TWO DIAMONDS ARE MORE THAN ONE[*]

Transitivity and the Factivity of Feasible Knowability

Elia Zardini
Northern Institute of Philosophy, Department of Philosophy, University of Aberdeen
elia.zardini@abdn.ac.uk

Abstract After introducing semantic anti-realism and the paradox of knowability, the paper considers a new threat posed to a restricted version of semantic anti-realism by a group of recent arguments. These arguments purport to achieve with respect to a restricted version of anti-realism what the paradox of knowability achieves with respect to an unrestricted version. Building on features of the informal notion of feasible knowability which occurs in the formulation of the anti-realist's thesis—in particular, on its factivity—the paper shows how the accessibility relation for the relevant modality must be relativized to particular facts, and how this relativization can lead to peculiar failures of transitivity. The paper also argues—again, on the basis of the factivity of feasible knowability—that the relevant modality cannot be identified with (non-transitive) epistemic possibility. Together, these results suffice to show the unsoundness of the arguments considered, and contribute to an appreciation of the distinctive semantics and logic underlying the anti-realist's notion of feasible knowability.

Keywords: anti-realism, epistemic possibility, modal epistemic logic, paradox of knowability, transitivity of accessibility

Introduction and Overview

Semantic anti-realism (henceforth simply 'anti-realism')[1] may neutrally be characterized as the doctrine that *there is a conceptual connection between truth and our recognition of it*. As qualifiedly applied to a particular discourse D, anti-realism is the doctrine that there is a conceptual connection between

[*]Earlier versions of the material in this paper were given in 2005 at the Arché Modality Seminar (University of St Andrews); in 2006, at the International Colloquium on *(Anti-)Realisms, Logic and Metaphysics* (University of Nancy 2); in 2007, at the Arché Epistemology Seminar (University of St Andrews). I would like to thank all these audiences for very stimulating comments and discussions. Special thanks go to Ross Cameron, Roy Cook, Bob Hale, Jonathan Lowe, Agustín Rayo, Sven Rosenkranz and Crispin Wright. In writing this paper, I have benefitted, at different stages, from an AHRC Doctoral Award and from a RIP Jacobsen Fellowship.

the truth of the sentences belonging to D and our recognition of it.[2] This conceptual connection is quite naturally supposed to be captured by the formulation of an *epistemic constraint* on the notion of truth operating over a certain discourse:

(EC$^-$) Necessarily, 'P' is true only if it is feasibly knowable that P,[3]

where 'necessarily' expresses metaphysical necessity.

Two features of the modal epistemic operator 'it is feasibly knowable that', as used by the anti-realist in stating her thesis, are worth remarking upon right at the outset. Firstly, the operator is intended to be *factive*: that is, its being feasibly knowable that P entails that P (but see Tennant, 2009, pp. 230–237 for a recent disagreeing voice). In the presence of the *disquotational* schema:

(D) Necessarily, 'P' is true iff P[4]

(where, again, 'necessarily' expresses metaphysical necessity), (EC$^-$) can therefore be strengthened to:

(EC) Necessarily, 'P' is true iff it is feasibly knowable that P.

Secondly, that the knowability in question is a *feasible* one means that the relevant metaphysically possible situations which verify a claim of feasible knowability are situations:

(i) which concern epistemic subjects endowed with our actual cognitive powers or, at most, with some *finite* extensions thereof;

(ii) in which the available evidence is at least nomologically constrained by the *present* situation of the world (that is, by the situation at which the claim of feasible knowability is to be evaluated).

(i) ensures that (EC) is not made trivial e.g. for arithmetical discourse by the metaphysical possibility of a being endowed with an *infinite* extension of our actual cognitive powers, being which would thus be in a position to survey "at a glance", as it were, the whole number series in its infinity and so in a position to decide every sentence concerning it (just as we are in a position to decide every quantifier-free sentence concerning it). (ii) ensures that (EC) requires e.g. nothing less than nomologically available evidence in order for a past-tensed sentence to be true, and so, presumably, *traces in the present* of the past facts the sentence is about.[5]

Moreover, I henceforth assume that the modality expressed by ‘it is feasibly knowable that’ results from the composition of the epistemic modality of knowledge with a *sui generis* alethic modality, and call the latter ‘*feasible possibility*’ (and its dual ‘*feasible necessity*’). Feasible knowability is then just feasible possibility of knowledge.[6] Since inspection of (i) and (ii) reveals that the possible situations relevant for a claim of feasible possibility result in effect from a *restriction* on the collection of all metaphysically possible situations, the collection of all feasible possibilities will properly be included in the collection of all metaphysical possibilities.

