Not Every Truth Can Be Known (at least, not all at once)

Greg Restall*

Philosophy Department
The University of Melbourne
restall@unimelb.edu.au

VERSION 1
April 10, 2007

Abstract: According to the "knowability thesis," every truth is knowable. Fitch's paradox refutes the knowability thesis by showing that if we are not omniscient, then not only are some truths not known, but there are some truths that are not knowable. In this paper, I propose a weakening of the knowability thesis (which I call the "conjunctive knowability thesis") to the effect that for every truth p there is a collection of truths such that (i) each of them is knowable and (ii) their conjunction is equivalent to p. I show that the conjunctive knowability thesis avoids triviality arguments against it, and that it fares very differently depending on another thesis connecting knowledge and possibility. If there are two propositions, inconsistent with one another, but both knowable, then the conjunctive knowability thesis is trivially true. On the other hand, if knowability entails truth, the conjunctive knowability thesis is coherent, but only if the logic of possibility is weak.

§1 There are many things that we don't know to be true. Ignorance is a fact of life. However, it is tempting to think that of the things that are true but not known to be true, each of them *could* be known. If the significance of a proposition is to be explained in terms of its verification conditions, for example, then if it is true, there must *be* some verification conditions, and it is tempting to say that we could (at least *potentially*, *in theory*) have access to them. So, it is tempting to endorse the claim

$$\textit{Every truth is knowable.} \tag{$\Diamond$}$$

which has come to be known as the *knowability* thesis, and its formalisation

$$(\forall p)(p \supset \Diamond Kp) \tag{1}$$

*See http://consequently.org/writing/notevery/ for the latest version of the paper, to post comments and to read comments left by others. ¶ Thanks to Conrad Asmus, Allen Hazen, Lloyd Humberstone, Nick Smith and Timothy Williamson, to audiences at Monash University, the University of Melbourne, and Oxford University, and to commentators at http://consequently.org/writing/notevery/ for helpful discussions. Feedback from anonymous reviewers for this volume was useful in cleaning up the presentation.

for any truth, it is *possible* ($\Diamond$) that it be *known* (K). This position is tempting to many, but Fitch has shown that the temptation comes at a very high cost. Using only inference principles that are very tough to reject, we can show that, given the knowability thesis, every truth is, indeed, known.[1]

1·2 Here is Fitch's proof that the knowability thesis fails if we are not omniscient. Suppose, for a *reductio*, that we are ignorant of some truth, so suppose that p is true but not known to be true. Then $p \wedge \neg Kp$ is true. So, by the knowability thesis, *this* is possibly known: $\Diamond K(p \wedge \neg Kp)$. Now, this is very hard to take. How could we know that $p \wedge \neg Kp$? If knowing a conjunction entails knowing the conjuncts, then $K(p \wedge \neg Kp)$ entails Kp and $K\neg Kp$.[2] Now knowledge entails truth, so $K\neg Kp$ entails $\neg Kp$, a contradiction. So, by a *reductio*, it is *not* possible that $K(p \wedge \neg Kp)$, and we have (using the knowability thesis), refuted the hypothesis of ignorance. If the knowability thesis holds, a much stronger thesis holds too: every truth is not merely knowable, but *known*.

1·3 This phenomenon has come to be called *Fitch's Paradox*, after F. B. Fitch, who first formulated it [3]. This paradox has generated a vast literature, including on the one hand "search and rescue" missions designed to find the true principle underlying knowability thesis, and to save them from similar paradoxical fate, and "seek and destroy" campaigns aimed at hammering more nails into the coffin of any so-called principle of knowability that shows signs of life.[3] This paper contains elements of *both* kinds of discussion. I shall present and motivate yet another revision of the knowability thesis, and then show that this revision is consistent and not subject to Fitch-paradoxical refutation — that is the search and rescue part of the story — and then I will show that this revision is not only consistent but either it is also *almost trivially true* and therefore, it is not likely to do the work that a verificationist or anti-realist might require of a principle of knowability, or it's an interesting, controversial thesis about knowledge, which is coherent under certain conditions.

1·4 The knowability thesis, cast as the statement (1), is dead. Fitch's paradox is a conclusive refutation, and even though many interesting moves are possible with the

[1] I indicated before that Kp should be read as "it is known that p." But known by whom? Kp can be read either as "α knows that p" for some fixed agent α, or "*someone* knows that p" without too much strain in what follows. The existential reading, which requires that p merely be known by *someone*, ensures that K is nothing like a normal modal operator. We do not have $Kp, Kq \vdash K(p \wedge q)$, since it may well be that someone knows that p and someone (else) knows that q without *anyone* knowing that $p \wedge q$. In what follows, however, we soon move from reading Kp as the straightforward "p is known" to the idealisation "p is a logical consequence of what is known," and this does satisfy the principle that distributes knowledge over conjunctions: if p and q are consequences of what is known (by someone or other) then so is their conjunction. For any of these readings, the knowability thesis has some bite. It seems like a substantial claim that any truth is knowable by someone or other. It seems like a more substantial claim that any truth is knowable by *you*.

