8

Knowability and Possible Epistemic Oddities

JC Beall

1. Non-omniscience and the Knowability Rule

Our world is non-omniscient. Nobody knows all truths, and nobody ever will. Does it follow that there are unknowable truths? Frederic Fitch (1963) 'proved' the affirmative. In short, if some truth is unknown, then that it is unknown is itself unknowable; hence, given non-omniscience, there is some unknowable truth.

Verificationists, who tie truth to verifiability, are committed to the so-called knowability rule (henceforth, KP).[1] Let K be the epistemic operator it is known by someone at some time that . . . , and $\Diamond$ the aletheic it is possible that. . . . KP is the following rule.

$$\frac{\alpha}{\Diamond K\alpha}$$

Non-omniscience gives us $\alpha \wedge \neg K\alpha$, for some α. KP, in turn, gives us $\Diamond K(\alpha \wedge \neg K\alpha)$. A few related rules governing $\Diamond$ and K quickly yield $\Diamond(K\alpha \wedge \neg K\alpha)$, the possibility of 'true contradictions'. (For the relevant rules, see Section 2) Fitch's proof qua *reductio* makes the final step: KP is unsound.

In this paper, my concern is not so much with Fitch's 'proof' against KP (or the conditional version). The proof is blocked in familiar 'paraconsistent' and 'paracomplete' logics (see Sections 2 and 3), both of which are independently motivated and, hence, available to verificationists. Rather, my concern is with the apparent commitment to 'possibly true contradictions'.

For discussion I am grateful to Colin Caret, Carrie Jenkins, Graham Priest, Greg Restall, and various members of the AHRC Arché Centre for Logic, Language, Mathematics, and Mind. Thanks also to Joe Salerno for editing the volume.

[1] The key idea is often given as a (universally quantified) conditional principle, but, to simplify current discussion, I focus on the rule form. (The conditional version is often called 'KP', short for the knowability principle, but little confusion should result.)

Regardless of its effect on verificationism, Fitch's 'proof' highlights the oddity of epistemic optimism in a non-omniscient world. Given non-omniscience, KP involves something out of the ordinary. My main aim is to briefly explore a few options for cashing out the given oddity.

The paper runs as follows. In Section 2, I briefly set out the relevant rules (and one corresponding premise) on which the Fitch argument relies. Section 3, in turn, briefly reviews the main point of Beall, which suggests a paraconsistent response to Fitch's 'proof' qua *reductio* of KP; however, the same considerations also motivate a similar paracomplete, non-paraconsistent, response, on which I will focus.[2] In Section 4 I set out the main focus: what to make of Fitch's (initial) argument for the 'possibility of gluts'. While the paracomplete response undercuts Fitch's 'proof' qua *reductio*, the issue of 'possible, true contradictions' remains open—especially in a non-paraconsistent framework, which is the target. I explore two options. In Section 5 I suggest but reject a flat-footed option: living with possible—but merely possible—inconsistency. Section 6, in turn, explores the other option: avoiding even the 'mere possibility' of 'true inconsistencies'. As a sort of synthesis, Section 7 briefly sketches another option: a paracomplete and paraconsistent framework. Section 8 closes with some general comments and (brief) responses to objections.

2. Fitch's Proof, in Short

For present purposes, the basic rules, involved in Fitch's proof, may be divided into four categories: epistemic, (aletheic) modal, modal–epistemic (viz., KP), and 'background'. The rules run as follows.[3]

1. Epistemic rules

Veridicality (KV). The idea is that 'knowledge implies truth'.

$$\frac{K\alpha}{\alpha}$$

Distribution (KC). That a conjunction is known implies that its conjuncts are known.

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

[2] See Priest's (Ch. 7, this volume) for a development of the *LP*-based paraconsistent (indeed, dialetheic) position, and Section 7 for an alternative paraconsistent framework.

[3] To facilitate comparison with Priest's (Ch. 7, this volume), which discusses the related paraconsistent—indeed, dialetheic—response, I use a natural deduction version of the rules. As in Priest's paper, $[\alpha]$, in the context of rule (or proof), indicates a discharged assumption.

2. Aletheic modal rules

Non-contradiction (LNC). It is false that it's possible that $\alpha \wedge \neg\alpha$ is true, for any α.[4]

$$\frac{}{\neg\Diamond(\alpha \wedge \neg\alpha)}$$

Closure (CP). That α is possible and that α implies β implies that β is possible.

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

3. Modal–Epistemic rule

Knowability (KP). The idea is that 'truth implies knowability', the key verificationist position.

$$\frac{\alpha}{\Diamond K\alpha}$$

4. Background Rules

Adjunction and Simplification. Conjunction behaves normally.

$$\frac{\alpha \quad \beta}{\alpha \wedge \beta} \qquad \frac{\alpha \wedge \beta}{\alpha \quad \beta}$$

Contraposition. That β is not true and that α implies β implies that α is not true.

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

Fitch's Proof, in short: Suppose, for *reductio*, $\alpha \wedge \neg K\alpha$, for some α. KP yields $\Diamond K(\alpha \wedge \neg K\alpha)$. Given KC, we have that $K(\alpha \wedge \neg K\alpha) \vdash K\alpha \wedge K\neg K\alpha$. Simplification, VK, and Adjunction (and transitivity of implication), yield that $K(\alpha \wedge \neg K\alpha) \vdash K\alpha \wedge \neg K\alpha$. But, then, CP gives $\Diamond(K\alpha \wedge \neg K\alpha)$. LNC, in turn, delivers $\neg\Diamond(K\alpha \wedge \neg K\alpha)$. Contradiction.

[4] Given the inter-definability of $\Box\alpha$ and $\neg\Diamond\neg\alpha$, which is assumed, LNC (as here put) amounts to the validity of 'inferring' $\Box\neg(\alpha \wedge \neg\alpha)$, for all α, from no premises—i.e., as theorem.

3. Verificationism, the Knower, and Fitch's 'Proof'

While I am not a verificationist, I agree with those who think that verificationism is not undermined by the 'proof'. The result shows only that, given non-omniscience, verificationists cannot consistently endorse all of the rules involved in Fitch's 'proof'. Since the rules in question are largely 'classical' (e.g., LNC, Contraposition), verificationism is best understood in a non-classical framework, one in which some of the given rules are invalid.

One might worry that going non-classical is ad hoc. Were there no independent reason to reject some of the given rules, the worry would be warranted. But there are independent reasons to reject some of the given rules. Familiar semantic paradoxes, cases of vagueness, or other commonly 'deviant' phenomena, motivate familiar logics in which some of the given rules fail—notably, the LNC or Contraposition.

