ACTUAL TRUTH, POSSIBLE KNOWLEDGE

Wlodek Rabinowicz and Krister Segerberg

Uppsala University

1. Introduction

F.B. Fitch is credited with a simple argument purporting to show that the verificationist claim:

(Ver) Truth implies Knowability,

leads to the unacceptable conclusion:

Truth implies Knowledge.

The argument, as it is usually presented, rests on the following formalization of (Ver):

(FVer) $\phi \rightarrow \Diamond K\phi$,

with $\Diamond$ and K standing for the possibility and knowledge operators, respectively.

Given that knowledge implies truth and distributes over conjunctions, (FVer) immediately implies that for no ϕ is it true that $\phi \wedge \neg K\phi$. Otherwise, it would be by (FVer) possible to know this conjunction, but the knowledge of the first conjunct is incompatible with the truth of the second conjunct.[1]

[1] Cf. Fitch (1963). Fitch credits the argument to an anonymous referee of one of his earlier papers. The suggested formalization of (Ver) is actually not to be found in Fitch's paper. Fitch himself did not introduce any special symbol for the possibility operator. Furthermore, he pointed out that a similar argument applies not just to knowledge, but to every concept that implies truth and distributes over conjunction. The argument got the attention of the public when it was revived by W. D. Hart in a footnote (!) in Hart (1979).

This way of formalizing (Ver) was criticized by Dorothy Edgington. According to her interpretation of verificationism, the actual truth of a proposition implies that it could be known in some possible situation that the proposition holds in the *actual* situation. In other words, if ϕ holds in the actual situation w_0, then there is some possible situation v in which it is known that ϕ holds in w_0. Thus, if $\phi \wedge \neg K\phi$ holds in w_0, then in some possible v this conjunction should be known to hold in w_0. It is not required that the conjunction in question is to be known in v to hold in v itself.

Suppose that our object language contains the operator A - "it is actually the case that ..." - with the following truth condition:

$$\models_v A\phi \text{ iff } \models_{w0} \phi,$$

where w_0 is the designated world of the model - the actual world.

Then we can formalize the verificationist claim as follows:

(EVer) $A\phi \rightarrow \Diamond KA\phi$. [2]

How is it possible to have in v an appropriate knowledge concerning an alternative situation, such as w_0? How is it possible to know in v that this or that holds in in a *particular* possible situation?

Edgington suggests one way in which we can know such things: our knowledge of counterfactuals provides us with knowledge of counterfactual situations. In particular, this is how I can know of a certain possible situation that something or other would have been true in that situation without being known there. Here is Edgington's example:

> Suppose I am fortunate enough to chance upon a discovery which no one else is in a position to make. I am an astronomer, and am the only person to observe a supernova before it disappears for ever, say. Then I can know (or reasonably believe), that if I had not been star-gazing last night, this celestial body would not have been discovered; but it would still have been there. (*ibid.*, p. 563)

Let us re-construct the example. Suppose that in the actual situation, w_0, the supernova appears for a short time but is never observed. Thus:

True in w_0: Supernova $\wedge$ $\neg$K(Supernova).

---

[2] Cf. Edgington (1985). A somewhat similar suggestion was independently made by George Schlesinger in Schlesinger (1985), pp. 104-5. Schlesinger's discussion is, however, much less clear than Edgington's.

There is a possible situation, v, in which supernova happens to be observed, but which otherwise differs as little as possible from w_0. In that situation, the following may well hold in v:

True in v: K(No Observation $\Box\!\!\rightarrow$ Supernova $\wedge$ $\neg$K(Supernova)).

Now, if we accept a Stalnaker-Lewis-style semantics for counterfactuals, then, insofar as the closest No Observation-situation to v is w_0 itself, we can conclude that

In v, it is known of w_0 that *Supernova* $\wedge$ $\neg$*K(Supernova)* holds in w_0.[3]

---

[3] There are serious problems in making this step that we do not really know how to deal with. Here are some of them.

(i) What kind of knowledge are we here talking about? *De re* or *de dicto*? The *de re* interpretation might seem appropriate but can we have a *de re* knowledge of merely possible situations? Doesn't *de re* knowledge consist in a relation between the knower and the known? But how can a relation obtain if one of the relata is a merely possible object?

(ii) We cannot accept the following simple rule as a general sufficient condition for knowledge of possible situations:

In v, it is known of w that ϕ holds in w *if* there is some proposition ψ such that
(1) In v, it is known that $\psi \Box\!\!\rightarrow \phi$;
(2) w is one of the closest ψ-worlds to v.