[2] Maybe we could know a conjunction without knowing the conjuncts. No problem: Just interpret Kp as "p is a logical consequence of what I know." If the knowability thesis works for *knowledge*, it works for this K too. So, from now on, K will allow for deductive closure: if Kp and $p \vdash q$ then Kq.

[3] Brogaard and Salerno's "Fitch's Paradox of Knowability" [2] is a fine guide to this literature.

logic in which these principles are couched, defeating the inference from (1) to omniscience [1, 5], these answers do not address the question I take to be asked by Fitch's paradox. I say this because upon reflection, the principles motivating a knowability thesis in fact *undercut* its application in a case such as $p \wedge \neg Kp$. Consider any truth p, of which we are ignorant. *Given* the knowability thesis we can, indeed, imagine coming to know that p. This is all well and good, but any way we can go about knowing that p makes it no longer the case that $\neg Kp$. But, $\neg Kp$ was true, and so, maybe it too could be known. If we do not inquire as to the status of p (so we don't come to know that p is true) but rather take ourselves to consider whether or not Kp, it seems plausible to suppose that we could confirm that $\neg Kp$. In other words, it's quite coherent to suppose that there is nothing that we can see that makes it impossible for us to know that $\neg Kp$. But the conjunction $p \wedge \neg Kp$ is true, and none of the ways we have considered, of coming to know p, or coming to know $\neg Kp$ will provide a way to know *both* p and $\neg Kp$. The conjunction $p \wedge \neg Kp$ cannot be known "all at once." Fitch's proof, it seems, is not a *trick* to be avoided or to be explained away but a *result* to be understood.

§2 This reasoning points the way to a possible answer: the Fitch-paradoxical conjunction $p \wedge \neg Kp$ cannot be known "all at once" but it can be known "in pieces." In particular, the first conjunct can be known (or rather, there seems to be nothing preventing us knowing it), but it cannot be known if the second conjunct is known. Similarly, the second conjunct can be known, but it cannot be known if the first conjunct is known. They cannot be known together. (1) is refuted, but it begs for a reformulation. Instead of saying that any truth could be known, let's attempt to maintain instead that every truth can be known "in pieces." That is, for any truth p, there is some collection of truths, each of which *could* be known, and when taken together, entail the original truth p. In other words, p can be *factored* into components, each of which is knowable.

2·2 If we could defend the knowability thesis in this weaker form, according to which unknowable truths could be factored into knowable pieces, then we may be able to provide some comfort to the anti-realist who takes meaningfulness to be a matter of knowablility. For the fact that $p \wedge \neg Kp$ is unknowable is no counterexample to its meaningfulness any more than the unknowability of $p \wedge \neg p$ renders it meaningless. No, $p \wedge \neg p$ is meaningful when p is meaningful, because we can understand p and its negation and its conjunction, even if to understand this is to come to see that it can never be known for it can never be true. The same kind of process can be seen in $p \wedge \neg Kp$, though now we have a conjunction which we can see that we will never *know* even though it may be true. It is meaningful because it is a conjunction of meaningful claims.

2·3 In fact, one *could* say that in the original naïve formulation ($\Diamond$) didn't mean what is expressed by (1), at least in its application to the statement $p \wedge \neg Kp$. For the $p \wedge \neg Kp$ is not, in itself, *one* truth that is knowable, but *two*. There are two knowable truths here, not one. (This is altogether too tendentious a reading to take seriously, however. Nothing in this paper hangs on the idea that conjunctive knowability is what we *really wanted* in the first place.)

2·4 Now consider what it is for a sentence to be a conjunction of knowable sentences. (In what follows, I will call these 'knowables' for short). From the perspective of pure *logic* it matters not whether the original sentence is complex or atomic. For whatever may be expressed by a complex sentence may be expressed by an atomic sentence too. In whatever model theory we like, if we have a model in which a complex sentence is interpreted in some way, then as far as *logic* is concerned, any *simple* sentence may be interpreted in just that way too.[4] But suppose that our original sentence was a complex sentence like $p \wedge \neg Kp$. This sentence is unknowable. If this is because it expresses an unknowable sentence, then if we interpret the atomic sentence q as "meaning the same thing" as $p \wedge \neg Kp$, then it seems that q will (relative to this interpretation, of course) be unknowable as well. But q has no conjuncts at all: it is a simple sentence. Have we sunk the "factoring" analysis before it could set sail?

2·5 This factoring analysis may survive if we are prepared to agree that while the sentence q from our example has no *explicit* conjuncts, it may have conjuncts *implicitly*. The sentence q is *equivalent* (relative to this model, again) to the conjunction $p \wedge \neg Kp$. As far as *logic* is concerned this will suffice for a factorisation. We will say that q is *conjunctively knowable* (relative to a model) if it is equivalent (relative to that model) to a conjunction, each of whose conjuncts are knowable (relative to that model).