The rest of the paper is organized as follows. Section 1 expounds the problem posed by a simple well-known argument to a certain version of anti-realism, and considers a different version of it. Section 2 presents a novel challenge posed to the modified version, challenge which calls for a deeper understanding of the semantics of the operator ‘it is feasibly possible that’. Section 3 undertakes this task, showing how the attendant accessibility relation for the operator must be relativized to particular facts, and how this relativization can lead to peculiar failures of transitivity. In the lights of this result, section 4 considers (non-transitive) epistemic possibility as candidate for playing the role of feasible possibility, and finds it wanting. Section 5 draws the conclusions which follow for anti-realism and, more generally, for our understanding of the modality expressed by the operator ‘it is feasibly possible that’.

1. The Paradox of Knowability and the Restriction Strategy

In general, *unrestricted* anti-realism is the doctrine that, for whichever schema (such as (EC)) is to constrain the truth of sentences belonging to a discourse, *every* instance of it holds.[7] Under some natural assumptions, *(EC)*-unrestricted anti-realism is refuted by the following simple argument, originally published in Fitch, 1963, pp. 138–139[8] and long known as ‘*the paradox of knowability*’:

P_1	(P_1)	$\mathcal{K}(\varphi \wedge \neg\mathcal{K}\varphi)$	A
P_1	(P_2)	$\mathcal{K}\varphi$	P_1 $\mathbf{IKT}_{\mathcal{K}}$
P_1	(P_3)	$\mathcal{K}\neg\mathcal{K}\varphi$	P_1 $\mathbf{IKT}_{\mathcal{K}}$
P_1	(P_4)	$\neg\mathcal{K}\varphi$	P_3 $\mathbf{IKT}_{\mathcal{K}}$
P_1	(P_5)	$\mathcal{K}\varphi \wedge \neg\mathcal{K}\varphi$	P_2,P_4 $\mathbf{IP}$
	(P_6)	$\neg\mathcal{K}(\varphi \wedge \neg\mathcal{K}\varphi)$	P_1,P_5 $\mathbf{IP}$
	(P_7)	$\Box\neg\mathcal{K}(\varphi \wedge \neg\mathcal{K}\varphi)$	P_6 $\mathbf{IKT}_{\Box}$,

where $\Box$ expresses metaphysical necessity, $\mathcal{K}\varphi$ formalizes ⌜Someone knows at some time that φ⌝, $\mathbf{IP}$ is intuitionist propositional logic, $\mathbf{IKT}_{\Box}$ an intuitionistically acceptable version of a $\mathbf{KT}$ logic for $\Box$[9] and, analogously,

$\mathbf{IKT}_{\mathcal{K}}$ an intuitionistically acceptable version of a **KT** logic for $\mathcal{K}$. Given that metaphysical necessity (of lack of knowledge) entails feasible necessity (of lack of knowledge), metaphysical necessity of lack of knowledge intuitionistically entails the negation of feasible possibility of knowledge.[10] Thus, (EC)-unrestricted anti-realism in conjunction with (P_7) entails that, for every φ, $\varphi \wedge \neg\mathcal{K}\varphi$ is not true. Taking φ to be 'P', disquoting with (D) and generalizing in sentence position, we obtain from this the result that there are no unknown truths.[11] Classically, this result is logically equivalent to the statement that every truth is known. If you accept classical logic and think that it is not the case that every truth is known, you had better reject (EC)-unrestricted anti-realism. If you accept intuitionist logic and think that there are unknown truths, you too had better reject (EC)-unrestricted anti-realism (and, of course, you had better reject it even if you only accept that it is not the case that there are no unknown truths).

Anti-realism is a deep and difficult philosophical doctrine. It would be very strange, to say the least, if it were once and for all refuted by such a simple piece of reasoning as (P_1)–(P_7) in conjunction with (D) and the *epistemic-modesty* principle:

(EM) For some P, P and it is not known that P.