It is easily shown that such a condition is excessively permissive. When (2) is satisfied by some ψ, then it can be shown for *any* ϕ true in w, that (1) and (2) are satisfied for ϕ and ψ, at least provided that (i) logical truths are known in v, and (ii) the closest-world choice-function, which for any ψ picks out those worlds among the ψ-worlds that are closest to a given world, satisfies the so-called property α. (As is well-known, a choice function C satisfies α iff for all non-empty sets X and Y in the range of C, if $X \subseteq Y$, then $X \cap C(Y) \subseteq C(X)$.) In other words, if the above condition would be sufficient for knowledge, then it could be shown that if w is the closest ψ-world to v, then *everything* is known in v of w!

Proof: Suppose that ψ satisfies clause (2). Since the following formula:

$$\psi \wedge \phi \Box\!\!\rightarrow \phi,$$

is a logical truth, we can assume that

(1') In v, it is known that $(\psi \wedge \phi) \Box\!\!\rightarrow \phi$.

Also, since (2) holds, then for any ϕ true in w, the property α implies that

(2') w is one of the closest $(\psi \wedge \phi)$-worlds to v.

But then the following holds in w_0:

True in w_0: $A(\text{Supernova} \wedge \neg K(\text{Supernova})) \wedge \Diamond KA(\text{Supernova} \wedge \neg K(\text{Supernova}))$.

While we are not verificationists, we would like to understand the verificationist position. We would like to understand what is being claimed when one says that truth implies knowability. We do not believe that such a claim is so incoherent as Fitch's paradox might suggest. This is why we consider Edgington's alternative formalization of verificationism worth exploring.

There is another reason why her formalization seems interesting. Some people might want to *restrict* the verificationist claim to such propositions that do not themselves contain any epistemic components. This would allow them to avoid Fitch's paradox while keeping his formalization of verificationism. They might point out that we still could hold to

(FVer) $\phi \rightarrow \Diamond K\phi$,

provided that we restrict the range of ϕ to *non-epistemic* propositions.

Q.E.D.

Essentially the same argument appears in Williamson (1987).

It might be objected that the logically true counterfactual $(\psi \wedge \phi) \;\Box\!\!\rightarrow\; \phi$, which has been used in this trivialization proof, is itself "too trivial" to yield any knowledge of the counterfactual situation. Perhaps then we should qualify the suggested sufficient condition by a demand that the relevant counterfactual should not be logically true.

Unfortunately, however, this qualification does not help. As Timothy Williamson has pointed out (in private communication), the trivialization threat still exists. Thus, suppose that ψ satisfies clause (2), and "let χ state something utterly bizarre, logically quite independent of both ψ and ϕ, such that it is obvious in v that there are worlds much closer to v in which $\psi \wedge \phi$ is true than any in which χ is true." Then the substitution of $(\psi \wedge \phi) \vee \chi$ in place of ψ in the sufficient condition above will still make clauses (1) and (2) true even though the counterfactual $(\psi \wedge \phi) \vee \chi \;\Box\!\!\rightarrow\; \phi$ is not a logical truth.

Is it possible to qualify the condition above in some other way, so as to avoid all the trivialization threats? We are not sure.

(iii) Another thing we are unclear about is whether clause (2) in the condition should not be strenghtened so as to demand that w is *the* (unique) closest ψ-world to v. Is the case for the knowledge of possibilia stronger given such a uniqueness assumption?

However, the verificationist position, as it is usually presented, does not merely amount to the claim that Truth implies Knowability. Verificationists also want to make the converse claim:

(cver) Knowability implies Truth.

According to them, Truth and Knowability are equivalent notions. It is obvious, however, that this converse claim cannot be formalized as

$$\Diamond K\phi \rightarrow \phi,$$

even if we only allow non-epistemic propositions in the range of ϕ. Surely, if ϕ happens to be a *contingent* proposition, nothing prevents ϕ to be false even though it is possible that ϕ could be true and known to be true!

This difficulty disappears after the introduction of the actuality operator. The converse of Edgington's (Ever),

(Ecver) $\Diamond KA\phi \rightarrow A\phi,$

is unobjectionable from any point of view, whether one is a verificationist or not.

