppose that there could be some kind of idealized agent—a 'superagent'—who is never wrong about what she doesn't believe. The operator 'it is believed by a superagent that. . . ' will then be weakly factive, and so Fitch's Theorem would seem to show, bizarrely, that there are some true propositions that could not be believed by a superagent.[8]

For a less fanciful example of an application of Fitch's Theorem using a weakly factive operator, consider Robert Stalnaker's explanation of the phenomenon of speaker presupposition in terms of the notion of a proposition's being taken for granted by a group of conversational participants as part of the "common ground" among them.[9] It is not in general true that if so-and-so is part of the common ground then so-and-so is in fact the case: the operator, 'it is part of the common ground that. . . ' is not a factive operator. But in the special case where so-and-so is itself a fact about what is not part of the common ground, it seems that this should be so.[10] For suppose that all of us in a conversation are taking it for granted that it's not part of the common ground that p. That is to say, we are each accepting that it's not part of the common ground that p; we each believe this about the others; we each believe that we each believe this about the others; and so on up.[11] Our mutual belief system would have to be very defective—too defective for a successful conversation to ensue—if, despite what we are taking for granted, it is in fact part of the common ground that p. So the operator, 'It is part of the common ground that. . . ' is weakly factive. Since this operator also distributes over conjunctions, and since in any conversation there are true propositions that are not being taken for granted as true, Fitch's Theorem apparently yields the astonishing result that there are true propositions that *could not* be taken for granted in a conversation.

---

skeptical remarks. The idea of exchanging belief for knowledge in the context of Fitch's argument was first suggested to me by Gideon Rosen.

[8]I am assuming that a superagent's beliefs distribute over conjunctions and that there could be a superagent who doesn't in fact believe every true proposition. See Fara (2007) for details.

[9]See Stalnaker (1970), Stalnaker (1973), Stalnaker (1974) and Stalnaker (2002).

[10]Indeed, it is provably so in Stalnaker's (2002) logic of common belief, which he assumes will be the same as the logic of common ground.

[11]See Stalnaker (2002, p. 716). Note that 'accept' is being used here in Stalnaker's technical sense; see Stalnaker (1984).

We have now seen enough examples, I take it, to be confident that Fitch's Theorem has applications that extend considerably beyond the dispute between realists and anti-realists about the limits of possible knowledge. This fact, as well as being interesting in its own right, has a bearing on what might count as a solution to the paradox of knowability. In particular, rejectionist and restrictionist solutions appear far less plausible than they would be if Fitch's Theorem were just about knowability.

According to the rejectionist, Fitch's Theorem amounts to a demonstration that there are truths that could not be known, and hence that anti-realism of the appropriate sort is simply mistaken.[12] Perhaps, indeed, Fitch's Theorem does show this; but that cannot be the end of the matter. For the rejectionist will also have to maintain, if her solution to the paradox is to be a uniform one, that Fitch's Theorem is a demonstration that some truths could not be taken for granted in a conversation, and that Hume was wrong to think that any truth is either a matter of fact or a relation of ideas, and so on for the other cases.[13] Although such a sweepingly dismissive view is not provably mistaken, it is highly unattractive.

Restrictionist solutions to the paradox of knowability agree with the rejectionist that Fitch's Theorem shows that not every truth could be known, but they add that the anti-realist should never have thought that *every* truth could be known, only that every truth from some appropriate class of truths could be known. The challenge here, of course, is in specifying this appropriate class of truths. The restrictionist must ensure that the class is large enough for interest in the realist/anti-realist debate to be preserved (even the realist will happily grant that every truth expressed by a sentence in today's newspaper could be known), yet not so large as to include those truths that, the rejectionist concedes, are shown by Fitch's Theorem to be unknowable.

The most widely discussed restrictionist proposal is that of Neil Tennant (1997), who argues that the anti-realist's view should be construed as the claim that every *Cartesian* truth could be known, where p is Cartesian iff Kp is consistent. There has been considerable debate

[12]See, for example, Williamson (2000, chptr. 12).

[13]Of course, one could be a rejectionist in the case of knowability and offer a different kind of account for the other cases. But it is hard to see what would motivate such a hybrid view.

about whether Tennant's proposal amounts to anything more than the apparently tautologous claim that every *knowable* truth could be known,[14] but the present discussion raises concerns of a different kind. How is Tennant's proposal about knowability to be extended to the other cases we have seen in which Fitch's Theorem can be applied to yield counterintuitive results?

Fitch's Theorem, in the case of the knowledge operator K, shows that the schema

(V) $P \to \Diamond KP$

has false instances. Tennant's strategy exploits the fact that if the schematic letter 'P' in (V) is restricted to range over sentences for which KP is consistent, then Fitch's Theorem does not show this; and the strategy succeeds as a solution to the paradox of knowability to the extent that it is plausible to suppose that the verificationist, in saying that every truth could be known, means that every *Cartesian* truth could be known. In the general case of an operator O satisfying principles 2, 3 and 4 of Fitch's Theorem, the theorem shows that the schema

(V′) $P \to \Diamond OP$

has false instances. A generalization of Tennant's strategy would presumably proceed by pointing out that if 'P' is restricted in (V′) to range over sentences for which OP is consistent, then Fitch's theorem does not show this; and, again, the strategy will succeed to the extent that it is plausible to suppose that in saying that "every truth could be O-ed" one means that every truth for which OP is consistent could be O-ed. The trouble is, it is not at all plausible in general to suppose this.