The correct reaction to the paradox of knowability has therefore seemed to be to many commentators that of concluding that what the paradox shows is, at most, that a theorist who endorses (EC)-unrestricted anti-realism has unnecessarily overstated the idea characteristic of anti-realism, and thereby committed only, as it were, a tactical mistake. After all, it should have come as no surprise, given the level of *complexity* induced on a primitive language by the introduction of logical operators, that their interactions with the atomic sentences of the language and with an epistemic operator of some kind or other ends up yielding logically or conceptually unknowable truths. For, in general, an epistemic subject s's coming to be in some epistemic relation E (not necessarily the one of knowledge)[12] to a content (c_0) may very well logically or conceptually entail that E does not hold between s and the content (c_1) that E does not hold between s and c_0. Thus, assuming that E's holding between s and a conjunction logically or conceptually entails E's holding between s and each conjunct, E cannot, on purely logical or conceptual grounds, hold between s and the conjunction of c_0 and c_1, even if such a conjunction may very well be true.

These considerations may seem to point towards a very weak *restriction* on (EC) (along the lines of that favoured in Tennant, 1997, pp. 245–279; 2001; 2010, pp. 11–21 and defended—successfully, in my opinion—in Tennant, 2002 against the stricter restriction to atomic sentences favoured in Dummett, 2001):[13] namely, the restriction that, if 'P' is to be a legitimate *substituens* in the (EC)-schema, it be metaphysically possible to know that P (see the sustained exchange Williamson, 2000b; Tennant, 2001; Williamson, 2009; Tennant, 2010 for further discussion). The gist of anti-realism may seem to be that, if something is true *and it is metaphysically possible to know it*, then we as we actually are (or some finite extension of us as we actually are) can know it in our present evidential situation. Of course, the challenge remains for such and similar restriction strategies to locate, in the arguments given for anti-realism, the (implicit and) principled restriction which makes it the case that these arguments, if successful at all, only establish the favoured restricted version of (EC) rather than unrestricted (EC) itself (cf Williamson, 2000b, pp. 112–113). This challenge is certainly very pressing and I myself would actually favour a different approach on behalf of the anti-realist (as detailed in Zardini, 2010a). Be that as it may, in the rest of this paper I would like to address instead another recent objection raised against restriction strategies.

2. A New Threat of Collapse of Feasible Knowability on Actual Knowledge

Berit Brogaard and Joe Salerno have recently argued that, roughly, even versions of anti-realism relying on *any* kind of restriction on (EC) are threatened with inconsistency with (EM) interpreted as quantifying over the same range to which the restricted version of (EC) is still supposed to apply.[14] No progress would then have been made from the original situation of inconsistency between (EC)-unrestricted anti-realism and (EM) revealed by the paradox of knowability (see Brogaard, Salerno, 2002; 2009, pp. 17–18). Such a conclusion is, I submit, *incredible*: how could it ever be that our notion of feasible knowability is *per se* so strong as to rule out that *any* restricted version of anti-realism applies—even only *de facto*—to a class of sentences for which (EM) also holds?[15] In other words, how could it ever be that our notion of a sentence's being feasibly knowable inevitably collapses (classically) on the notion of that sentence's being known or at least (intuitionistically) on the notion of that sentence's not being unknown?

Any such case is clearly bound to overkill: if the notion of a feasibly knowable sentence which is not known is nothing but a *contradictio in adiecto*, then *no one*—i.e. neither the anti-realist *nor the realist*—can hold that it is not the case that every unknown truth is not feasibly knowable and, *a fortiori*, that there are feasibly knowable truths which are not known (cf Rosenkranz, 2004, p. 70). But that is highly counterintuitive, to say the least. For instance, it seems very plausible that there is a metaphysically possible world w_0 where, for some P, at an earlier time t_0 we have not yet known that P while at a later time t_1 we discover that P: knowledge develops in time. However, assuming that the laws of w_0 are indeterministic, it also seems very plausible that there is a metaphysically possible world w_1 nomologically (accessible from w_0 and) accessing w_0 where every relevant epistemic subject is killed by, say, a natural catastrophe soon after t_0 and before t_1, and where no relevant epistemic subject is to come into existence in the further history of the world. Clearly, w_0 at t_1 constitutes a metaphysically possible situation satisfying the conditions (i) and (ii) for verifying a claim of feasible knowability at the situation constituted by w_1 at t_0. Thus, w_1 would be a world where it is feasibly knowable that P even though it is not known by anyone at any time that P. What Brogaard and Salerno's case, if sound, would show is that, surprisingly enough, we can rule out *a priori* that such a w_1 is a genuine possibility!