Before we continue, we should make one final comment, concerning the logical framework we are going to use. In what follows, we assume classical logic and a possible-worlds semantics for modal notions. Both these assumptions may well be questioned but we want to see how much can be done with these somewhat old-fashioned formal tools.[4]

---

[4] Edgington herself suggests that we replace complete possible worlds with partial "possibilities" or "possible situations", in the style of Humberstone (1981). Humberstone assumes that a proposition may be true, false, or else lack a truth value in a given possibility, and that possibilities are partially ordered by a refinement relation with respect to which both truth and falsity are taken to be persistent (i.e.,hereditary). The model is very similar to the Kripke-semantics for intuitionistic logic. Humberstone gives intuitionistic truth-conditions to his primitive sentence connectives - negation and conjunction. (In particular, he takes $\neg\phi$ to be true in a possibility X iff ϕ is never true in any refinement of X; but he then defines disjunction as the negated conjunction of the negated disjuncts, and implication as the negated conjunction of the antecedent with the negated consequent. As a result, his truth-conditions for disjunction and implication are less demanding than the standard intuitionistic ones.) Then, however, he retrieves classical logic by imposing the condition of refinability: If ϕ lacks a truth-value in X, then there is a refinement of X in which ϕ is true and another refinement of X in which ϕ is false.

This move from worlds to possibilities may have some advantages in the present context: It is more plausible to claim that knowledge of counterfactuals gives us knowledge of counterfactual possibilities, partly because particular possibilities seem generally easier to know due to their smaller size, and partly because it is more reasonable to assume *unique* counterfactual possibilities than to assume unique counterfactual worlds. It is easier to assume the existence of a unique closest possibility in which the antecedent of the counterfactual is true than to make the corresponding assumption about complete worlds. (On the other hand, as we have pointed out in the previous note, it is not clear whether this kind of uniqueness strenghtens the case for the knowledge of possibilia.)

Note that the verificationist claim does *not* automatically commit one to a non-classical logical framework. Philosophers like Bolzano or logical empiricists took verificationism seriously[5] but they still kept the law of excluded middle. In short, we are not interested here in verificationism seen as a part of the intuitionistic research programme.

2. Paradox Regained

We move now to the central problem of this paper. While Edgington's introduction of the actuality operator dissolves Fitch's paradox, our troubles are not yet over. As we remember, the following clause is supposed to specify the truth-condition for the actuality operator:

$$\models_v A\phi \text{ iff } \models_{w0} \phi.$$

Let us see what happens when we combine this truth-condition with the standard truth-clauses for necessity and knowledge, formulated in terms of appropriate accessibility relations between worlds:

$$\models_v \Box\phi \text{ iff } \models_w \phi, \text{ for every w such that vNw;}$$

$$\models_v K\phi \text{ iff } \models_w \phi, \text{ for every w such that vEw.}$$

N and E stand here for the alethic and the epistemic accessibility relations, respectively. Thus, intuitively, vNw iff all the propositions that are necessary in v are true in w.[6] Analogously, vEw iff all the propositions that are known in v are true in w.

The models are of the form:

$$\langle W, w_0, \Pi, N, E, V\rangle$$

where W is a set of worlds, $w_0 \in W$ (w_0 is the designated world of the model), Π a non-empty family of subsets of W (the set of propositions of the model); N, E $\subseteq$ W $\times$ W, N and E are both reflexive relations, and a valuation V is an assignment of truth-values to atomic sentences at

---

To keep things simple, however, we prefer to work with possible worlds in this paper. A move from worlds to possibilities must await another occasion.

[5] As for the reference to Bolzano, see his *Wissenschaftslehre*, para. 314. Bolzano argues there that one cannot specify any limit to our knowledge. He admits that such a limit may still exist even though it is unspecifiable, but he ends his discussion with an optimistic declaration of faith: "I am more inclined to suspect that we are not in a position to identify such a limit because in fact there *is* none, in that the sum total of human knowledge admits of being enlarged to an infinite degree." Cf. Bolzano (1973), p.351.

[6] Here we may choose the notion of necessity and, correspondingly, the notion of possibility that are appropriate for the purposes at hand. Depending on the choice, we get different versions of the verificationist claim.

different worlds. Since we want every sentence to express a proposition, we demand that (i) for every atomic sentence, the set of worlds at which V assigns truth to this sentence belongs to Π; (ii) Π is closed under complement and finite intersection; (iii) Π is closed under the two operations that correspond to $\Box$ and K, respectively: if π belongs to Π, then the same applies to $\{w \in W: \forall v\,(wNv \rightarrow v \in \pi)\}$ and $\{w \in W: \forall v\,(wEv \rightarrow v \in \pi)\}$.