Consider just one example, with O taken to be Williamson's operator 'it is T-conceived that...', where a proposition is T-conceived iff it is both true and conceived of by someone. Since some true propositions have never been conceived of by anyone, Fitch's Theorem apparently shows that there are truths that could not be T-conceived. A proponent of Tennant's strategy would presumably point out that although Fitch's Theorem does show this, it does not show that there is a true but non-T-conceivable sentence p for which "it has been T-conceived that p" is consistent. But why is this relevant? The application of Fitch's Theorem in this case is surprising

[14] See, for example, Hand & Kvanvig (1999) and DeVidi & Kenyon (2003), as well as Tennant (2001).

because it seems obvious—indeed, almost definitional—that *any* true proposition could be conceived of by someone without thereby becoming false. Perhaps careful reflection will show that this apparently obvious claim is false. But surely no amount of reflection should convince us that in thinking that any truth is T-conceivable we really meant to be thinking that any truth whose T-conception is consistent is T-conceivable.

The general lesson of the discussion in this section is the following. The so-called "paradox of knowability", the fact that there is an extremely simple and apparently sound argument that seems to refute an apparently substantive claim about what can be known, is just one instance of a more general phenomenon that needs to be explained. The more general phenomenon is that the same simple argument, appropriately reinterpreted, also seems to refute a whole range of apparently substantive claims: about what can be rationally believed, what can be T-conceived, what can be discovered *a posteriori*, what can be taken for granted in a conversation, and more. Any solution to the original paradox, therefore, is incomplete unless it can be extended to cover these other cases. While these observations do not entirely rule out rejectionist or restrictionist solutions to the paradox, they do make them appear unattractive, and they motivate exploring a different strategy. Since we are not here countenancing revisionism (according to which either the logical background or the factivity or distribution principles in Fitch's Theorem are mistaken), we are left with reinterpretivism: The claim that every truth could be known, or more generally that every truth could be O-ed, is not correctly represented by the schema

(V′) $P \rightarrow \Diamond OP$.

In the next section I present and discuss Dorothy Edgington's (1985) reinterpretivist strategy. Although I will argue that Edgington's solution cannot succeed, I will extract a lesson from it that will be put to work in my own reinterpretivist proposal. I will largely confine myself, for convenience, to the original version of Fitch's Theorem in which O is taken to be the operator K, 'it is known that…'; but I will return at the end to address the more general form of the paradox.

3. Knowability and actuality

The considerations given above motivate a reinterpretivist strategy for solving the paradox of knowability. But there is an additional, independent reason for thinking that the claim that any truth could be known cannot be correctly represented by the schema (V). For whether or not this claim is correct, it is surely certain that no *falsehood* could be known. No-one could know a falsehood, since knowing a falsehood would violate the incontestable principle that anything known is true. Accordingly, whatever we say about the verificationist claim that any truth could be known, we must at least accept its converse, that anything that could be known is true.[15] Yet the converse of schema (V) is manifestly wrong:

(CV) $\Diamond KP \rightarrow P$.

The mere fact that it is *possible* that it is known that p does not show that it is in fact true that p. It is, for example, possible that I know that I have fifteen toes (since it is possible for me to *have* fifteen toes without there being any impairment of the usual process leading from counting one's toes to knowledge of how many toes one has); it does not follow from this that I in fact have fifteen toes. In short, 'it is knowable that…' or 'it could be known that…' is a factive operator; '$\Diamond K$' is not;[16] and so the claim that any truth could be known is not to be represented by schema (V). A reinterpretation is required.

Dorothy Edgington (1985) finds further motivation for reinterpretation by comparing the verificationist claim that any truth could be known to a non-modal analog, the false claim that each truth is *in fact* known at some time. By analogy with the schema (V), we might try representing this non-modal claim as the schema

(F) $P \rightarrow SK_TP$,

where 'S' is the temporal operator 'at some time…', and 'K_T' is the tensed operator 'it *is* known by someone that…'. Fitch's Theorem, in this context, shows that schema (F) is incompatible with the fact that there are some truths which are not known *now*. But we should not conclude from

[15] This point is mentioned by Edgington (1985), Rabinowicz & Segerberg (1994), and Williamson (2000); it is discussed in more detail by Tennant (2000), Rosenkranz (2004) and Brogaard & Salerno (2006), among others.

[16] If '$\Diamond K$' were factive then we would have $\Box(\Diamond KP \rightarrow P)$, which together with (V) would imply that there are no contingent truths.

this that if any truth is eventually known then every truth is already known. Instead, Edgington points out, we should conclude that schema (F) was an inadequate representation of the claim that each truth is in fact known at some time.