Consider, for example, the Knower paradox, which involves a sentence κ that says of itself (only) that it is not known.[5] Given LEM, κ is either known or not. In the latter case, κ is not known, and hence true. In the former case, KV gives us that κ is true, in which case κ is not known. Either way, κ is not known, and, so, κ is true. But, now, we have a proof that κ is true, and hence—on the basis of our proof—we know that κ is true. The upshot: there is a sentence, namely, κ, such that we know that κ is true but, as κ says, do not know that κ is true.

In response to the Knower (or many such paradoxes), one might, as in Beall (2000), take the Knower to independently motivate 'dialetheism' with respect to knowledge—that $K\alpha \wedge \neg K\alpha$ is true, for some α. On such a line, a paraconsistent logic, in which such inconsistency is 'harnessed', is motivated. But, then, at least Contraposition is invalid—and, hence, Fitch's 'proof' fails. Graham Priest (Ch. 7, this volume) advocates just such a line in the context of *LP*, a paraconsistent logic in which both LEM and LNC are valid.

On the other hand, a paracomplete (and non-paraconsistent) response to the Knower is equally natural. A paracomplete logic is one in which LEM is invalid.[6] Unlike the dialetheic response, a paracomplete theorist rejects the Knower instances of LEM. One familiar paracomplete framework is K_3, the Strong Kleene framework.[7] In such a logic, not only is Contraposition invalid,

[5] Here, I simply use 'is known' rather than the operator. In a suitably non-classical framework, we can enjoy a genuine (intersubstitutable) truth predicate that, in turn, diminishes the importance of distinguishing between operators and predicates.

[6] Accordingly, Intuitionistic logic counts as paracomplete. I will not discuss the Intuitionistic options, as these are well known. (Besides, Intuitionistic logic does not afford viable options for the broader class of paradoxes—e.g., Liars, etc.)

[7] I set issues of a suitable conditional aside, and concentrate mostly on the given 'rules'. A suitable conditional is an important issue, but it is one that would take the discussion too far afield.

but LNC is also invalid.[8] Accordingly, a verificationist who, for independent reasons, endorses K_3 (or some suitable extension), need not worry about Fitch's proof qua *reductio* of KP.

While I am (very) sympathetic with the paraconsistent—indeed, dialetheic—framework, I will focus on non-dialetheic and, except, for Section 7, non-paraconsistent but paracomplete responses. Either way, verificationists have independent reason—e.g., Knower or the like—to endorse a non-classical logic in which Fitch's 'proof', qua *reductio* of KP, fails.

4. The Real Issue: Possibly True Gluts?

Though verificationists needn't worry about Fitch's proof qua *reductio* of KP, there is more to Fitch's argument than the final few (*reductio*) steps. As in Section 1, Fitch's argument highlights the oddity of epistemic optimism in a non-omniscient world. Contraposition and LNC aside, the remaining rules (see Section 2) still leave a curiosity: the apparent commitment to 'possibly true contradictions'.

To see the issue, we assume an extension of K_3 in which the relevant rules remain, except, of course, for Contraposition and LNC. (See Section 5 for the natural semantics.) The initial steps of Fitch's proof still go through: non-omniscience gives us $\alpha \wedge \neg K\alpha$, for some α. KP gives us $\Diamond K(\alpha \wedge \neg K\alpha)$. But $K(\alpha \wedge \neg K\alpha) \vdash K\alpha \wedge K\neg K\alpha$ from KC. Simplification, VK, and Adjunction (and transitivity of implication), yield that $K(\alpha \wedge \neg K\alpha) \vdash K\alpha \wedge \neg K\alpha$. But, then, CP gives $\Diamond(K\alpha \wedge \neg K\alpha)$. So, losing Contraposition or LNC still leaves the noted oddity.

In a non-paraconsistent setting, 'possibly true contradictions' are at least curious. The real issue, then, is what to make of the given oddity in a non-paraconsistent, paracomplete setting. How should a paracomplete, non-paraconsistent verificationist—or KP theorist, in general—respond to the apparent 'possibly true contradictions'?

As in Section 3, I will explore two salient options. The first option is to simply live with the given 'oddity'. The second option, rejecting even the 'mere possibility' of 'true inconsistency', involves expanding one's space of possibilities while restricting one's account of validity. After discussing such (non-paraconsistent) options, I turn to a brief sketch of a 'compromise', a non-dialetheic but nonetheless paraconsistent and paracomplete framework. I will then close (in Section 8) with general comments, briefly answering two objections.

[8] According to K_3, 'Explosion' is valid: $\alpha \wedge \neg\alpha \vdash \beta$. But LEM is invalid: $\nvdash \neg(\alpha \wedge \neg\alpha)$. Were Contraposition valid, we'd immediately have $\neg\beta \vdash \neg\alpha \vee \alpha$, for any β and α. But we don't have that in K_3, since, as said, we do not have LEM.

5. Living with Merely Possible Gluts

Given non-omniscience, verificationists—and KP theorists, in general—are apparently committed to the possibility of 'true contradictions'. In a non-paraconsistent context, which is the chief concern in this paper, such a commitment is curious. The question is: what to make of it?

The flat-footed option is to just live with it. On the surface, the 'possibility of true contradictions' is startling. Upon inspection, though, the situation is in many respects mundane, especially if there's exactly one—non-actual, merely possible—such 'possibility'. The flat-footed response acknowledges a (unique) trivial world, and she learns to live with it.

To make the idea clearer, I briefly sketch a basic—paracomplete but non-paraconsistent—semantics. I then return to the flat-footed response.

5.1. Paracomplete semantics with the trivial world

We are considering a paracomplete and non-paraconsistent framework for verificationism (or KP theorist, in general), one in which there's a unique possibility of 'true contradictions'. By way of contrast with the natural *LP*-based paraconsistent framework,[9] I focus on an extension of K_3, the Strong Kleene framework.

Our set of semantic values, namely, $\mathcal{V} = \{1, .5, 0\}$, is ordered in the standard way. $\mathcal{D}$, our designated values, comprises exactly 1. In addition to our usual extensional connectives, we add two unary connectives, the epistemic K and the aletheic $\Diamond$. (We define $\Box$ as $\neg\Diamond\neg$.) K and $\Diamond$ are intended to be modal connectives. Accordingly, we pursue a modal extension of K_3.