We can define validity in two ways:

(1) *Weak* validity: truth in the designated world in every model;

(2) *Strong* validity: truth in all worlds in all models.

Here is an example that illustrates the distinction between the two notions of validity. The formula

$$\phi \rightarrow A\phi$$

is weakly, but not strongly valid.

Another example: while

(Ecver) $\Diamond KA\phi \rightarrow A\phi$

is strongly valid, the following formula

(cv) $\Diamond KA\phi \rightarrow \phi$

is weakly valid.

In what follows, we are going to interpret validity as *weak* validity - mainly because weak validity seems to be the more interesting and the more natural notion.

Now, it is easy to see that something is seriously wrong with the semantic framework that has been sketched here. In particular,

$$\phi \rightarrow KA\phi$$

turns out to be valid!

Proof: We have to show that, if $\models_{w0}\phi$, then $\models_{w0}KA\phi$. But if $\models_{w0}\phi$, then $A\phi$ is true in *every* world in the model. Therefore, in particular, $\models_v A\phi$ for every v such that w_0Ev. But then the truth-condition for K implies that $\models_{w0}KA\phi$.

We get equally counter-intuitive results for strong validity. It can be shown that

$$A\phi \rightarrow KA\phi$$

is strongly valid. In fact, one can show that

$$A\phi \rightarrow \Box KA\phi$$

is strongly valid, while

$$\phi \rightarrow \Box KA\phi$$

is (weakly) valid, and this is much more than anyone can stomach.

We thus get the old paradox with a vengeance. Without making any verificationist assumptions at all, we are led to the absurd conclusion that whatever happens to be true is known to be actually true. Once again, Truth turns out to imply Knowledge (Necessary Knowledge, this time!).

Surely, something is seriously wrong. The standard truth-conditions for the actuality operator and the epistemic operator do not mix: when we try to combine them, they yield absurdities. It seems we should re-think the whole issue, perhaps start anew, if we want a semantics that allows mixing of epistemic notions with the concept of actuality.[7]

---

[7] Perhaps it should be pointed out that the problem raised in this section cannot be solved by a move from complete worlds to Humberstone's possibilities. While Humberstone does not discuss the epistemic operator in his paper, he has some revealing things to say about one of its close relatives: the belief operator B. He points out that the truth-condition for B may be formulated in two different ways in his semantics:

(1) We may choose the standard approach, that is, we may assume the existence of a dyadic doxastic accessibility relation between possibilities and then take $B\phi$ to be true in a given possibility iff ϕ is true in all doxastically accessible possibilities.

(2) Alternatively, we may postulate the existence of a doxastic function d that assigns to each possibility X its 'belief-possibility' d(X). Intuitively, d(X) comprises all that is believed in X. Then we can take $B\phi$ to be true in X iff ϕ is true in d(X).

Whichever alternative we choose, however, we are going to end up with the problem we have encountered above. Suppose we pick out certain possibility X_0 in the model as the designated one and give the standard truth-condition for the actuality operator in terms of that designated possibility: $\models_X A\phi$ iff $\models_{X_0}\phi$. If ϕ happens to be actually true (that is, if ϕ is true in the designated possibility of the model) then $A\phi$ is going to be true in every possibility in the model. Therefore, we get the absurd result that $BA\phi$ is going to be true everywhere in the model.

But perhaps we should somehow modify the standard truth-condition for A - in such a way as to allow that $A\phi$ may lack truth-value in some possibilities, even though ϕ is actually true? Unfortunately, this solution is simply unavailable. Given Humberstone's refinability condition, $A\phi$ can lack truth-value in X only if it is false in some refinement of X. This is, however, impossible given that ϕ is actually true.

3. Paradox Lost - Two-Dimensional Framework

What follows is a rough sketch of an alternative semantic proposal for the proper mix of the concept of knowledge with the concept of actuality. This is by no means the only possible solution. In particular, after the present paper had been written, Sten Lindström has suggested a competing approach (cf. Lindström, 1993). However, the discussion of Lindström's proposal must await another occasion.

What is distinctive for the present proposal is that we give up the idea of a *fixed actual world* (the designated point of the model) and replace it with a *variable perspective*. In general, a formula will be said to be true *at* a *reference*-world v *from* a *perspective* w. In symbols:

$$^{w}\models_{v} \phi.$$

Here, w and v are the perspective- and the reference-world, respectively.