Moreover, reflection on the interaction of the temporal operator S and the tense in K_T shows why schema (F) gets things wrong. As a simple illustration, suppose someone gives me a flower that I know grows an inch each day. Today it is in a box, so I don't know that it is six inches tall. But tomorrow I will unwrap it, see that it is seven inches tall, and thereby come to know that today it was six inches tall. If p is the (tensed) sentence "the flower is six inches tall," then p now expresses a truth that will in fact be known by me tomorrow, and *a fortiori* will be known by someone at some time. Yet if we put p into schema (F), we get the claim that if the flower is six inches tall then it is at some time known that the flower is six inches tall. And this claim is false, for there is *no* time at which the sentence "it is known that the flower is six inches tall" expresses a truth.

The mechanism of failure of schema (F) as a representation of the claim that any truth will in fact be known can be diagnosed if, following Edgington, we replace the tense operators in (F) by explicit quantifiers over moments of time. So understood, we can think of schema (F) as expressing an object-language version of the metalinguistic

(QF) $(P$ is true at $t) \rightarrow (\exists t')($It is known at t' that P is true at $t')$.

Yet the claim we are interested in, the claim that any truth will in fact be known, intuitively amounts to the quite different

(QF′) $(P$ is true at $t) \rightarrow (\exists t')($It is known at t' that P is true $\underline{\text{at } t})$.

We can get an object-language version of (QF′) if we help ourselves to a temporally rigid indexical operator N, 'it is now the case that…'. Using that operator, we can represent the claim that any truth will in fact be known as the schema

(F2) $NP \rightarrow SK_TNP$.

Intuitively, this schema is the correct representation of our target claim; at the very least, it overcomes the problems raised for schema (F).[17] Moreover, (F2) is compatible with the fact that some truths are not known now—Fitch's Theorem simply does not apply.

Edgington's thought is that the very same considerations are at work with respect to the original, modal verificationist claim that any truth *could* be known. Recall the initial representation of that principle

(V) $P \to \Diamond KP$.

If we replace the operator $\Diamond$ by a quantifier over possible worlds then (V) is effectively an object-language version of

(QV) (P is true at w)→($\exists w'$)(It is known at w' that P is true at w').

Edgington suggests that, just like in the temporal case, this is not what we want. Instead, the claim that any truth could be known intuitively says something like

(QV′) (P is true at w)→($\exists w'$)(It is known at w' that P is true $\underline{\text{at } w}$).

Just as the strong temporal verificationist principle says that for everything that is true *now* there is a time at which it is known to be true *now*, so the modal verificationist principle says, according to Edgington, that for everything that is *in fact* true it is possible that it is known to be *in fact* true.

As before, we can get an object language version of (QV′) by using a modally rigid indexical operator A, 'it is actually the case that…'. With this operator, we can represent the verificationist principle as

(V2) $AP \to \Diamond KAP$.

And as before, Fitch's Theorem does not apply.[18] Edgington's schema (V2), it seems, has several advantages over the original schema (V) as a representation of the claim that any truth could be

[17] For an argument against the introduction of N in this context, as well as a discussion of the disanalogies between the temporal and modal cases, see Burgess (forthcoming).

[18] If in (V2) we substitute our problematic sentence '$p \wedge \neg Kp$', and apply modus ponens, we get '$\Diamond KA(p \wedge \neg Kp)$', which is consistent. Applying the two principles about knowledge, and distributing the actuality operator over the conjunction, instead of deriving '$\Diamond(Kp \wedge \neg Kp)$' we can only derive '$\Diamond(KAp \wedge A\neg Kp)$'.

known: It seems a better candidate for capturing the meaning of that claim; it respects the fact that 'it is knowable that. . . ' is a factive operator; and, in blocking Fitch's Theorem, it constitutes a solution to the paradox of knowability.

Nevertheless, these advantages of (V2) come at a significant cost, of which Edgington is well aware. If we follow the usual practice of treating the modal operators $\Diamond$ and A in (V2) as quantifiers over possible worlds, then it is not clear that (V2) can ever be satisfied in the cases that interest us, namely those cases involving a proposition that is actually true but unknown. For in such cases the truth of (V2) will require that there are non-actual possible worlds in which people know things about the goings-on of this, the actual world. And this kind of counterfactual knowledge of actuality seems starkly impossible.

Consider, for example, the proposition that there are exactly *n* french fries on my plate. For the appropriate *n* this proposition is actually true, so the relevant instance of schema (V2) holds that it is possible that someone knows that actually there are exactly *n* french fries on my plate. In terms of possible worlds, that is to say that there is a possible world *w* in which it is known that in the actual world there are exactly *n* french fries on my plate. Since, in the actual world, no-one knows how many french fries are on my plate, this possible world *w* must be distinct from the actual world. So if schema (V2) is correct, then there is a non-actual possible world in which it is known that in the actual world there are exactly *n* french fries on my plate. But how could this be so?