2·6 This is my proposed revision of the knowability principle:

$$\textit{Every truth is conjunctively knowable.} \tag{$\heartsuit$}$$

In the rest of this paper, we will examine the fate of this principle.[5]

§3 For the proposal to be formally evaluated, it must be stated more sharply. This thesis assumes a number of exotic elements of logical vocabulary, such as propositional equivalence, propositional quantification, epistemic and modal operators. To properly state this thesis will require a great deal of machinery. The *syntax* of the claim is straightforward enough. We may formalise one version of this claim as follows:

$$(\forall p)\big(p \supset (\exists q, r)\big((p = q \wedge r) \wedge \Diamond Kq \wedge \Diamond Kr\big)\big) \tag{2}$$

[4] This phenomenon underlies one of the substitutional properties of formal logics. If ϕ is a tautology containing the atomic sentence p, then ϕ', found by replacing p everywhere by another formula B is also a tautology.

[5] After writing a draft of this paper, Joe Salerno brought to my attention Risto Hilpinen's paper "On a Pragmatic Theory of Meaning and Knowledge," [4] in which he argues that a Peircean pragmatism motivates a conjunctive knowability principle just like this. I must leave it to the reader to determine whether or not the results of the investigation below are congenial to the pragmatist project.

This version posits a very *strong* version of conjunctive knowability: every proposition may be factored into a conjunction of *two* propositions, each of which is knowable. [6]

3·2 The work comes in when we are required to characterise the logical properties of propositional quantification, propositional identity and the modal and epistemic operators. Thankfully, for our purposes, we need not attempt to pin down the *right* principles governing $\Diamond$, K, $(\forall p)$ and $=$. Qua *logician* my job is to investigate the consistency of (♡) and its formalisation (2). Fitch showed that (1) is inconsistent with deeply plausible modal and epistemic principles. I will show that (2) does not suffer *that* fate. (2) is consistent, and compatible with the strong principles of modal and epistemic reasoning. To do this, we need not find the *right* principles of such reasoning. In doing this, it is acceptable to overshoot and require too much. I will provide a class of models that show that the revised knowability thesis (♡) and its formalisation (2) can be absolutely unrestrictedly true at no cost to ignorance or to many other epistemic or modal principles. (There will, however, be an important caveat to be discussed in Section 5.)

3·3 Our logic will be the incredibly *strong* modal epistemic logic in which $\Diamond$ and K are both governed by the principles of the logic s5. This is unrealistic in the extreme, for it commits us to wild epistemic principles such as the claim that if p is true then we must know that we don't know that $\neg p$ (if p then $K\neg K\neg p$) and even that if we don't know something we know that we don't know it (if $\neg Kp$ then $K\neg Kp$). Neither of these principles is particularly plausible (even if we take Kp to mean that p is a consequence of what we know) but we will use such a strong logic nevertheless, since nearly every epistemic or modal principle endorsed by someone or other is valid in this logic: $\mathrm{s5}_\Diamond \oplus \mathrm{s5}_K$.

3·4 This logic has models of the usual kind for modal logics. Here a model is a quadruple $\langle W, R_\Diamond, R_K, [\![\cdot]\!]\rangle$ where W is a non-empty set of worlds, $R_\Diamond$ and R_K are accessibility relations on W and $[\![\cdot]\!]$ is a function assigning to each atomic sentence (for example, p) an interpretation—a set of worlds (in this case, $[\![p]\!]$). In this modal logic we place no restrictions on which sets can be used to interpret sentences. All sets may be *propositions* in our model. The set $[\![p]\!]$ is the set of worlds in which p is true. In the usual way, the interpretation function is extended to assign sets

[6] It could be that something could be factored into *three* knowable conjuncts but not two. As far as I can see, there is no natural upper limit to the number of conjuncts one could require in a formalisation of (♡), so a perfectly *general* formulation would perforce be quite complex indeed, as it would have to quantify over *collections* of propositions (there are *some* propositions which factorise p). This seems to be the right condition:

$$(\forall p)\big(p \supset (\exists \mathfrak{P})((\textstyle\bigwedge \mathfrak{P} = p) \wedge (\forall q)(\mathfrak{P}q \supset \Diamond Kq))\big) \tag{2'}$$

where $\mathfrak{P}$ is the second order propositional variable used in the second-order propositional existential quantifier $\exists\mathfrak{P}$ ranging over classes of propositions (which is "really" a third-order quantifier over objects, since propositions are "really" zero-place properties) and the "connective" $\bigwedge$ sends a class of propositions to its conjunction.

to arbitrary sentences in the language. For conjunction, disjunction, the material conditional and negation we use the standard boolean operations. A conjunction $p \wedge q$ is true at the worlds where both p and q are true: $[\![p \wedge q]\!] = [\![p]\!] \cap [\![q]\!]$, and similarly for the other Boolean connectives.