The intuitive idea behind the symbolism may be expressed as follows: a formula ϕ says something *about* the reference-world, but *what* it says is partially determined by the world of perspective. In other words, a formula ϕ is being interpreted at a reference-world from a given perspective. In the old model, the perspective was fixed: w_0 was the fixed world from which all the formulas were being evaluated. Now, we allow perspective variation in the model. This leads to the following truth-condition for the actuality operator:

$$^{w}\models_{v} A\phi \text{ iff } ^{w}\models_{w} \phi.$$

In other words, $A\phi$ is true at v from w's perspective iff, from w's perspective, ϕ is true at w. (From w's perspective, it is w that constitutes the actual world.)

We have here what is sometimes called a two-dimensional semantics (cf. Segerberg, 1973): a formula is being evaluated not just at one point, v, but at an *ordered pair* of points, <w, v>, with w being the point of perspective and v the point of reference. In what follows, we are going to refer to such pairs of worlds as *states*. Since we do not impose any conditions on states, the set S of states equals W × W. Pairs of the form <w, w>, in which the reference coincides with the perspective, will be called *self-centered* states. The set of self-centered states is what replaces the fixed designated world of the old model. Thus, by (*weak*) *validity* we now mean truth in all self-centered states in all models. *Strong* validity equals truth in *all* the states in all models.[8]

What about the accessibility relations that correspond to modal operators? Well, we can distinguish between two operator-types: *fixed-perspective* and *variable-perspective* operators.

---

[8] Note that the set of valid formulas is easily definable in terms of strong validity: ϕ is valid iff $A\phi$ is strongly valid. There is no equally simple definition of strong validity in terms of validity.

An operator is of a fixed-perspective type if its truth-condition keeps the perspective-world fixed. Necessity is an example of such a fixed-perspective operator:

$$^{w}\models_{v} \Box\phi \text{ iff } ^{w}\models_{v'} \phi \text{ for every v' such that vNv'.}$$

Another example of a fixed-perspective operator is A itself.

We cannot, however, treat the epistemic operator K in the same way. Otherwise, we would get back our original problem. The formula

$$\phi \rightarrow KA\phi$$

would again turn out to be valid. The reason is that, whenever ϕ is true in w from w's own perspective, $A\phi$ is true in all the reference-points provided we keep the perspective w fixed. As a consequence, if K had been a fixed-perspective operator as well, $KA\phi$ would have to be true at every v from the perspective w.

We must therefore allow for perspective-variation in our interpretation of K. We give K's truth-condition in terms of a dyadic epistemic accessibility relation between *states* rather than worlds:

$$^{w}\models_{v} K\phi \text{ iff } ^{w'}\models_{v'} \phi \text{ for every w' and v' such that } <w, v> E <w', v'>.$$

We do *not* assume that, in order for E to obtain between the two pairs, w' must coincide with w. In other words, we do not assume the fixity of the perspective.

Before we say something about the intended interpretation of E, let us first introduce the notion of a model. The models we work with are of the form:

$$\langle W, \Pi, N, E, V\rangle,$$

where W is non-empty, $N \subseteq W \times W$ and $E \subseteq S \times S$ (with S defined as $W \times W$), Π (the set of propositions) is a non-empty family of subsets of S, N and E are both reflexive, and a valuation V is an assignment of truth-values to atomic sentences at different states in S. Since we want all sentences express propositions, we demand that

(i) for every atomic sentence, the set of states at which V assigns truth to this sentence belongs to Π;
(ii) Π is closed under finite intersection and under complement;
and
(iii) Π is closed under the three operations that correspond to the modal operators of the language (necessity, knowledge and actuality): if π belongs to Π, then the same applies to *Necessarily* $\pi = \{<w,v> \in S: \forall v'(vNv' \rightarrow <w,v'> \in \pi)\}$, *It is known that* $\pi = \{<w,v> \in S: \forall w,v'(<w,v>E<w',v'> \rightarrow <w',v'> \in \pi)\}$ and *It is actually the case that* $\pi = \{<w,v> \in S: <w,w> \in \pi\}$.

What is the intuitive interpretation of the epistemic accessibility relation between states? When we consider a state <w, v>, the set of states to which <w, v> bears the epistemic relation E delimits the range of knowledge obtaining in <w, v>. When this set is smaller, the range of knowledge is larger - the epistemic uncertainty decreases. Now, in a given state, there are *two sources* of epistemic uncertainty. Uncertainty concerns in part the reference-point of the state, in part the world that constitutes the point of perspective. Our knowledge of the world that is referred to (= described) and of the the world that constitutes the perspective is more or less limited. The relation E is meant to model both types of uncertainty. This is why, when <w, v> E <w', v'>, we must allow not only that v' may differ from v but also that w' and w may differ. The knowledge in <w, v> may not be sufficient to identify the perspective-point of that state uniquely.