Interpretations are structures $\langle \mathcal{W}, R, E, v, w_\perp \rangle$, where $\mathcal{W} \cap \{w_\perp\}$ comprises 'worlds', with $w_\perp \not\in \mathcal{W}$ the trivial world. R and E are binary relations on $\mathcal{W} \cup \{w_\perp\}$ (each at least reflexive), and $v : \mathcal{A} \times \mathcal{W} \to \mathcal{V}$ is a valuation from atomics and worlds into $\{1, .5, 0\}$. For convenience, we let $v_w(\alpha) = v(\alpha, w)$, this being the value of α at w.

The value of any sentence at any $w \in \mathcal{W}$ is achieved via the following clauses.[10]

$$\begin{aligned}
v_w(\neg\alpha) &= 1 - v_w(\alpha)\\
v_w(\alpha \wedge \beta) &= \min\{v_w(\alpha), v_w(\beta)\}\\
v_w(\alpha \vee \beta) &= \max\{v_w(\alpha), v_w(\beta)\}\\
v_w(\Diamond\alpha) &= \max\{v_{w'}(\alpha) : wRw' \text{ for any } w' \in \mathcal{W} \cup \{w_\perp\}\}\\
v_w(K\alpha) &= \min\{v_{w'}(\alpha) : wEw' \text{ for any } w' \in \mathcal{W} \cup \{w_\perp\}\}
\end{aligned}$$

[9] See Priest (Ch. 7, this volume).

[10] Except for the clause concerning $w_\perp$, the following are the standard Kleene clauses for modal connectives.

With respect to $w_\perp$, the trivial world, the clause for any interpretation is the obvious (trivial!) one:

$$v_\perp(\alpha) = 1$$

Finally, we define validity as 'truth-preservation' over all worlds of all interpretations.

A few features of the framework are notable. To begin, the semantics is clearly paracomplete in that $\alpha \vee \neg\alpha$ is invalid. Just consider a model in which $v_w(\alpha) = 0.5$, for some $w \in \mathcal{W}$. Since, as one may verify, $\alpha \vee \neg\alpha$ and $\neg(\alpha \wedge \neg\alpha)$ are equivalent in the semantics, the same (counter-) model serves to invalidate LNC (see Section 3). Similarly for Contraposition.[11]

On the other hand, it is clear that Adjunction and Simplification are valid. Moreover, and more to the current point, the remaining epistemic and aletheic-modal rules are all validated. (See Section 2.)

KV. Suppose that $v_w(K\alpha) = 1$, for some $w \in \mathcal{W}$ and α. Since R is reflexive, we have it that $v_w(\alpha) = 1$.

KC. Suppose that $v_w(K(\alpha \wedge \beta)) = 1$. Then $v_{w'}(\alpha \wedge \beta) = 1 = v_{w'}(\alpha) = v_{w'}(\beta)$ for all $w' \in \mathcal{W}$ such that wRw'. But, then, $v_w(K\alpha) = 1 = v_w(K\beta)$.

CP. Suppose that $\alpha \vdash \beta$ and, for some interpretation, $v_w(\Diamond\alpha) = 1$. Then $v_{w'}(\alpha) = 1$, for some w such that wRw'. But, by supposition, there's no world, in any interpretation, at which α is true and β not true. Hence, $v_{w'}(\beta) = 1$, and so $v_w(\Diamond\beta) = 1$.

The question, of course, turns to our essential modal–epistemic rule KP, which is not valid on the current semantics. Can the semantics be tweaked to ensure the validity of KP? Yes. Indeed, the whole point of invoking $w_\perp$, which has thus far played no role, is to ensure the validity of KP. To achieve KP we stipulate that $wRw_\perp$, for all worlds w (including $w_\perp$). That KP is now valid is obvious; it is vacuously so.[12]

So, except for Contraposition and LNC—which, as in Section 3, are suspect for independent reasons—the current framework preserves all of the key rules, including KP. Because KP is preserved, (actual) non-omniscience forces an oddity: the trivial world. Indeed, so long as validity is defined as 'all points validity' (e.g., truth-preservation at all worlds), then, unless one goes with a paraconsistent framework, I see no way to avoid the trivial world without giving up KP.[13] The issue, to which I now return, is whether such oddity is too odd.

[11] The corresponding *LP*-based paraconsistent framework validates LNC but, as here, not Contraposition. See Priest's (Ch. 7, this volume).

[12] Compare the *LP*-based dialetheic model (Ibid.), which likewise invokes $w_\perp$ for the same job. (One difference, of course, is that the trivial world naturally falls out of the *LP* framework, whereas here it is at least curious.)

[13] If one endorses a paraconsistent logic, a more natural paracomplete framework might be had. See Section 7 for a sketch.

5.2. The flat-footed response

KP, the reflection of high epistemic optimism, produces an oddity in a non-omniscient world. The apparent oddity, given the going rules (except Contraposition and LNC), is the possibility of true contradictions. The flat-footed response to such oddity is to accept it, but accept it as merely possible and, importantly, a unique case. Given non-omniscience, KP is preserved in virtue of the unique trivial world—the possibility in which 'true contradictions' occur.

Is the trivial world too high a price to pay for KP? The answer is not obvious. Admittedly, it may be very difficult to fully understand the trivial world. While one can easily understand that the trivial world is the world at which every sentence is true, it is not easy to understand what such a world is like. Still, there are a few things that can be said on the trivial world's behalf.

1. KP! As in Section 5.1, if validity is to be understood as all points validity (e.g., truth-preservation at all worlds), it is difficult to retain KP without the trivial world—unless one goes with a paraconsistent logic, which is set aside at this stage. (See Section 7.) So, one virtue of the trivial world is that it affords the chief desideratum for a non-classical verificationism: it preserves KP.
2. Concrete Explosion! In the current semantics we have 'explosion', that is, $\alpha, \neg\alpha \vdash \beta$. In many (most) non-paraconsistent logics, explosion itself is vacuously achieved: it is valid in virtue of no interpretation in which the premises are true. Here, we have 'concrete evidence' of explosion: any world in which $\alpha \wedge \neg\alpha$ is true is the explosive one in which everything is true. There is something to say for such 'concrete evidence' (although I wouldn't put too much weight on this).
3. Merely possible! Similarly, while the possibility of 'true contradictions' sounds startling at first, the current proposal is rather mundane. After all, in discussing the possibility of 'true contradictions', one may quickly point out that we're talking about a unique and limit case—the merely possible trivial world.