The relations $R_\Diamond$ and R_K are used to model the operators $\Diamond$ and K respectively. In our case $R_\Diamond$ and R_K are both equivalence relations. $R_\Diamond$ is the equivalence relation governing $\Diamond$:

» $\Diamond\phi$ is true at w iff ϕ is true at a world in w's $R_\Diamond$ equivalence class.

We will call the worlds in w's $R_\Diamond$ equivalence class the *modal alternatives* of w.[7] Another way to understand the interpretation is as a function from propositions to propositions. $[\![\Diamond\phi]\!]$ is the union of all modal equivalence classes *overlapping* $[\![\phi]\!]$. Think of approximating the proposition $[\![\phi]\!]$ by equivalence classes, counting in our approximation any equivalence class that at least partly overlaps the original proposition: $[\![\phi]\!]$ is approximated by its *closure*. Similarly, an equivalence relation R_K governs K:

» $K\phi$ is true at w iff ϕ is true at all worlds in w's R_K equivalence class.

We will call the worlds in w's R_K equivalence class the *epistemic alternatives* of w. $[\![K\phi]\!]$ is the set containing all epistemic equivalence classes totally included in $[\![\phi]\!]$. Here, $[\![\phi]\!]$ is approximated by its epistemic *interior*. The modal alternatives of w need not be the same worlds as the epistemic alternatives: a modal alternative need not be an epistemic alternative (we can know things that are not necessary) and an epistemic alternative need not be a modal alternative (we can be ignorant of some necessary truths). Now for the propositional quantifier:

» $(\exists p)\phi$ is true at a world w if and only if for some set X of worlds, ϕ is true at w when we take the formula p occurring unbound in ϕ to be true at exactly the worlds in X.

» For identity, we will say that $\phi = \psi$ is true at a world just when $[\![\phi]\!] = [\![\psi]\!]$.

This suffices to ensure that we may, for example, infer from $\phi = \psi$ that $\theta(\phi) = \theta(\psi)$. Any formula containing ϕ (for example, $\Diamond K\phi$) is true in the same worlds as the formula found by replacing those ϕs by ψs (in this case, $\Diamond K\psi$).

An argument is $\text{s5}_\Diamond \oplus \text{s5}_K^{\exists p,=}$ *valid* if for every model, if the premises are true in a world in that model, the conclusion is true in that world too. A sentence is an $\text{s5}_\Diamond \oplus \text{s5}_K^{\exists p,=}$ *tautology* if and only if it is true in every world in every model.

3·5 Let me reiterate: This model theory is not to be endorsed as giving us the "true picture" of knowledge, possibility, propositional quantification and propositional identity. It is intended as a grab-bag sizeable enough to catch all principles thought

[7] This leads to a slight infelicity: w counts as a modal alternative of *itself*.

to govern epistemic modal logic with propositional quantifiers. If we can find models that both validate the conjunctive knowability principle (♡) and allow for ignorance, then they show that no principle *true* in these models collapses conjunctive knowability into omniscience.

3·6 Here is a simple model in which (♡) holds. There are four worlds $\{a_1, a_2, b_1, b_2\}$.

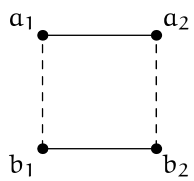

Figure 1: A simple $\text{s}5_\Diamond \oplus \text{s}5_K^{\exists p,=}$ frame

The modal accessibility relation $R_\Diamond$ relates a_1 to b_1 and a_2 to b_2; the epistemic accessibility relation R_K is orthogonal to the modal relation: it relates a_1 to a_2 and b_1 to b_2. So, a world's modal alternatives are those worlds sharing a *number*, and its epistemic alternatives are those sharing a *letter*. In Figure 1 (and in all other diagrams) solid lines join epistemic alternatives, and dashed lines join modal alternatives.

3·7 (♡) says that any proposition true at the world of evaluation is a conjunction of two propositions which, at the world of evaluation, are knowable. What *are* the propositions in our model that are knowable at any world? Any proposition true at both a_1 and a_2 is knowable at all worlds, since it is known at a_1 and a_2 (and hence, it is possibly known there), and at b_1, the world a_1 is possible, and at b_2, a_2 is possible, so at b_1 or at b_2, this proposition is also possibly known. So, if $\{a_1, a_2\} \subseteq [\![\phi]\!]$, then ϕ is knowable at any point in the model. Similarly, any proposition true at both b_1 and b_2 is knowable at every world. And these propositions are the *only* propositions knowable at any world. The propositions which can *not* be known are $\emptyset$, each singleton proposition $\{a_1\}$, $\{b_1\}$, etc., and the two *diagonal* propositions $\{a_1, b_2\}$ and $\{a_2, b_1\}$ and the modal alternative propositions $\{a_1, b_1\}$ and $\{a_2, b_2\}$. All other propositions are knowable, from the point of view of every world.