We should, however, add a qualification concerning self-centered states. In a self-centered state, it may be unclear what is the reference-point and what is the point of perspective. But whatever they are, it is certain that they are the same. Thus, for example, when I make a claim about the (reference-)world, I may be uncertain about many things, but I know that the world of my perspective is just the world about which I am making my claim. This suggests the following condition on E:

(Coincidence) If <w, v> E <w', v'> and w = v, then w'= v'.

This condition on E corresponds to the "necessitation" rule for the operator K:

if a formula ϕ is (weakly) valid, then Kϕ must be (weakly) valid as well.

Proof of sufficiency: A formula ϕ is valid iff it is true in all self-centered states. But, by Coincidence, all states that are epistemically accessible to a self-centered state are themselves self-centered. Therefore, Kϕ must be true in every self-centered state.

Proof of necessity: Note that, for every ϕ, A$\phi \rightarrow \phi$ is (weakly) valid. Suppose now that, in a particular model M, Coincidence is violated: for some w, w', v' such that w'$\neq$ v', <w,w>E<w',v'>. (For the precise definition of a model, see below.) Now, consider a model M' that is based on the same set of worlds as M and involves the same accessibility relations, but differs from M in (at most) two respects: in M', the family of propositions (which in the two-dimensional framework are represented by sets of states) is extended if necessary, so as to contain the set I of all the self-centered states, and an atomic formula p is evaluated as true in exactly those states that belong to I. It is easy to see that Ap $\rightarrow$ p is then also true in exactly those states that are self-centered. Therefore, it will be false in <w', v'>. In consequence, K(Ap $\rightarrow$ p) will be false in the self-centered state <w, w> in M'. Thus, K(A$\phi \rightarrow \phi$) will not be valid.

The proof above may be easily adjusted so as to show that Coincidence is also both necessary and sufficient for the weak validity of the formula:

$$K(A\phi \rightarrow \phi).$$[9]

In other words, in a self-centered state it is always known that whatever is actually true is true.

It is easy to show that the present framework does not validate the troublesome formula:

$$\phi \rightarrow KA\phi.$$

In particular, nothing precludes that ${}^{w}\models_{w} \phi$, while at the same time there is some w' such that <w, w> E <w', w'> and not ${}^{w'}\models_{w'} \phi$. Therefore, it is not the case that ${}^{w'}\models_{w'} A\phi$, nor that ${}^{w}\models_{w} KA\phi$. We avoid our original problem precisely because E allows perspective-variation.[10]

On the other hand, the converse of verificationism,

(cv) $$\Diamond KA\phi \rightarrow \phi,$$

is still (weakly) valid - just as we wanted it to be. (The converse of Edgington's formulation of the verificationist claim,

(Ecver) $$\Diamond KA\phi \rightarrow A\phi,$$

is also only weakly valid. In this respect, the two-dimensional approach differs from the one-dimensional one.)

Proof: Suppose that ${}^{w}\models_{w} \Diamond KA\phi$. Then, for some v, ${}^{w}\models_{v} KA\phi$. Therefore, since E is reflexive, ${}^{w}\models_{v} A\phi$. But then ${}^{w}\models_{w} \phi$.

---

[9]Here we are indebted to Timothy Williamson for proving this and for clarifying the relationship between $K(A\phi \rightarrow \phi)$ and the "necessitation" rule for K. That the latter implies the validity of the former is clear: $A\phi \rightarrow \phi$ is valid, and therefore, given necessitation, the same must apply to $K(A\phi \rightarrow \phi)$. To derive the necessitation rule, we note first that, even in the absence of Coincidence, if ϕ is valid, then $KA\phi$ must be valid. (In fact, if ϕ is valid, then $A\phi$ is strongly valid, so that $KA\phi$ is also strongly valid.) Second, since K distributes over implication, the validity of $K(A\phi \rightarrow \phi)$ implies the validity of $KA\phi \rightarrow K\phi$. Thus, it follows by modus ponens that $K\phi$ is valid.

[10] The necessity operator, on the other hand, keeps the perspective fixed. Therefore, the present approach does validate:

$$\phi \rightarrow \Box A\phi.$$

But this is, we believe, as it should be.