In addition to (1)–(3), there is another—perhaps the strongest—point to consider. As throughout, KP's validity in a non-omniscient world is indeed odd. One reason we might think it odd is that it clashes against the 'normal' behaviour of our connectives—which behaviour, perhaps, is by and large classical. The trivial world, which, on the current proposal, is the result of KP's validity (and a non-omniscient world), might best be seen as a world in which our connectives 'go on holiday'. Clearly, the connectives are not behaving normally at $w_\perp$. Perhaps such abnormal behaviour is the price of KP's validity, given non-omniscience.

5.3. Trouble with flat-footedness

Despite its virtues (if virtues they be), the trivial world is nonetheless disappointing in the current context. While I do not think the trivial world itself is terribly objectionable, its role in the current context is prima facie problematic. The heart of verificationism is KP, a rule that, at least traditionally, has served to distinguish verificationists from non-verificationists. On the current proposal, KP is preserved—indeed, its validity achieved—solely in virtue of the trivial world. But, now, the traditional role of KP cannot be served. After all, it is obvious that anyone—even a classical logician—could acknowledge the trivial world, at least in the fashion in Section 5.1. But if anyone can have the trivial world, anyone can have KP. Surely verificationism is more demanding than that.[14]

The flat-footed response, then, is ultimately unsatisfactory. Unfortunately, without going paraconsistent (though not necessarily dialetheic), there is no obvious way to preserve KP without the trivial world, at least if validity remains 'all points validity'. Giving up such a notion of validity provides an alternative paracomplete approach, to which I now briefly turn.

6. Abnormal Epistemic Possibilities

In Section 5, I suggested—but found wanting—the 'flat-footed' paracomplete response to the verificationist's apparent commitment to possibly true contradictions. Might an alternative paracomplete response do away with the 'possibly true contradictions' altogether? In this section, I briefly explore one route towards doing as much.[15] On this approach, the oddity of KP in a non-omniscient world is not 'possibly true contradictions', but rather the sheer oddity of possibly knowing an unknown truth.

I will first give a philosophical sketch of the idea, followed by a slightly more formal sketch, and then offer a few comments on the overall framework.

6.1. The philosophical story

Verificationists tie truth to verification. An essential ingredient of the connection is reflected in (at least) KP. What Fitch seemed to show is that, given non-omniscience, KP leads to possibly true contradictions. But perhaps another

[14] I should say that this point might affect Priest's *LP*-based proposal (Ch. 7, this volume). A more natural (paraconsistent) approach, not subject to the same problem, is briefly sketched in Section 7.

[15] Other options are available, of course, if the extensional connectives (negation, conjunction, disjunction) behave non-standardly, but I am chiefly interested in 'normal' behaviour for such (extensional) connectives—i.e., classical input, classical output. (One could go weaker than Strong Kleene, but independent motivation for such logics is more difficult to find than for the K_3 case.)

lesson may be drawn. In particular, what verificationists are committed to is not some possibly true contradiction; rather, they're committed to epistemically abnormal—but none the less entirely (aletheically) possible—worlds, worlds in which, for example, knowing an unknown truth happens.

Verificationists are committed to $\Diamond(\alpha \wedge \neg K\alpha)$, for some α, but the possibility in question is epistemically abnormal, a world, perhaps, in which 'epistemic fictions' transpire.[16] At such worlds, the normal behaviour of K breaks down in various respects. In particular, given that possibilities are one and all consistent (though not necessarily complete), the normal distributive behaviour of K breaks down. Such abnormal worlds are precisely where the oddity—but not inconsistency—of KP's clash with non-omniscience emerges.[17]

The proposal, then, is to avoid 'possibly true gluts' via expanding one's range of possibilities. Specifically, the verificationist acknowledges epistemically abnormal possibilities in which K is deviant. At the same time, the verificationist is committed, on the whole, to the validity of standard K-rules. While knowledge might deviate from its normal behaviour at odd points, the validity of standard K-rules ought to remain intact. Accordingly, in addition to expanding her range of possibilities, the verificationist narrows her account of validity—or, what comes to much the same, keeps her account of validity focused on the non-deviant, normal possibilities.

A formal—and, in some respects, familiar—picture will be helpful. I will return to philosophical discussion in Section 6.3.

6.2. A formal picture

The basic idea can be modelled along 'non-normal lines'.[18] We make a distinction among worlds—the normal and non-normal (or abnormal, as I will say). In turn, we define validity as 'truth-preservation' over only one sort of world, not as 'all points (worlds) validity'. The behaviour of target operators at the abnormal worlds is recognized, but such behaviour is (in effect) ignored for purposes of defining validity. A simple account is as follows.

Let $\mathcal{V}$ and $\mathcal{D}$ be as in Section 6 (Strong Kleene base). Our interpretations are structures $\langle \mathcal{W}, \mathcal{N}, \mathcal{N}^*, R, E, v, \varepsilon \rangle$, where $\mathcal{W} = \mathcal{N} \cup \mathcal{N}^*$, with $\mathcal{N}$ (normal worlds) and $\mathcal{N}^*$ (abnormal) non-empty, and $\mathcal{N} \cap \mathcal{N}^* = \emptyset$. R and E are as before, each being at least reflexive on $\mathcal{W}$. ε, to which I'll return, has the job

[16] Compare Priest (1992).

[17] Admittedly, if one acknowledges 'abnormal epistemic worlds' in which, e.g., K's normal distributive behaviour breaks down, there may be no strong reason to reject other such abnormal worlds in which more radical deviance occurs (such as knowing a contradiction!). Even so, the current proposal aims at avoiding 'possibly true inconsistency' altogether.

[18] The idea behind 'non-normal semantics' comes from Kripke (1965), wherein the aim was to model Lewis systems weaker than $S4$. Arguably more significant philosophical use of non-normal semantics has emerged in literature on 'relevant logics'. See Dunn and Restall (2002) and references therein.

of evaluating K-claims at abnormal worlds. $v : \mathcal{A} \times \mathcal{W} \longrightarrow \mathcal{V}$ assigns values to all atomics at all worlds, normal and not. Interpretations are extended to all sentences at all worlds via the following clauses.