3·8 It will be helpful to consider why for any interpretation of p, the proposition denoted by $p \wedge \neg Kp$ is not knowable at any world. If $[\![p]\!] = X \subseteq W$, then $[\![Kp]\!]$ consists of the interior epistemic approximation of X, and $[\![\neg Kp]\!]$, then, is the union of all equivalence classes not totally inside X. So, its intersection *with* X (the set $[\![p \wedge \neg Kp]\!]$ consists of the union of all X-overlapping parts of epistemic equivalence classes that overlap X but do not fall completely inside X.[8] In the case where $[\![p]\!] = \{a_1, a_2, b_1\}$, $[\![Kp]\!] = \{a_1, a_2\}$ and so $[\![\neg Kp]\!] = \{b_1, b_2\}$, and

[8]In topological terms it is the part of the (epistemic) boundary of X that is also inside X.

$[\![p \wedge \neg Kp]\!] = \{b_1\}$. This proposition is not knowable, because it contains no epistemic equivalence classes as a subset.

3·9 In this model, every true proposition is a conjunction of two knowable propositions. The singleton $\{a_1\}$ is the conjunction of $\{a_1, a_2\}$ and $\{a_1, b_1, b_2\}$. The same goes for each other singleton. The pair $\{a_1, b_1\}$ is the conjunction of $\{a_1, a_2, b_1\}$ and $\{a_1, b_1, b_2\}$. The same goes for each other pair. It follows that (2) is true in our model.

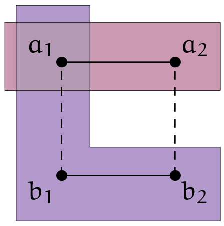

Figure 2: $\{a_1\}$ as a conjunction of knowables

3·10 So, we have shown that (2) survives consistently and coherently. In this model there is much ignorance (a proposition true at a_1 alone is true there but not known to be true), yet every proposition is a conjunction of two *knowable* propositions. The principle (♡) of conjunctive knowability is secure. Truth and knowability can be intimately connected, even if not every truth is knowable. What the friend of knowability cannot have *whole* she is allowed to have if she will accept it in two pieces.

3·11 At this point, the story takes a different turn. Conjunctive knowability is secure, but it either is almost certainly not what the verificationist wants, or it is too high a price for the verificationist to pay. In the rest of this paper I shall show that if knowability does not entail truth (so some falsehoods are knowable while not being known), then conjunctive knowability, in the form of (2) is not only true, but it's very hard to *refute* in an epistemic modal logic. It puts precious few constraints on knowledge or necessity, and so, is not useful as criterion for favouring one theory over another. If principles acceptable to the *realist* lead them to accept (2) while maintaining their realism, then (2) will do no good as a principle designed to *favour* the anti-realist. On the other hand, if knowability entails truth, then any non-trivial account of conjunctive knowability is inconsistent with plausible modal principles. (In particular, with the modal principle of transitivity: $\Diamond\Diamond p \vdash \Diamond p$.)

§4 So, consider what we have done so far. We have a model of $S5_\Diamond \oplus S5_K^{\exists p,=}$ in which conjunctive knowability is satisfied. It turns out that this is not a one-off affair. In an epistemic modal logic like $S5_\Diamond \oplus S5_K^{\exists p,=}$, and its much weaker cousins in which the modal and epistemic accessibility relations satisfy fewer constriants, (2) turns out to be *very* easy to validate. Not only are there many models in which

(2) is true, it turns out that (2) is a consequence of other, unproblematic modal and epistemic principles. In particular, (2) follows from the following thesis about possible knowledge, satisfied in the models we have seen:

$$(\exists q)(\Diamond Kq \wedge \Diamond K\neg q) \tag{3}$$

This is relatively uncontroversial, given one understanding of how possibility and knowledge (or the consequences of what is known) are connected. Provided that, for some q, both q and $\neg q$ are possibly true (and this is not too difficult to imagine) then it is not much more difficult to conclude that for some q, both q and $\neg q$ are possibly *known.* Of course, a circumstance in which one knows q is, perforce, one in which $\neg q$ is not known, and vice versa, for there to be a q such that $\Diamond Kq$ and $\Diamond K\neg q$, there must be at least *two* distinct modal alternatives, one at which q is known, and the other at which $\neg q$ is known. All that requires is that we have two modal alternatives whose epistemic closures do not intersect, like so.

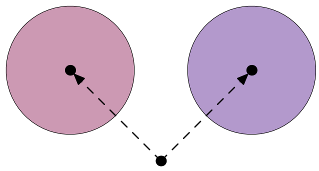

Figure 3: Two inconsistent knowables

Given a q such that both it and its negation are knowable, we can prove that for *any* true p, there are knowable p_1 and p_2 where $p = p_1 \wedge p_2$ and both p_1 and p_2 are knowable. If q is such that $(\Diamond Kq \wedge \Diamond K\neg q)$, then we may choose p_1 to be $p \vee q$ and p_2 to be $p \vee \neg q$. Then by simple boolean reasoning, $p_1 \wedge p_2 = (p \vee q) \wedge (p \vee \neg q) = p$. However, since $\Diamond Kq$, we have $\Diamond K(p \vee q)$.[9] Similarly, since $\Diamond K\neg q$, we have $\Diamond K(p \vee \neg q)$. Both $p \vee q$ and $p \vee \neg q$ are knowables, regardless of how unknowable p might be! (2) is a trivial consequence of the trivial truth $(\exists q)(\Diamond Kq \wedge \Diamond K\neg q)$. It looks as if (2) tells us little about the connection between truth and knowability.