Notice, first, that it is not enough for someone at *w* to know that the *sentence* "Actually there are *n* french fries on Fara's plate" expresses a truth. For given the indexical character of the actuality operator, that sentence expresses at *w* a claim about french fries at *w*, which is not what concerns us. For someone to know at *w* that in the *actual* world there are exactly *n* french fries on my plate, it seems that there are only two alternatives. Either someone at *w* somehow has *acquaintance* with the actual world, and thereby knows about french fries *there*; or someone possesses a uniquely identifying description D of the actual world, and knows that there are exactly *n* french fries on my plate at the world satisfying D. But neither of these alternatives is plausible. Perhaps we can in fact have acquaintance with the actual world; but whatever our conception of possible worlds may be, surely an agent could not *counterfactually* have *de re* knowledge of the actual world, still less acquaintance with it. But nor does it seem possible for an agent to counterfactually possess a uniquely identifying description of the actual world. A uniquely identifying description would have to specify, for each proposition, whether or not it is true; and it is hard to see how an agent could be in possession of such a description.[19] Edgington herself agrees, when she writes

> No knowledge is specific enough to be knowledge of the actual world, as opposed to countless others which coincide with the actual world in relevant respects.
> (Edgington 1985, p. 566)

Edgington's solution to this problem is to treat the modal operators in (V2) as quantifiers not over possible worlds but over what she calls "possible situations", where a possible situation is a way things might be that, unlike a possible world, is less than maximally specific. As an example, Edgington has us consider the possibility that a die land six-up when thrown. There are indefinitely many possible worlds in which the die is thrown and lands six-up; yet, according to Edgington, there is only one possible *situation* in which it lands six-up. Interpreted in these terms, schema (V2) requires that if a proposition p is actually true, then there is a possible situation in which someone knows that p is true in the actual situation. And for someone in another possible situation to know this, Edgington argues, she need not be acquainted with, or possess a uniquely identifying description of, the actual *world*; she need only be able to uniquely identify the actual *situation*, a task which is manageable because the actual situation is highly non-specific. It seems that one can know that p is true in the actual situation just by knowing that p is true in the possible situation in which α, where α is some description that is detailed enough to pick out the actual situation, yet is sufficiently simple for someone to grasp.

It is not at all clear, however, that the move from possible worlds to possible situations helps in understanding schema (V2). For what does Edgington mean by *the* actual situation? There are

[19] For a similar point in a quite different context, see Soames (1998).

six possible situations regarding how a die lands, one of which is an actual situation; there are twenty-six possible situations regarding what the first letter of *Naming and Necessity* is, one of which is an actual situation; and so on. These are different actual situations, since they differ in the specificity of certain facts. Which of these indefinitely many actual situations is *the* actual situation of which someone in another possible situation knows that, in it, *p*? In other words, which of these actual situations is the one *uniquely* characterized by the descriptive sentence α?

However this question is answered, notice that α must have *p* as a logical consequence. For suppose not. Then how *could* our agent know that *p* is true in the situation in which α? By hypothesis, α fully characterizes some possible situation, yet it does not specify whether or not *p* is true in that situation. And given that α *uniquely* characterizes a situation, it must be that *p* is neither true nor false there. For otherwise there would be two situations, alike except for the truth value of *p*, each of which satisfies the characterization given by α, violating the assumption of uniqueness. But if α specifies a situation in which *p* is neither true nor false, then how could it be known that in the possible situation in which α, *p* is *true*? It couldn't; and so we must conclude that *p* is a logical consequence of α.

But if *p* is a logical consequence of α, then the relevant instance of Edgington's schema (V2) is vacuous:

$$Ap \rightarrow \Diamond KAp.$$

For we are being asked to understand this as meaning something like:

> If *p* is actually true, then it's possible that it's known that *p* is true in the possible situation in which α.

But if *p* is a logical consequence of α, then *of course* it's possible that it's known that *p* is true in the situation in which α. This kind of knowledge requires no more than trivial knowledge of a logical truth. Yet the claim that (13) is supposed to represent, the verificationist claim that if *p* is true then the truth of *p* is something that could be known, is surely not vacuous. So we are no further forward. Even if we interpret $\Diamond$ using a quantifier over possible situations instead of worlds, we still don't know how to understand ascriptions of counterfactual knowledge of actuality.[20]

We seem to have reached an impasse. First, consideration of the generalized form of Fitch's Theorem suggests that the most promising strategy for solving the paradox of knowability will be one of reinterpretation: The claim that any truth could be known should not be represented by the schema

(V) $P \to \Diamond KP.$

Second, this suggestion is bolstered by noticing that 'it could be known that…' is a factive operator—whatever could be known is true. The factivity of this operator can be respected if we insist that in saying that it could be known that *p* we really mean to be saying that it could be known that *actually p*. This suggestion is made independently plausible by noticing that its analog in the temporal case seems true: In saying that it will be known that *p*, we *do* really mean to be saying that it will be known that it is *now* the case that *p*. Yet unlike future knowledge of the present, counterfactual knowledge of actuality seems unattainable. Edgington's solution to the paradox of knowability, although tempting, is apparently hopeless.