1. Extensional. For any $w \in \mathcal{W}$,

$$v_w(\neg\alpha) = 1 - v_w(\alpha)$$
$$v_w(\alpha \wedge \beta) = \min\{v_w(\alpha), v_w(\beta)\}$$
$$v_w(\alpha \vee \beta) = \max\{v_w(\alpha), v_w(\beta)\}$$

2. Possibility. For any $w \in \mathcal{W}$,

$$v_w(\Diamond\alpha) = \max\{v_{w'}(\alpha) : wRw' \text{ for any } w' \in \mathcal{W}\}$$

3. Knowledge.
   (a) Normal worlds. For any $w \in \mathcal{N}$

   $$v_w(K\alpha) = \min\{v_{w'}(\alpha) : wEw' \text{ for any } w' \in \mathcal{W}\}$$

   (b) Abnormal worlds. For any $w \in \mathcal{N}^*$

   $$v_w(K\alpha) = \varepsilon_w(K\alpha)$$

The job of ε, as above, is to give values to K-claims at our abnormal worlds. ε may be viewed as an 'arbitrary evaluator' of K-claims at abnormal worlds, though the arbitrariness, to avoid inconsistency at abnormal worlds, is subject to the following constraint.

$$\varepsilon_w(K\alpha) = 1 \Rightarrow v_w(\alpha) = 1$$

Finally, validity is defined as 'truth-preservation' over all normal worlds of all interpretations.

So given, the semantics delivers some, but not all, of the target principles. Importantly, we do not get KP. For example, consider an interpretation in which $\mathcal{N} = \{w\}, \mathcal{N}^* = \{w^*\}$, and, in addition to reflexivity, we have only wEw^*. Now let $v_w(\alpha) = 1$ and $v_{w^*}(\alpha) = 0 = \varepsilon_{w^*}(K\alpha)$. Figure 8.1 shows diagram of the counter-example.

R-accessibility

R	w	w^*
w	√	
w^*		√

E-accessibility

E	w	w^*
w	√	√
w^*		√

Values at worlds

	α	$K\alpha$	$\Diamond K\alpha$
w	1	0	0
w^*	0	0	0

Figure 8.1.

The trouble, of course, is that our class of interpretations is too big.

Towards narrowing our class of interpretations, let a narcissistic world—an n-world, for short—be any world w (normal or abnormal) such that, for any $w' \in \mathcal{W}$,

$$wRw' \text{ or} wEw' \Rightarrow w = w'$$

N-worlds see only themselves, in either relevant sense of 'see'. Now, define a V*-model to be any interpretation (as above) such that the following holds.

V*. For any normal w, if $v_w(\alpha) = 1$, then there is some abnormal n-world w^* such that $\varepsilon_{w^*}(K\alpha) = 1$ and wRw^*.

In turn, validity is defined as 'truth-preservation' over all normal worlds of all V*-models. That there are V*-models may be seen by tweaking the previous counter-example to get the results shown in Figure 8.2 (where 'starred' worlds are abnormal).[19]

R-accessibility

R	w	w_1^*	w_2^*	w_3^*
w	√	√		√
w_1^*		√		
w_2^*			√	
w_3^*				√

E-accessibility

E	w	w_1^*	w_2^*	w_3^*
w	√		√	
w_1^*		√		
w_2^*			√	
w_3^*				√

Values at worlds

	α	$K\alpha$	$K(\alpha\wedge\neg K\alpha)$	$\Diamond K\alpha$
w	1	0	0	1
w_1^*	1	0	1	0
w_2^*	0	.5	.5	.5
w_3^*	1	1	.5	.5

Figure 8.2.

The corresponding picture is shown in Figure 8.3.[20]

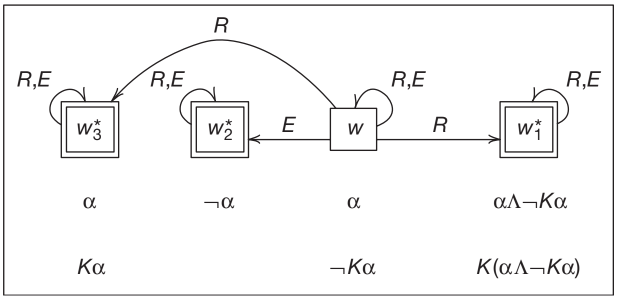

Figure 8.3.

[19] .5 is not forced at w_2^*. One could also give $K\alpha$ the value 0.
[20] In general, a doubly squared world is abnormal.

Notice that each of the abnormal worlds except for w_2^* serves as an n-world for w. The V*-model above also serves to invalidate Fitch's chief inference—from knowability to known. In the model, w is a non-omniscient (normal) world with respect to α, but—thanks to the abnormal worlds—it is possible to know α.

6.3. Comments

I turn to a few comments about the 'abnormal' approach. I begin with a few salient virtues of the semantics, and then briefly turn to the broader, philosophical picture (returning to the topic in Section 8).

As expected, Contraposition and LNC are invalidated, and the regular extensional connective remain normal (as in Strong Kleene). More importantly, the semantics validate each of the standard K-rules, including the essential KP.

KV. Let $v_w(K\alpha) = 1$ for some $w \in \mathcal{N}$. Then $v_{w'}(\alpha) = 1$ for all $w' \in \mathcal{W}$ such that wEw'. Since E is reflexive, $v_w(\alpha) = 1$.

KC. Let $v_w(K(\alpha \wedge \beta)) = 1$ for some $w \in \mathcal{N}$. Then $v_{w'}(\alpha \wedge \beta) = 1 = v_{w'}(\alpha) = 1 = v_{w'}(\beta)$ for all $w' \in \mathcal{W}$ such that wEw'. Hence, as E is reflexive, $v_w(K\alpha) = 1 = v_w(K\beta)$.[21]

KP. Let $v_w(\alpha) = 1$. Then, by V*, there's some abnormal n-world w^* such that wRw^* and $v_{w^*}(K\alpha) = \varepsilon_{w^*}(K\alpha) = 1$.

On the other hand, not everything is retained. Not surprisingly, the deviation from 'all points validity' to 'all normal points' invalidates certain inferences, notably, CP (see Section 2). For example, as above, KC is valid, and so $K(\alpha \wedge \neg K\alpha)$ implies $K\alpha \wedge K\neg K\alpha$. Moreover, KP is valid (as above). Yet, $\Diamond K(\alpha \wedge \neg K\alpha) \nvdash \Diamond(K\alpha \wedge K\neg K\alpha)$, since the relevant world—the world at which $K(\alpha \wedge \neg K\alpha)$ is true—might be abnormal.[22] In abnormal worlds, K can deviate from its normal behaviour. One might think of such worlds not only as 'odd epistemic possibilities' but, further, as worlds at which valid K-behaviour breaks down.

Turning to the broader philosophical picture, a few virtues of the current account may be noted.

1. Consistency. A main motivation behind the 'abnormal' approach was to avoid even the possibility of 'true contradictions'. While achieving as much requires constraints on ϵ, the aim seems to be realized, for what that is worth.