4·2 Where can the fan of conjunctive knowability resist this analysis? It might be thought that a friend of *relevance* would quail at the identification of p with $(p \vee q) \wedge (p \vee \neg q)$, as well they should. The inference from p to $(p \vee q) \wedge (p \vee \neg q)$ is valid in almost every logic you care to mention, as it is found by composing the inferences from p to $p \vee q$, and from p to $p \vee \neg q$ and from these to their conjunction. All are simple lattice moves. The problem with relevance is in the other direction. To get from $(p \vee q) \wedge (p \vee \neg q)$ we need $q \wedge \neg q \vdash p$, and *this* is

[9] By distribution of both K and $\Diamond$ over logical consequence: since $q \vdash p \vee q$, then $Kq \vdash K(p \vee q)$ (remember, we read $K(p \vee q)$ as "$p \vee q$ is a consequence of what is known") and so, $\Diamond Kq \vdash \Diamond K(p \vee q)$: all are reasonable principles.

relevantly invalid. Nonetheless, rejecting this identity[10] is not going to stop this argument from getting off the ground. The crucial premise in the argument was that q and its negation were both knowable, and could be used in the factorisation of p. There is no requirement that q and its negation be used for this purpose. Provided that we are given two *incompatible* propositions (say, q_1 and q_2) that are knowable—so $q_1, q_2 \vdash \bot$ for the trivial proposition $\bot$, and $\Diamond Kq_1$ and $\Diamond Kq_2$—then even in relevant logics, the sentences p and $(p \land q_1) \lor (p \land q_2)$ are true in exactly the same situations. Blocking the inference from $(p \land q_1) \lor (p \land q_2)$ (with the rule $q_1, q_2 \vdash \bot$) to p requires blocking the distribution of conjunction over disjunction, not any odd behaviour about negation or relevance. So, pleading relevance or paraconsistency will not give the fan of conjunctive knowability (or its enemy, for that matter) a straightforward way out of the problem.[11]

4·3 Denying that one can infer $K(p \lor q)$ from Kp is not going to help, either, for as we have seen, we can replace talk of what is known by talk of what is a consequence of what is known (at least in decidable logics), and clearly, if p is a consequence of what is known, so is $p \lor q$. Furthermore, possibly being a logical consequence of what is known is not very far removed from being possibly known, so reading Kp throughout as "p is a consequence of what is known" does little violence to the principles in question, and it validates the inferences used in our deduction. So requiring high standards for knowledge, so high that logical consequence can lead you from what is known to what is not, is also not a way out of the problem.

4·4 So, conjunctive knowability is not only consistent, but it is *trivially* so, if possibility and knowledge are connected as given by (3), that is, if possible knowledge can sometimes outrun truth, just as Fitch's paradox has shown us that truth can sometimes outrun possible knowledge.

§5 Nonetheless, (3) is by no means uncontroversial.[12] What if we reject (3) and hold, instead, that only truths may be possibly known? So, let us embrace (4):

$$\Diamond Kp \vdash p \tag{4}$$

5·2 Before proceeding, I wish to do away with a *bad* argument for (4). No-one should argue as follows: "What is possibly known must be true, because of necessity, what is known is true. It is, therefore, impossible for what is known to be false. It follows that if something is possibly known, it is true."

---

[10] Which, it must be said, is *not* the same identity as that between p and $(p \land q) \lor (p \land \neg q)$, a factorisation seen again and again in different kinds of reasoning.

[11] A not-quite-straightforward way out of the problem is to deny that there are any incompatible pairs of propositions. To be sure, in many relevant logics, there is no way to construct formulas ϕ and ψ such that $\phi, \psi \vdash \bot$. Nonetheless, in most *models* for such logics there are ways to interpret ϕ and ψ such that their conjunction is absolutely inconsistent. Merely take $[\![\phi]\!]$ and $[\![\psi]\!]$ to have empty intersection, so their conjunction is true nowhere.

[12] Thanks to Nick Smith for pressing me on this point.

This contains a modal fallacy. We have attempted to infer from the innocuous "it is impossible for what is known to be false" ($\neg\Diamond(Kp \wedge \neg p)$) to the much stronger "if something is possibly known, it is true" (which as a material implication is $\neg\Diamond Kp \vee p$). In the first, the truth of p (or its not being false) is under the scope of the possibility operator, and in the second it is not.

That is a bad argument for (4). If you contemplate (4), do not do so for that reason.