4. Different senses of 'could'

In our original presentation of the paradox of knowability we understood 'it could be known that *p*' as the claim that it is (metaphysically) possible that it is known that *p*. Edgington argued that this understanding ignores an implicit reference to actuality, and that "it could be known that *p*" should be understood as the claim that it is (metaphysically) possible that it is known that *actually p*. These two proposals share the assumption that the 'could' in 'it could be known…' should be understood in terms of (metaphysical) possibility, as the operator 'it is possible that…'. I claim that the lesson to be drawn from our discussion is that this shared assumption is false.

There is considerable precedent for thinking that the modal auxiliary 'could' should not uniformly be understood as the modal operator 'it is possible that…'. One obvious example is the

[20]The argument given here closely follows one given by Williamson (1987). See also Williamson (2000, chptr. 12).

use of ‘could’ in its epistemic sense, to express ignorance on the part of an agent. If I am asked to quickly solve an algebra problem, I might say “The answer could be 17, and it could be 102” to express the fact that 17 and 102 are among the answers that I do not know to be incorrect. In saying this I certainly do not express the proposition that it is *possible* that the answer is 17 and that it is *possible* that it is 102. For I know that at most one of these answers is correct, and that which one is correct is not a contingent matter; and so I know that if it is indeed possible that the answer is 17 then it is *not* possible that the answer is 102.[21]

Another, more subtle example of a use of ‘could’ to express something other than simple possibility is found in sentences like “Michael Jordan could make three shots in a row from the free-throw line”. This sentence attributes an ability or skill to Michael Jordan, an ability or skill that most of us lack. In the very same sense of ‘could’, I could *not* make three shots in a row from the free-throw line—I would be very lucky to make even one. Yet it is of course *possible* that I make the three shots, even in a restricted sense of ‘possible’. When we say that Michael Jordan could but I could not perform this feat, we are not distinguishing the two of us by means of possibilities, ways things might be. Instead, we are distinguishing us on the basis of ways things *are*: Michael Jordan has an ability that I lack. This “ability-sense” of ‘could’ is not plausibly interpreted as the modal operator ‘it is possible that. . . ’.[22]

The sense of ‘could’ which I wish to focus on here, and which I will argue is at work in the paradox of knowability, is the one used to attribute a *capacity* to a subject. I take capacities to be very much like abilities—indeed, the two terms are often used interchangeably—although there are some important differences between them. First, although this will not be important in what follows, inanimate objects typically possess capacities but not abilities. Abilities, as I understand

[21]I do not mean to suggest that the apparatus of possible worlds should not be used to interpret ‘could’ in its epistemic uses. I mean only to suggest that epistemic uses of ‘could’ are among the cases where “it could be that *p*” and “it is possible that *p*” are not equivalent.

[22]In a quite different context, Peter van Inwagen (1983) makes some remarks that amount to a proof of this fact, although he certainly did not intend them as such. Van Inwagen points out that his principle β is a *theorem* provided his operator ‘N’ obeys a normal modal logic. If we read ‘Np’ as ‘p, and no-one has or ever had the ability to make it the case that $\neg p$’, then on the assumption (which van Inwagen would deny) that there is a compatibilist sense of ability—that is, a sense of ability according to which determinism is compatible with having the ability to do other than one in fact does, and therefore one according to which principle β is false—this result shows that “*S* has the ability to *F*” does not express the claim that it is possible that *S F*s, in *any* sense of ‘possible’, however restricted.

them, are always abilities to perform intentional actions, whereas something can have the capacity to behave in any way at all. A sponge has the capacity to soak up water, for example, but not the ability to do so.

More importantly for our purposes, abilities differ from capacities in that if one currently has the ability to do something, and if one has the opportunity to exercise that ability in the sense of being in circumstances conducive to the exercising of the ability, then if one tries to exercise the ability one will most likely succeed. For example, since Michael Jordan has the ability to make three shots in a row from the free-throw line, he will most likely succeed in making them if he is given a ball and sufficient incentive to try his best. This is not so for one's capacities. I have the capacity, but not the ability, to make three shots in a row from the free-throw line. I have the capacity because I have sufficient strength, coordination and control over my movements for there to be some chance of my making the three shots—in this regard I differ from rocks, mice, computers and other things. But even if I tried my best, I would most likely fail, and so I lack the *ability* to make the shots.

This (somewhat stipulative) distinction between abilities and capacities can be further illustrated by a memorable example of David Lewis's:

> An ape can't speak a human language—say, Finnish—but I can. Facts about the anatomy and operation of the ape's larynx and nervous system are not compossible with his speaking Finnish. . . But don't take me along to Helsinki as your interpreter: I can't speak Finnish. (Lewis 1976)

Lewis rightly concluded from this observation that the auxilliary 'can't' in the first and last sentences is being used equivocally, to mean first one thing and then another. One diagnosis of the equivocation is as follows. When, at the beginning, we say that an ape can't speak Finnish, we mean that an ape lacks the capacity to speak Finnish. When, at the end, we say that I can't speak Finnish, we mean that I lack the *ability* to speak Finnish. But I do have the *capacity* to speak Finnish (and so in that sense I can speak Finnish), since the anatomy and operation of my larynx and nervous system do not rule out my speaking Finnish. More generally, one has the capacity to perform some feat provided one's internal constitution does not rule out the performance of that feat; to have the ability to perform that feat one must, in addition, be disposed to succeed in performing the feat if one tries.[23] So having the ability to do something requires having the capacity to do it, but not conversely.