What about the verificationist claim itself? Well,

(v) $\phi \rightarrow \Diamond KA\phi$

is not valid, but we can (weakly) validate it, if we wish, by imposing a corresponding extra condition on our models. What this condition says is that for every w in W and every proposition π such that π holds in <w,w>, there exists an alethically accessible v such that, in the state <w, v>, it is known that π holds in <w,w>. Formally:

> For every w in W and π in Π, if <w,w> $\in \pi$, then there exists some v in W such that wNv and, for every w' and v' in W, if <w, v> E <w', v'>, then <w',w'> $\in \pi$.

This condition is both sufficient and necessary for the validity of (v). It also validates Edgington's formula:

(EVer) $A\phi \rightarrow \Diamond KA\phi$.

This should not be surprising. Since $\phi \leftrightarrow A\phi$ is a valid schema, (v) and (Ever) are equivalent claims in our framework: the former is valid iff the latter is valid.[11]

At the same time, it can be shown that there are models that satisfy the relevant condition and in which Truth still does not imply Knowledge. That is, in some such models, for some formulas ϕ and some self-centered states, $\phi \wedge \neg K\phi$ is true in those states. So it is clear that Fitch's paradox is definitely dissolved in the framework under consideration.

4. Two sources of epistemic uncertainty: decomposition of the epistemic accessibility relation

The epistemic uncertainty that E is meant to express has, as we have seen, two sources: lack of knowledge concerning the reference-world and the world of perspective. This suggests that we might try to decompose E into two "components" that correspond to these two sources of uncertainty.

As a first step, let us define two triadic relations on worlds in terms of E:

The perspectival relation: $wP_v w' =_{df}$ for some v', <w, v> E <w', v'>.

---

[11] It should perhaps be noted that the condition under consideration does *not* make Edgington's formula *strongly* valid. If w and u are distinct, it may well be the case that $A\phi$ is true in <w,u> without there being any world v alethically accessible to *the reference-world u* such that it is true in <w,v> that $KA\phi$. The condition under consideration only guarantees that such a world v must be alethically accessible to *the perspective-world w*. Therefore, the corresponding strongly valid formula would instead look as follows: $A\phi \rightarrow A\Diamond KA\phi$.

The referential relation: $vR_wv' =_{df}$ for some w', <w, v> E <w', v'>.

The perspectival relation determines the range of uncertainty concerning the world of perspective in a given state <w,v>; w' ranges over those worlds that in <w,v> are epistemically indistinguishable from the perspective-world w. Analogously, the referential relation determines the range of uncertainty concerning the reference-world in a given state <w,v>.

Now, consider what happens if we assume that these two types of uncertainty are relatively *independent* of each other. In particular, consider what happens if we impose the following extra condition on E:

(Combinability) If wP_vw', vR_wv' and $w \neq v$, then <w,v> E <w', v'>.

In other words: if, relative to a state <w, v>, w' is an epistemically possible perspective-point while another world v' is an epistemically possible reference-point, then w' and v' may be combined into an epistemically possible state relative <w,v>. The demand that w and v should be distinct from each other is motivated by Coincidence. Coincidence restricts combinability.

If we are prepared to accept this condition on E, then E can actually be *defined* from P and R, as follows:

(Equivalence 1) <w, v> E <w', v'> iff (i) wP_vw' and vR_wv', and (ii) if w = v, then w'= v'.

Here is the proof that this equivalence follows from the definitions of P and R in terms of E, given Coincidence and Combinability: The left-to-right direction, clause (i), immediately follows from the definitions in question. Clause (ii) follows by Coincidence. Moving now to the right-to-left direction, we consider two cases: (1) $w \neq v$. Then clause (i) implies the left-hand side of the equivalence by Combinability. (2) w = v. Clause (2) implies then that w' = v'. But if w = v and w' = v', then, by Coincidence and the definitions of P and R, each conjunct in clause (1) implies the left-hand side.

As we have mentioned, one might look at Equivalence 1 as the definition of E in terms of P and R. It is easy to see that such a definition immediately implies that E satisfies both Coincidence and Combinability. Also, it implies that E is reflexive, provided that P and R are reflexive.

Could one go further in this reduction process so as to reduce P and R to just one triadic relation on worlds? The answer is yes, provided we are prepared to impose one additional condition on E.