5·3 In the case of an epistemic modal logic modelled with an accessibility relation $R_\Diamond$ for possibility and R_K for knowledge, (4) is straightforward to guarantee: we need simply that

$$(\forall x)(\forall y)(xR_\Diamond y \supset yR_K x) \tag{4'}$$

for then, when we are at x and we have *some* modally accessible world (call it y) in which *every* epistemically accessible world has p true, p is true at x, since x is epistemically accessible from y. Conversely, if we have *some* x and y where $xR_\Diamond y$ but not $yR_K x$, then if p is true at everywhere *other* than x (or, if you like, everywhere epistemically accessible from y, if everywhere other than x seems like overkill) then at y, Kp is true, and hence at x, $\Diamond Kp$ is true. However, at x, p is false. So, if we are allowed to assign the extension of a proposition at whim in our models (and it is hard to see why not) then condition (4′) corresponds precisely to the validity of (4).

5·4 Similarly, we can say, precisely, what condition on R_K and $R_\Diamond$ corresponds to conjunctive knowability in its weakest possible form. First note that, in a given model, if p is conjuctively knowable when $[\![p]\!]$ is a singleton set (so p is true at one world only) then *every* proposition is conjunctively knowable. (If ϕ is true at x, and if p is true at x alone, then consider the propositions, each knowable, which jointly entail p. These jointly entail ϕ—relative to that model—too, which shows that ϕ is also conjunctively knowable.) So, what does it take for a proposition true at x alone to be conjunctively knowable? Well, we must find for any world y distinct from x, a proposition which is knowable but not true at y. If that is not found, the conjunction of *all* knowable propositions will not entail p, since it will also be true at y, where p is not true. So, we require the following condition

$$(\forall x)(\forall y)(x \neq y \supset \exists z(xR_\Diamond z \wedge \neg zR_K y)) \tag{5}$$

for if (5) does not hold, then any z modally accessible from x will include y as epistemically accessible, so no proposition false at y will be know*able* from x, as it will not be known at any modally accessible worlds.

5·5 It follows that normal epistemic modal models for conjunctive knowability satisfy (5). Alas if (4′) and (5) both hold, then if $R_\Diamond$ is transitive, it is trivial in the sense that $xR_\Diamond y$ only if $x = y$. Here is why: if (4′) holds, then $\neg zR_K y$ means that $\neg yR_\Diamond z$, which when substituted in (4′) gives

$$(\forall x)(\forall y)(x \neq y \supset \exists z(xR_\Diamond z \wedge \neg yR_\Diamond z))$$

but if x and y are non-identical and $yR_{\Diamond}x$, then whenever $xR_{\Diamond}z$ by transitivity $yR_{\Diamond}x$, which contradicts what we have assumed.

5·6 We have a *syntactic* proof of this modal collapse as well. We can show that (4), $\Diamond\Diamond p \vdash \Diamond p$ and conjunctive knowability (in the most general form (2′)), ensure that $\Diamond p \vdash p$, in the presence of propositional quantification.

5·7 Here is the proof: Suppose $\Diamond p$. So there is some possible circumstance in which p is true. Consider one. In this circumstance p is true, so there are propositions r_1, $r_2, \ldots$, which together entail p, and each of which are possibly known. So, we have $\Diamond Kr_1$, $\Diamond Kr_2$, etc., and $r_1, r_2, \ldots \vdash p$. Now consider the actual circumstance in which $\Diamond p$ is true. In *this* circumstance, each $\Diamond Kr_i$ is *possible*: that is, $\Diamond\Diamond Kr_i$. But possible possibility is (we assume) possibility, so we have $\Diamond Kr_i$ for each i. But by (4), $\Diamond Kr_i \vdash r_i$, so each r_i is true. But $r_1, r_2, \ldots \vdash p$, so p is true too. In other words, we have inferred p from $\Diamond p$.

5·8 So, we cannot have (4), (5), $\Diamond\Diamond p \vdash \Diamond p$ and the non-triviality of $\Diamond$. One, at least, must go. Which one is to go? I am tempted to do away with (4), but we have already seen what can be done without (4): it makes conjunctive knowability all too easy. Making $R_{\Diamond}$ trivial is unacceptable, for then the only possibilities will be *truths*, so if every proposition is a conjunction of knowables, it will be a conjunction of *knowns*, and hence, every truth will be a consequence of what is known, making all ignorance vanish. We avoid Fitch's paradox and its heirs by denying the premise that we are not omniscient. To do away with (5) is to give up the task of exploring the consequences of conjunctive knowability. The only remaining option here (given the machinery of normal epistemic modal logics and their possible worlds models) is to explore the rejection of the transitivity of $R_{\Diamond}$. As a result, we will examine what follows if we deny the inference from $\Diamond\Diamond p$ to $\Diamond p$.