I have already argued that when 'could' is used to ascribe an ability to someone, that use of 'could' should not be understood as the modal operator 'it is possible that…'. The same holds when 'could' is used to ascribe a capacity to someone or something. Clearly, to say that something has the capacity to do so-and-so is not to say that it is *metaphysically* possible that it does so-and-so. It is metaphysically possible for an ant to destroy an island—it is metaphysically possible for an ant to grow explosive tentacles that are detonated while the ant is walking on the island—but no ant has the capacity to do that. But might there be some restricted notion of possibility, less easily obtained than metaphysical possibility, according to which something has the capacity to do so-and-so if and only if it is possible in that sense for it to do so-and-so? I will argue that the answer to this is "no", and for a reason that is initially surprising: Sometimes, one might have the capacity to do so-and-so even though it is not possible that one does so-and-so in *any* sense of 'possible' obtained by restriction of metaphysical possibility. Let us work up to that.

Most of our capacities are, as a matter of fact, unexercised. Right now, for example, I am refraining from exercising all manner of capacities: my capacity to sing, my capacity to dance, my capacity to drink gasoline, and so on. Addtionally, many of our capacities will in fact *never* be exercised. I expect this holds of my capacity to drink gasoline and my capacity to speak Finnish, for example.

More interestingly, some of our capacities are not just unexercised, they are *unexercisable*; or at least they could turn out to be unexercisable if certain states of affairs obtained. For example, consider my capacity to swim across mile-wide rivers. We can imagine a frightening scenario in which global warming proceeds to the point that there are no mile-wide rivers left (and water is a scarce enough commodity that no artificial mile-wide rivers are created). If this scenario were to

[23]I do not mean this as an *analysis* of either capacities or abilities, but merely as an indication of the distinction between the two that I have in mind.

obtain, I would still have the capacity to swim across mile-wide rivers. But due to the circumstances I found myself in, this would be an unexercisable capacity. In one very restricted sense of ‘possible’, where it means something like “compatible with current topography”, it would not be possible for me to swim across mile-wide rivers; yet I would still have the capacity to do just that.

We have imagined a case in which facts about the earth have led to an absence of mile-wide rivers. Stretching credibility, we can further imagine that mile-wide rivers turn out to be physically impossible. We can imagine that the chemistry of water and the laws of fluid dynamics evolve to a state where they together preclude the flowing of any body of water wider than a hundred meters. Assuming that a body of water must be flowing to count as a river, this would be a case in which mile-wide rivers were physically impossible. Yet I would still have the capacity to swim across mile-wide rivers, as long as nothing changed about my physical make-up, motor control, coordination, and the rest. My capacity to swim across mile-wide rivers depends on the way I am; it does not depend on the availability of mile-wide rivers to swim across.[24]

We have described a case in which I have the capacity to do something even though it is not physically possible that I do that thing. Another example of David Lewis’s should persuade us that there are (or in principle could be) cases in which someone has the capacity to do something even though it is not even *metaphysically* possible that they do that thing. Lewis (1976) describes a version of the “grandfather paradox” of time travel, an argument that purports to show that at least some kinds of travel backwards in time are inconsistent. (Lewis’s aim is to refute the argument.) The argument begins with the supposition that travel backwards in time is possible, and concludes that on that supposition it both is and is not possible for someone to travel back in time and kill their own grandfather. We are to imagine that a certain character (Tim) has travelled backwards to a time before his own birth, with the intention of killing his grandfather

[24]This is not to suggest that capacities are intrinsic properties, depending for their possession *only* on the natures of the people or things that possess them. My capacity to swim across mile-wide rivers depends, for example, on the fact that the density of water is close to my own density. The point being made in the text is just that we *sometimes* retain the capacity to do so-and-so even in circumstances where doing so-and-so is not only not an option, but is nomologically precluded from being an option.

(Grandfather). Lewis writes:

> Tim can kill Grandfather. He has what it takes. Conditions are perfect in every way: the best rifle money could buy, Grandfather an easy target only twenty yards away, not a breeze, door securely locked against intruders. . . What's to stop him? The forces of logic will not stay his hand!. . . Tim is as much able to kill Grandfather as anyone ever is to kill anyone.

That is to say, Tim can kill Grandfather in the sense that he has the capacity to kill Grandfather (as I am understanding the notion of capacity). Yet, of course, in another sense of 'can', Tim cannot kill Grandfather. Indeed, it is metaphysically impossible that Tim kills Grandfather (even if we assume that travel backwards in time is itself metaphysically possible), since Grandfather's continued existence is a precondition for Tim's birth, and hence for Tim's doing anything at all, including trying to kill Grandfather. So we have a case in which someone has the capacity to do something, even though it is metaphysically impossible that they do that thing.