[21] But see further discussion below!

[22] This is not surprising given non-normal semantics. Indeed, as mentioned, Kripke's original motivation beyond non-normal worlds semantics was to model Lewis systems weaker than *S*4, systems in which Necessitation fails. Moreover, in subsequent non-normal approaches to conditionals, the aim is often to model conditionals for which there is no (standard) deduction theorem.

2. Fitch's Lesson. Fitch's 'proof', as in Section 3, points to a genuine oddity in combining KP and non-omniscience. In the current 'abnormal' case, the oddity is fully acknowledged; it shows up as 'abnormal possibilities', possibilities that seem inconsistent but, in the end, avoid outright inconsistency via deviant K behaviour.
3. Failure of CP. While CP's failure is odd, the current story comes with an explanation: distribution of K fails inside (aletheic) modal contexts because such contexts are pointing to epistemically deviant worlds.

By my lights, there is a coherent story along the 'abnormal' lines—odd, but coherent. Some oddness, as Fitch highlighted, is inevitable, at least given KP and non-omniscience.

The question, of course, is whether the 'abnormal' approach to the inevitable oddness is overly odd. Ultimately, that is an issue for verificationists. As far as I can see, there is nothing in verificationism that either rules out or implausibly conflicts with (something like) the foregoing 'abnormal' approach.

Whether, in the end, the abnormal approach is ultimately viable is something that I leave for debate. Doing away with even the possibility of 'true contradictions' is difficult. Perhaps, ultimately, verificationists are better off accepting Fitch's argument for apparently possible 'true contradictions' in a broader paraconsistent (but non-dialetheic) framework. I turn now to a brief sketch of such an approach. In Section 8 I (very briefly) return to the overall philosophical viability of the canvassed approaches.

7. Synthesis: Gaps and Merely Possible Gluts

If, as in Section 5, one goes with 'all points validity' in a (normal) paracomplete but non-paraconsistent framework, the verificationist seems to be stuck with the trivial world—and a vacuous KP. Dropping 'all points validity', as in Section 6, affords more options, but one is forced to give up a few more rules (e.g., KC in the context of aletheic modalities). While each option may hold promise (especially the second), a further option is worth noting.

In this section, I briefly sketch—without arguing for—another option: an 'all points validity' approach that is both paracomplete and paraconsistent but nonetheless non-dialetheic.

The paraconsistent verificationist blocks Fitch's 'proof' at the same place(s) that K_3 does—either LNC or Contraposition. With respect to the 'oddity' of KP in a non-omniscient world, the paraconsistent response is straightforward: non-omniscience and KP generate an inconsistent possibility—knowing an unknown truth.

But such possible inconsistency needn't generate actual inconsistency, at least in a paracomplete paraconsistent framework. In a paracomplete framework, the paraconsistent verificationist may acknowledge 'possible gluts' without thereby accepting dialetheism—the view according to which there is actual inconsistency.[23] Here, I sketch a basic four-valued framework for verificationists, an extension of the familiar Anderson–Belnap framework (1992).

7.1. The basic model

Interpretations are structures $\langle \mathcal{W}, R, E, v\rangle$, where $\mathcal{W}$, R, and E are as before (with R and E at least reflexive). Here, it is convenient to let $\mathcal{V}$, our semantic values, be $\mathcal{P}(\{1, 0\})$.[24] Then v is any function from $\mathcal{S} \times \mathcal{W}$ into $\mathcal{V}$ subject to the following constraints.

1. Negation
   (a) $1 \in v_w(\neg\alpha)$ iff $0 \in v_w(\alpha)$
   (b) $0 \in v_w(\neg\alpha)$ iff $1 \in v_w(\alpha)$
2. Conjunction
   (a) $1 \in v_w(\alpha \wedge \beta)$ iff $1 \in v_w(\alpha)$ and $1 \in v_w(\beta)$
   (b) $0 \in v_w(\alpha \wedge \beta)$ iff $0 \in v_w(\alpha)$ or $0 \in v_w(\beta)$
3. Disjunction
   (a) $1 \in v_w(\alpha \vee \beta)$ iff $1 \in v_w(\alpha)$ or $1 \in v_w(\beta)$
   (b) $0 \in v_w(\alpha \vee \beta)$ iff $0 \in v_w(\alpha)$ and $0 \in v_w(\beta)$
4. Possibility
   (a) $1 \in v_w(\Diamond\alpha)$ iff $1 \in v_{w'}(\alpha)$ for some w' such that wRw'.
   (b) $0 \in v_w(\Diamond\alpha)$ iff $0 \in v_{w'}(\alpha)$ for all w' such that wRw'.
5. Knowledge
   (a) $1 \in v_w(K\alpha)$ iff $1 \in v_{w'}(\alpha)$ for all w' such that wEw'.
   (b) $0 \in v_w(K\alpha)$ iff $0 \in v_{w'}(\alpha)$ for some w' such that wEw'.

Validity is defined as 'truth preservation' over all worlds of all interpretations.

With the expected exception of Contraposition and LNC, the semantics, with validity so defined, preserve most of the target principles: KV, KC, CP, Adjunction, etc. The question, of course, concerns KP.

[23] In Priest's alternative *LP* (dialetheic) setting, which is not paracomplete, the 'mere possibility' of 'true contradictions' immediately generates actual inconsistency. In *LP* (or the target extension), we have $\vdash \neg\Diamond(\alpha \wedge \neg\alpha)$. Hence, given any β such that $\Diamond(\beta \wedge \neg\beta)$ is actually true, we immediately have actual inconsistency. For further discussion, see Restall (1997), wherein Restall first discussed
● Q1 the point regarding the *LP* situation, and Beall and Restall (2006) for broader discussion.●)

[24] This idea is due to Dunn (1966, 1976).

Alas, KP is invalid. Just consider an interpretation in which $\mathcal{W} = \{w, w'\}$ and, in addition to the required reflexivity of R and E, we have wRw' and wEw', but also $v_w(\alpha) = \{1\}$ and $v_{w'}(\alpha) = \{0\}$.[25] This serves as a counter-example to KP. The diagram shown in Figure 8.4 may be useful.

R-accessibility

R	w	w'
w	√	√
w'		√

E-accessibility

E	w	w'
w	√	√
w'		√

Values at worlds

	α	$K\alpha$	$\Diamond K\alpha$	$\Diamond K(\alpha \wedge \neg K\alpha)$
w	{1}	{0}	{0}	{0}
w'	{0}	{0}	{0}	{0}

Figure 8.4.