There is, it seems, an important asymmetry between perspective and reference. When we consider knowledge in a given state, it is clear that this knowledge really *obtains* at the reference-point of that state: it is there that the knower and his knowledge are supposed to be located. This means, in particular, that, as far as the knowledge of the reference-world is concerned, the perspective-world does not really have any role to play: the relevant knowledge *obtains in* the reference-world and *concerns* just this world - the perspective-world is idle! We are therefore led to the following condition:

(Irrelevance of Perspective) For all w and w', if vR_wv' then $vR_{w'}v'$.

In other words, as far the referential uncertainty is concerned, the perspective does not matter.

Irrelevance of Perspective may also be be given the following (equivalent) formulation:

$$\text{For all } w,\ vR_wv' \text{ iff } vR_vv'.$$

I.e., the referential uncertainty in a given state <w,v> is exactly the same as the referential uncertainty in the corresponding self-centered state <v, v>.

But now the following should be noted: in a self-centered state, there is no difference between referential and perspectival uncertainty! By the definitions of P and R and Coincidence, we have:

$$vR_vv' \text{ iff } \langle v, v\rangle\ E\ \langle v', v'\rangle \text{ iff } vP_vv'.$$

We end up, therefore, with the following reduction of R to P:

$$vR_wv' \text{ iff } vR_vv' \text{ iff } vP_vv'.$$

Conclusion: E can be defined in terms of just one triadic relation P on worlds. Substituting vP_vv' for vR_wv' in Equivalence 1, we obtain:

(Equivalence 2) $\langle w, v\rangle\ E\ \langle w', v'\rangle$ iff (i) wP_vw' and vP_vv', and (ii) if $w = v$, then $w' = v'$.

Equivalence 2 may be seen as expressing the following claim: <w', v'> is epistemically possible relative to <w, v> if w' and v' are compatible with what is known in v of w and v, respectively, and if, in addition, w' and v' coincide provided that w and v coincide.

Equivalence 2 immediately implies that E satisfies Coincidence. It also implies that E is reflexive provided that P is reflexive in its first argument. (I.e., if for every w and v, wP_vw.)

Given Equivalence 2, the condition on N and E that corresponds to the verificationist claim,

(v) $\phi \rightarrow \Diamond KA\phi$,

reduces to the following condition on N and P:

For every world w and every proposition π such that $\langle w,w\rangle \in \pi$, there is some v such that wNv and for every w', if wP_vw', then $\langle w',w'\rangle \in \pi$.

Let us sum up the results we have reached. While verificationism may well be unacceptable, it is at least not an absurd position. It can be expressed within a framework that combines the notions of actuality, possibility and knowledge. All these notions may be interpreted in a two-dimensional possible-world semantics, provided we treat the concept of knowledge as a variable-perspective operator. Its truth-condition is formulated in terms of a dyadic accessibility relation between states (pairs of worlds), or - given a couple of simplifying assumptions - in terms of a triadic accessibility relation between worlds.[12]

References

Bolzano, B., 1973, *Theory of Science*, ed. by Jan Berg and translated by Burnham Terrell, D. Reidel Publ. Comp., Dordrecht, Holland.

Edgington, D., 1985, "The Paradox of Knowability", *Mind* 93, pp. 557 - 68.

Fitch, F. B., 1963, "A Logical Analysis of Some Value Concepts", *The Journal of Symbolic Logic* 28, pp. 135 - 42.

Humberstone, I. L., 1981, "From Worlds to Possibilities", *Journal of Philosophical Logic* 10, pp. 313 - 339.

Hart, W. D., 1979, "The Epistemology of Abstract Objects", *Proceedings of the Aristotelian Society*, suppl. vol. 53, 1979, pp. 153-65.

Lindström, S., 1993, "Situations, Truth and Knowability - A Situation-Theoretic Analysis of a Paradox by Fitch", draft.

Schlesinger, G., 1985, *The Range of Epistemic Logic*, Aberdeen University Press, Aberdeen.

Segerberg, K., 1973, "Two-Dimensional Modal Logic", *Journal of Philosophical Logic* 2, pp. 77 - 96.

Williamson, T., 1987, "On the Paradox of Knowability", *Mind* 95, pp. 256-61.

[12] This paper owes much to our heated but extremely enjoyable discussions with Sten Lindström. We have also been greatly helped by Timothy Williamson. The earlier versions of this paper have been presented at a truth-and-knowability symposium in Stockholm, and at a seminar in philosophical logic in Uppsala. We are grateful to the members of these audiences for their stimulating comments. In particular, we are thankful for the suggestions offered by Peter Pagin and Fredrik Stjernberg.