Perhaps the preceding example presupposed the metaphysical possibility of travel backwards in time, a contentious presupposition. So suppose that such travel is metaphysically impossible. Then it is also metaphysically impossible for me to interact with people who no longer exist (assuming that they could not travel *forwards* in time to meet me). Yet surely I have various capacities to interact with such people. For example, when he was alive I had the capacity to sit politely through my grandfather's stories. Nothing has changed in that regard: I *still* have the capacity to sit politely through my grandfather's stories. Yet it is metaphysically impossible that I sit politely through my grandfather's stories since it is metaphysically impossible (we are supposing) for him to tell me one of those stories.

It is important to note that although, on the conception of capacities I am appealing to here, we sometimes have the capacity to do certain things even while it is impossible that we do those things, it is the circumstances we find ourselves in, rather than the nature of those things themselves, that have this consequence. Speaking Finnish, swimming across rivers, killing a person, and listening to stories, are all of them types of events that it is routinely possible for a person to participate in, which is why it is unproblematic to suppose that most of us typically have the capacity to participate in these events. The point being made here is that we sometimes retain those capacities even if we find ourselves in circumstances that render the occurrence of those events—the *exercising* of the corresponding capacities—impossible. It is not being suggested that we ever have the capacity to do things which are *in themselves* impossible types of things to do—to square the circle, to change the atomic number of gold, or to know a conjunction, "p and it's not known that p".

Anything which is metaphysically impossible is thereby physically impossible, "topographically" impossible, and so on. Since, therefore, one can have a capacity to do something which it is metaphysically impossible to do, to say that someone has the capacity to do so-and-so is not to say that it is possible, in *any* restricted sense of 'possible', that they do so-and-so. So we have another case in which 'could' can be used to express something other than possibility, as, for example, in the claim that Tim could kill Grandfather.

It is this very sense of 'could'—the sense in which it expresses a capacity rather than a possibility—that I claim is at work in the paradox of knowability. Consider, for example, the (true) proposition that there were exactly *n* french fries on my plate this morning. No-one will ever know that this proposition is true, since no-one bothered to count my french fries. But this proposition *could* be known. That is to say (for the reasons given in section 3 above) that it could be known that *actually* there were *n* french fries on my plate this morning. This, I maintain, is a use of 'could' in its capacity-sense. To say that this proposition could be known is to say that someone had, has or will have the capacity to know that actually there were *n* french fries on my plate this morning. Indeed, *I* had the capacity to know this, since I had the capacity to count my french fries; and counting is, in cases such as these, a way of coming to know.

Notice that this is another case in which someone has the capacity to do something that is metaphysically impossible. As a matter of fact, I did not know that there were actually *n* french fries on my plate. So any possible world in which I know that there were actually *n* french fries on my plate is a non-actual world. Yet, for the reasons rehearsed in section 3 above, there is no non-actual world in which I know a proposition of the form 'actually *p*'. So there is no possible world in which I know that there were actually *n* french fries on my plate. Yet for all that, knowing this was something that I earlier had the capacity to do.

In general, the solution to the paradox of knowability being proposed is this. When the verificationist says that every truth could be known, she should be understood as meaning

(V3) $AP \rightarrow A\exists x(C_x K_x AP)$,

where 'C_x' abbreviates '*x* has the capacity to…' and 'K_x' abbreviates 'x knows that…'.[25] As an interpretation of the verificationist's claim, (V3) shares with Edgington's (V2) the advantage of blocking Fitch's argument: substituting a truth of the form 'p & $\forall y \neg K_y p$' in (V3) does not imply a contradiction.[26] But (V3) improves on Edgington's (V2) in that it does not require that inhabitants of non-actual possible worlds have knowledge of the goings-on of worlds that are foreign to them (*viz.* the actual world). Indeed, there is no quantification over possible worlds in (V3) at all. According to (V3), the knowability of a truth consists in the actual capacities of actual-world inhabitants, capacities which in some cases it is not possible for those inhabitants to exercise.

In addition to blocking Fitch's paradox, (V3) (like Edgington's (V2)) succeeds in preserving the factivity of knowability, since someone can have the capacity to know that actually p only if it is *true* that actually p. Although this is a benefit of (V3), it does leave it potentially exposed to the "pharaonic paradoxes" of Rosenkranz (2004) and Brogaard & Salerno (2006), which purport to show that the factivity of knowability is incompatible with the principle that one may infer Q from the conjunction of P with $\Box(P \rightarrow Q)$. A discussion of the details of these interesting paradoxes would take us too far afield; it is sufficient here to note that they are all blocked by the presence of the actuality operators in (V3).[27]

[25] Of course, "x has the capacity to x knows that actually P" is not grammatical, since "x has the capacity to" must be followed by a non-finite clause. I thus intend '$C_x K_x \alpha$' to be read as "x has the capacity to know that α", where the non-finite clause here expresses the proposition expressed by "x knows that α".

[26] The proposal is correspondingly open to an objection that Williamson (1987) and others have leveled against Edgington, namely that in providing a condition only on truths of the form '*actually* P' it has nothing to say about contingent truths. This objection loses its force, however, once we recognize that 'actually P', although metaphysically necessary, is a priori equivalent to P.