A picture of the counter-example is shown in Figure 8.5. I give only the value of α.

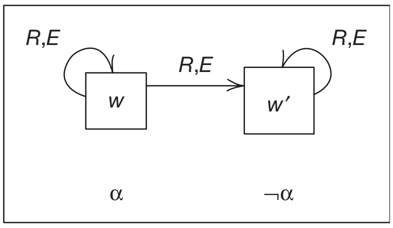

Figure 8.5.

So, KP fails. The trouble, of course, is that our class of interpretations is too big. To get the target interpretations, we need to pare down our class of interpretations.

7.2. The target: V-models

The natural remedy is to invoke n-worlds, as in Section 6 (but now without abnormal worlds). Let an epistemically narcissistic world—an n-world, for short—be a world w such that, for any $w' \in \mathcal{W}$,

$$wEw' \Rightarrow w = w'$$

Since E is reflexive, every world epistemically sees itself; n-worlds (epistemically) see only themselves. In turn, we define a V-model (for Verificationism model) to be any interpretation (as above) that conforms to the following.

[25] For that matter, you could let $v_{w'}(\alpha) = \emptyset$.

V. If $1 \in v_w(\alpha)$ then there's some n-world w' such that wRw' and $1 \in v_{w'}(\alpha \wedge \neg\alpha)$.[26]

Validity is defined as before, but now only over V-models. And with that we get KP.

That there are V-models is clear. In particular, we have V-models that invalidate Fitch's chief inference—from knowability to knowledge. A simple V-model—perhaps the simplest—in which the Fitch inference fails (viz., from knowable to known) is shown in Figure 8.6.

R-accessibility

R	w_0	w_1	w_2
w_0	√	√	
w_1		√	
w_2		√	√

E-accessibility

E	w_0	w_1	w_2
w_0	√		√
w_1		√	
w_2			√

Values at worlds

	α	$K\alpha$	$\Diamond K\alpha$	$\Diamond K(\alpha \wedge \neg K\alpha)$
w_0	{1}	{0}	{1,0}	{1,0}
w_1	{1,0}	{1,0}	{1,0}	{1,0}
w_2	{0}	{0}	{1,0}	{1,0}

Figure 8.6.

A picture of the model is shown in Figure 8.7.

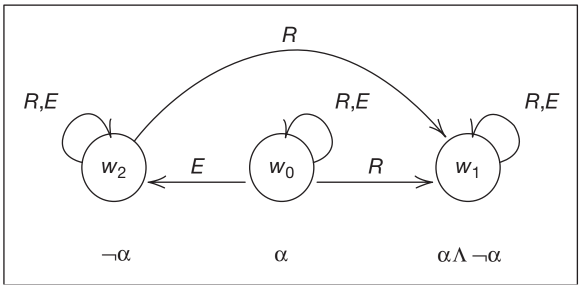

Figure 8.7.

This is a model in which the given Fitch inference fails, since α is knowable at w_0 but not thereby known. The trouble, of course, is that the paracomplete (but paraconsistent) V-models were supposed to afford an entirely consistent actual world while allowing for 'merely possible inconsistency'. In the simple model above, such a promise does not show up. After all, w_1 serves as an n-world for

[26] A simpler, perhaps more natural, route would be to add a distinguished actual world and impose V only on that, but I will go with the 'all worlds of all models' approach.

both w_0 and w_2 (and itself). Since there are only three worlds, the above model, in effect, is basically an *LP*-based model. (If we demanded LEM for all atomics, it would be an *LP*-based model.) The upshot is that, in the above model, the possibility of α-inconsistency—and, in particular, $K\alpha$-inconsistency—trickles back into actual inconsistency: $\Diamond K\alpha$ is true and false at all worlds, and hence the actual.

To get a consistent but non-omniscient 'actual world' (say, w_0), we simply add more worlds. The simplest addition is the null world $w_\emptyset$, shown in Figure 8.8.[27]

R-accessibility

R	w_0	w_1	w_2	$w_\emptyset$
w_0	√	√		√
w_1		√		
w_2		√	√	√
$w_\emptyset$				√

E-accessibility

E	w_0	w_1	w_2	$w_\emptyset$
w_0	√		√	
w_1		√		
w_2			√	
$w_\emptyset$				√

Values at worlds

	α	$K\alpha$	$\Diamond K\alpha$	$\Diamond K(\alpha \wedge \neg K\alpha)$
w_0	{1}	{0}	{1}	{1}
w_1	{1,0}	{1,0}	{1,0}	{1,0}
w_2	{0}	{0}	{1}	{1}
$w_\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$

Figure 8.8.

The corresponding picture can be seen in Figure 8.9.

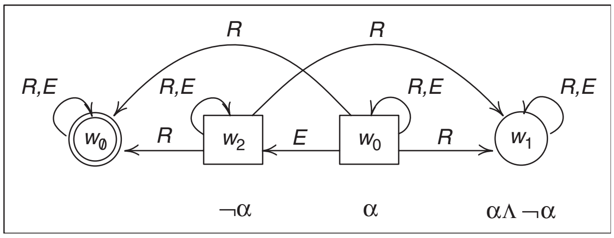

Figure 8.9.

This model is more attractive than the former, simpler model, as it leaves 'the actual world' consistent while nonetheless refuting the Fitch inference. I move to a few general comments.

7.3. Comments

There are various virtues of the V-models over the ***LP***-based paraconsistent approach. For present purposes, I list the salient ones.

[27] Note that the null world is not essential; one merely needs the appropriate 'incompleteness'.

1. Merely Possible Inconsistency. The foregoing approach shares the basic response common to any paraconsistent verificationism: namely, that KP (and the other set of K-rules) forces inconsistency in the face of non-omniscience. But since the current approach is also paracomplete, there's no threat that 'merely possible inconsistency' implies actual inconsistency—as is the case in an ***LP***-based approach.[28] The upshot is that a verificationist can admit that the possibility of knowing unknown truths forces inconsistency; however, it need only force inconsistency 'elsewhere' and only elsewhere—some merely possible world.
2. Trivial world. While the trivial world, without further constraints, certainly shows up in V-models, it isn't required to ensure KP (or, as discussed below, the countermodel to Fitch's basic inference). There are V-models, of course, in which LEM holds among all worlds of the given models (viz., ***LP***-models!); however, being based on a broader four-valued framework, LEM certainly isn't valid. In short, V-models allow for 'incomplete worlds', worlds in which neither α nor $\neg\alpha$ show up (as it were).
3. Not entirely inconsistent. Moreover, while KP is fully ensured, as above, by inconsistency 'elsewhere', V-models allow for 'local inconsistency' to do the work. In particular, the n-worlds, into which knowing 'non-omniscience truths' (e.g., $\alpha \wedge \neg K\alpha$) forces inconsistency, need not themselves be entirely inconsistent. Because of incompleteness, there can be many n-worlds throughout which the given inconsistency is distributed, and many of them can be perfectly consistent in proper quarters.