5·9 Denying transitivity of $R_{\Diamond}$ is a severe price to pay to save conjunctive knowability. It turns out that it is enough. In the remaining paragraphs of this section I will show that we may maintain (4), making every knowable a truth, and (5), making every truth conjunctively knowable, without concluding that every truth is known. A model showing this is relatively simple. The worlds are the (positive and negative) integers $\mathbb{Z}$. We have $xR_{\Diamond}y$ iff $y = x$ or $y = x+1$. (Notice that this is not transitive, since $0R_{\Diamond}1$ and $1R_{\Diamond}2$, but we don't have $0R_{\Diamond}2$. Nonetheless, it is reflexive, so at the very least, $p \vdash \Diamond p$.) We have xR_Ky iff $y = x$ or $y = x-1$. (Notice that this is not transitive *either*, so we do not have $Kp \vdash KKp$, but it is reflexive, so $Kp \vdash p$, as one would hope.)

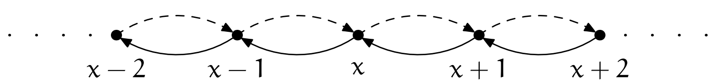

Figure 4: The Model

5·10 That is the model. Let's see how it manages to satisfy (4′) and (5). We have satisfied

($4'$) by fiat: if $xR_{\Diamond}y$, then $y = x$ or $y = x + 1$, in which case either $x = y$ or $x = y - 1$, ensuring that yR_Kx. So, ($4'$) is satisfied, ensuring that $\Diamond Kp \vdash p$.

Conjunctive knowability, in the form of (5), is satisfied too. If $x \neq y$, then there is always some z where $xR_{\Diamond}z$ but not zR_Ky. If we don't have xR_Ky, then choosing x for z will suffice (since $xR_{\Diamond}x$ always). If we do have xR_Ky, then since $x \neq y$, we have $y = x - 1$. Then choose $x + 1$ for z. We have $xR_{\Diamond}z$ (z is one step up from x) but we don't have yR_Kz (y is *two* steps down from z, which is just too far).

5 · 11 How does this model *work*? At every point, x, knowledge is a little limited because x is epistemically indistinguishable from $x - 1$. Only propositions true at both x and $x - 1$ may be known at x. Nonetheless, the world $x + 1$ is *modally* accessible from x, and at *this* world, $x - 1$ is not epistemically accessible but x is. This means that any proposition true at x is a conjunction of two knowable propositions. If p is true at the set X (including x) then consider two propositions q_1 and q_2, true at $X \cup \{x - 1\}$ and $X \cup \{x + 1\}$ respectively. q_1, true at $X \cup \{x - 1\}$, is known at x (and so is *possibly* known at x) and q_2, true at $X \cup \{x + 1\}$, is known at $x + 1$ (and so is also possibly known at x). In this case, as in our other models, every proposition true at a point is a conjunction of two knowable propositions. Nonetheless, not every proposition is known: at every point there is ignorance.

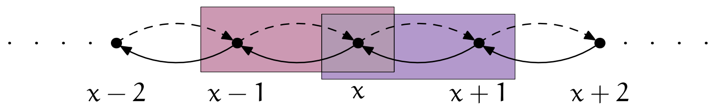

Figure 5: $\{x\}$ as a conjunction of two propositions knowable at x

5 · 12 So, if knowability entails truth then we can maintain the conjunctive knowability thesis in the form of (2), but only at the cost of rejecting the S4 principle for possibility: $\Diamond\Diamond p \nvdash \Diamond p$.

§6 Fitch's paradox shows us that not every proposition is knowable: at least all at once. Fitch's paradoxical sentence is an example of a proposition that cannot be known, but which can nonetheless be split into pieces, each conjunct of which can be known. It turns out that this modest fallback position is coherent. We may coherently hold that every proposition can be factored into a conjunction, each of which are knowable. Thinking of this in terms of possible worlds, it comes quite close to one original consideration in motivating of knowability. Propositions divide possible worlds into those that are *in* and those that are *out*. Conjunctive knowability tells us that for any world that a proposition takes to be *out*, we can know something that would rule out that world. Think of the discriminations that a proposition makes as constituted by all of the worlds inconsistent with it. According to conjunctive knowability, no proposition makes a discriminiation essentially beyond our grasp. This is coherent. If two inconsistent propositions are knowable, then conjunctive knowability is coherent but trivial. If knowability entails truth, then conjunctive knowability is both coherent and substantial.

§7 REFERENCES

[1] JC BEALL. "Fitch's Proof, Verificationism, and the Knower Paradox". *Australasian Journal of Philosophy*, 78(2):241–247, 2000.

[2] BERIT BROGAARD AND JOE SALERNO. "Fitch's Paradox of Knowability". In *Stanford Encyclopedia of Philosophy*. 2004.

[3] F. B. FITCH. "A Logical Analysis of Some Value Concepts". *Journal of Symbolic Logic*, 28(2):135–142, 1963.

[4] RISTO HILPINEN. "On a Pragmatic Theory of Meaning and Knowledge". *Cognitio: Revista de Filosofia*, 5(2):150–167, 2004.

[5] HEINRICH WANSING. "Diamonds are a Philosopher's Best Friends: The Knowability Paradox and Modal Epistemic Relevance Logic". *Journal of Philosophical Logic*, 31:591–612, 2002.