[27] Very briefly, one of the paradoxes works by noting that if a true conjunction p & $(Kp \rightarrow Kq)$ is possibly known

(V3) is being proposed as an interpretation of the verificationist claim that any truth could in principle be known. Analogous interpretations can be given of the other (more plausible) claims discussed earlier. For example, the claim that any truth could be T-conceived should be understood to mean that for any truth p there is someone with the capacity to T-conceive that actually p; the claim that every empirical truth could be discovered *a posteriori* should be understood to mean that for any empirical truth p there is someone with the capacity to discover *a posteriori* that actually p; and so on. Fitch's Theorem, recall, was the result that there is no sentential operator O satisfying all four of the principles

1. $P \rightarrow \Diamond OP$;
2. $\Box(OP \rightarrow P)$;
3. $\Box(O(P \,\&\, Q) \rightarrow (OP \,\&\, OQ))$; and
4. There are true sentences of the form '$(q \,\&\, \neg Oq)$'.

This theorem is an interesting—perhaps surprising—artifact of the behavior of the possibility operator $\Diamond$. The lesson of this paper is that Fitch's Theorem need not concern us once we interpret claims of the form "every truth could be Oed" not as (1) but as meaning that for every truth p there is someone with the capacity to O that actually p.[28]

(in a factive sense of "possibly known"), then it follows that q is true. But from the fact that someone has the capacity to know that *actually* $p \,\&\, (Kp \rightarrow Kq)$, it does *not* follow that q is true. Interested readers should work through the proofs given in Brogaard & Salerno (2006), subsituting '$A\exists x(C_x K_x AP)$' for '$\Diamond KP$'.

[28]Thanks to Joe Salerno for extremely helpful comments. For help with early ancestors of this paper, thanks are due to Delia Graff Fara, Gilbert Harman, David Lewis, Timothy Williamson, and audiences at Syracuse University and the University of Michigan.

References

Brogaard, Berit & Joe Salerno. 2004. Fitch's Paradox of Knowability. In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta.
URL: *http://plato.stanford.edu/archives/sum2004/entries/fitch-paradox/*

Brogaard, Berit & Joe Salerno. 2006. "Knowability and the Closure Principle." *American Philosophical Quarterly* 43.

Burgess, John. forthcoming. Can Truth Out? In *New Essays on Knowability*, ed. Joe Salerno. Oxford: Oxford University Press.

DeVidi, D. & T. Kenyon. 2003. "Analogues of Knowability." *Australasian Journal of Philosophy* 81:481–495.

Edgington, Dorothy. 1985. "The Paradox of Knowability." *Mind* 93:557–568.

Fara, Michael. 2007. "The Paradox of Believability." *Review of Contemporary Philosophy* 6:13–17.

Fitch, F. B. 1963. "A logical analysis of some value concepts." *Journal of Symbolic Logic* 28:135–142.

Hand, M. & J. Kvanvig. 1999. "Tennant on Knowability." *Australasian Journal of Philosophy* 77:422–428.

Hume, David. 1777. *An Enquiry Concerning Human Understanding*. Oxford University Press.

Kvanvig, Jonathan. 1995. "The Knowability Paradox and the Prospects for Anti-Realism." *Noûs* 29:481–499.

Lewis, David. 1976. "The Paradoxes of Time Travel." *American Philosophical Quarterly* 13:145–52.

Rabinowicz, Wlodek & Krister Segerberg. 1994. "Actual Truth, Possible Knowledge." *Topoi* 13:101–115.

Restall, Greg. 2004. "Not Every Truth Can Be Known (At Least, Not All At Once).".
URL: *http://consequently.org/writing/notevery/*

Rosenkranz, Sven. 2004. "Fitch back in action again?" *Analysis* 64:67–71.

Soames, Scott. 1998. "The Modal Argument: Wide Scope and Rigidified Descriptions." *Noûs* 32:1–22.

Stalnaker, Robert. 1970. "Pragmatics." *Synthese* 22:272–289.

Stalnaker, Robert. 1973. "Presuppositions." *Journal of Philosophical Logic* 2:447–457.

Stalnaker, Robert. 1974. Pragmatic Presuppositions. In *Semantics and Philosophy*, ed. M. Munitz & P. Unger. New York: New York University Press pp. 197–213.

Stalnaker, Robert. 1984. *Inquiry*. Cambridge, Mass.: MIT Press.

Stalnaker, Robert. 2002. "Common Ground." *Linguistics and Philosophy* 25:701–721.

Tennant, Neil. 1997. *The Taming of the True*. Oxford: Oxford University Press.

Tennant, Neil. 2000. "Anti-realist Aporias." *Mind* 109:825–854.

Tennant, Neil. 2001. "Is every truth knowable? Reply to Hand and Kvanvig." *Australasian Journal of Philosophy* 79:107–113.

van Inwagen, Peter. 1983. *An Essay on Free Will*. Oxford: Oxford University Press.

Williamson, Timothy. 1982. "Intuitionism Disproved?" *Analysis* 42:203–207.

Williamson, Timothy. 1987. "On the Paradox of Knowability." *Mind* 96:256–261.

Williamson, Timothy. 1993. "Verificationism and Non-Distributive Knowledge." *Australasian Journal of Philosophy* 71:78–86.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.