There are probably other notable virtues vis-à-vis the ***LP***-based (paraconsistent) approach, but I turn to one final matter.

One might think that V, which invokes suitable en-worlds to ensure KP, is ad hoc. Such charges are notoriously difficult to adjudicate, and I won't pursue the issue in any depth here. By my lights, V is not at all ad hoc. After all, V reflects the (paraconsistent and paracomplete) verificationist's chief tenet: that all truths are knowable—even those that reflect non-omniscience, and hence generate inconsistency elsewhere. Rather than being some ad hoc posit, the relevant en-worlds that V invokes might best be seen as an implicit feature of verificationism.

One thing is uncontroversial about Fitch's argument: verificationism's commitment to KP makes for some oddity in its confrontation with our actual non-omniscience. The paraconsistent-cum-paracomplete framework accepts that the given oddity is indeed as it appears: possibly true inconsistency. But verificationists are not thereby dialetheists; such inconsistency, in virtue of incompleteness, is harnessed at the merely possible.

[28] Of course, enriching the language might raise further problems, but the aim here is merely to sketch a beginning option.

8. Closing Remarks

I would like to end this paper by arguing for the supremacy of one of the canvassed options, but I cannot. As above, I am not a verificationist, and so not committed to KP via a prior theory of truth (or meaning, or etc.). Moreover, I know of no good arguments for KP.[29] Still, I find KP plausible and think that each of the canvassed approaches has merit. Instead of trying to settle which, if any, of the given approaches is best, I will close by answering the most salient worries that confront each of the two chief options—setting aside the 'flat-footed response' (see Section 6).

8.1. Abnormal epistemic possibilities?

The suggestion, here, is that KP holds in virtue of abnormal epistemic possibilities, where these are possibilities in which normal K behaviour breaks down. The chief worry about such a picture is that we are no longer talking about knowledge when we are talking about 'abnormal $\boldsymbol{K}$ behaviour'. Put differently (with echoes of Quine), the charge is that necessarily, $\boldsymbol{K}$ behaves like such and so—in particular, distributes over conjunction (and is such that, e.g., CP is valid). Hence, the 'abnormal epistemic possibility' framework is really one in which we are introducing two distinct epistemic operators, one reflecting our 'real knowledge operator/predicate', the other some 'deviant' (but distinct) operator/predicate. As such, verificationists—and KP theorists, in general—are still stuck with the original problems confronting our 'real' item.[30]

By way of reply, the way I look at the situation is (briefly) as follows. Verificationists are committed to some sort of oddity. If verificationists are likewise committed to the bulk of the given rules (see Section 3) and 'no possibly true inconsistency', then a natural suggestion, as in Section 6, is that $\boldsymbol{K}$ behaves differently at different sorts of worlds. Now, the charge, as above, maintains that we have two different $\boldsymbol{K}$s, rather than a single $\boldsymbol{K}$ that, as said, behaves differently at different points. I'm not sure how to adjudicate this. If the project is to give the verificationist—or KP theorist, in general—entirely consistent worlds across the board, while also retaining the bulk of the given rules, then it's unlikely that there's a distinct $\boldsymbol{K}$ of the sort presupposed in the charge (as opposed to a single $\boldsymbol{K}$ that behaves differently at different points, as per the proposal). After all, if 'the real $\boldsymbol{K}$' is like that (e.g., supports distribution inside the diamond), then the verificationist is stuck with inconsistent worlds.

Again, I don't know how to ultimately adjudicate the matter. I'm not a verificationist, but I think it worthwhile to see how the verificationist might

[29] I do not consider appeals to 'intuition' good arguments.

[30] I am grateful to Carrie Jenkins for pushing this point.

enjoy entirely consistent worlds and the bulk of the rules. Of course, she has to give up something, and the Section 6 proposal gives up 'all points validity' (and, in turn, CP). In the end, perhaps the resulting picture is too implausible to suit verificationists. I don't know. But I do not see why they can't have a single K that behaves differently at different points. Indeed, verificationists—or, again, KP theorists, in general—can take the lesson of Fitch's 'proof' to be that we were ignoring various possibilities, namely, ones in which our unique K behaves in very abnormal ways.[31]

8.2. Paraconsistent but non-dialetheic V-models

There may well be various worries about this approach, many of which might spring from general worries about 'possibly true contradictions'. This paper is not the place to address such broad worries.[32] Instead, I will assume a general openness to the idea of (merely) 'possibly true gluts'. There remains a salient worry for the V-model approach.[33]

The worry, in short, is that the proposal calls for too much. In particular, the proposal commits us to the possibility of $\alpha \wedge \neg\alpha$ for every true α. Even if one is prepared to acknowledge merely possible gluts, it is hard to accept that for every truth α, it is possible that α is true and false!

I think that, by way of reply, one needn't quite accept as much as the V-model approach yields. The V-models were so given as a simple example, but one should be able to restrict matters further so as to avoid the going worry. For example, one approach might be to restrict condition V to any 'non-omniscience truth', any truth of the form $\alpha \wedge \neg K\alpha$.[34] Whether this would immediately yield KP is not obvious, but it would at least deal with the main worry over KP—namely, the sort of 'non-omniscience' claims involved in Fitch's argument.

[31] One might also argue that the verificationist—or KP theorist, in general—ought to acknowledge possibilities in which the constraints on K (on knowledge, in general) vary. In the case of verificationism, it is not implausible to think that knowledge might be achieved in some (admittedly, abnormal or remote) possibilities in which verification criteria are weaker than normal. I think that this line is worth exploring, but for space reasons I omit further discussion.

[32] For discussion of such broader issues, see Priest, Beall, and Armour-Garb (2004).

[33] I am grateful to Greg Restall for pushing this concern.

[34] In this case, it might be easier to add a distinguished 'actual world' to the models, and define validity as 'truth-preservation' over actual worlds (of all such models), but I will leave details for another occasion.

Queries in Chapter 8

Q1. opening paranthesis is missing here. Please check.