Indeed, Brogaard and Salerno's case, if sound, would be a devastating case against what we may dub 'optimistic epistemic modesty' (the claim, usually endorsed both by realists and anti-realists, that there are unknown truths which are nevertheless feasibly knowable), underwriting either a form of "epistemic despair" (for every truth you think you have more reason to believe not to be known than to be feasibly knowable: each such truth would then not be feasibly knowable either) or a form of "epistemic arrogance" (for every truth you think you have more reason to believe to be feasibly knowable than not to be known: each such truth would then not be unknown either, and therefore, classically, would indeed be known).[16]

It is therefore very likely that Brogaard and Salerno's case is flawed. Their overarching case divides in three different independent arguments, all of which are aimed at establishing, in different ways, the result that any kind of restriction on (EC) is inconsistent with (EM) interpreted as quantifying over the same range to which the restricted version of (EC) is still supposed to apply. Maybe sensing the devastating force that their overarching case, if sound, would possess, they cautiously conclude that "[...] the restriction strategies proposed thus far are insufficient to treat the real problem. The paradoxes presented herein turn on the basic logic of $\Diamond$ and the ways in which $\Diamond$ operates on epistemic statements. If a restriction strategy can be vindicated, this will be known only after we have formally analysed the anti-realist's notion of possibility" (Brogaard, Salerno, 2002, p. 149). The rest of this paper is intended to be a first step in that direction. Assuming a broadly possible-worlds semantics for the feasible-possibility operator, I will set forth, in the lights of the informal characterization given at the outset of the paper, some features that any adequate formal semantics for the operator should have. This will already be sufficient to show the unsoundness of two of the arguments offered by Brogaard and Salerno and, more importantly, will provide the beginnings of an appreciation of the distinctive modality expressed by the operator 'it is feasibly possible that'.

3. Transitivity, Factivity, and the Relativity of Accessibility

In the following, $\blacklozenge$ formalizes the feasible-possibility operator (and hence $\blacklozenge\mathcal{K}$ formalizes the feasible-knowability operator), while $\blacksquare$ formalizes its dual. The first argument assumes $\mathcal{K}\varphi$ to be subject to (EC) if φ is, and runs as follows (adapted from Brogaard, Salerno, 2002, pp. 144–145):[17]

T_1	(T_1)	$\varphi \wedge \neg\mathcal{K}\varphi$	A
T_2	(T_2)	$\blacksquare(\mathcal{K}\varphi \equiv \blacklozenge\mathcal{K}\mathcal{K}\varphi)$	(EC), (D) $\mathbf{IKT}_{\blacksquare}$
T_2	(T_3)	$\mathcal{K}\varphi \equiv \blacklozenge\mathcal{K}\mathcal{K}\varphi$	T_2 $\mathbf{IKT}_{\blacksquare}$
T_4	(T_4)	$\Box(\varphi \equiv \blacklozenge\mathcal{K}\varphi)$	(EC), (D) $\mathbf{IKT}_{\Box}$
T_4	(T_5)	$\varphi \equiv \blacklozenge\mathcal{K}\varphi$	T_4 $\mathbf{IKT}_{\Box}$
T_1	(T_6)	$\neg\mathcal{K}\varphi$	T_1 $\mathbf{IP}$
T_1,T_2	(T_7)	$\neg\blacklozenge\mathcal{K}\mathcal{K}\varphi$	T_3,T_6 $\mathbf{IP}$
T_1	(T_8)	φ	T_1 $\mathbf{IP}$
T_1,T_4	(T_9)	$\blacklozenge\mathcal{K}\varphi$	T_5,T_8 $\mathbf{IP}$
T_2	(T_{10})	$\blacklozenge\mathcal{K}\varphi \equiv \blacklozenge\blacklozenge\mathcal{K}\mathcal{K}\varphi$	T_2 $\mathbf{IKT}_{\blacksquare}$
T_1,T_2,T_4	(T_{11})	$\blacklozenge\blacklozenge\mathcal{K}\mathcal{K}\varphi$	T_{10},T_9 $\mathbf{IP}$
T_1,T_2,T_4	(T_{12})	$\blacklozenge\mathcal{K}\mathcal{K}\varphi$	T_{11} $\mathbf{IKT4}_{\blacksquare}$
T_1,T_2,T_4	(T_{13})	$\blacklozenge\mathcal{K}\mathcal{K}\varphi \wedge \neg\blacklozenge\mathcal{K}\mathcal{K}\varphi$	T_{12},T_7 $\mathbf{IP}$
T_2,T_4	(T_{14})	$\neg(\varphi \wedge \neg\mathcal{K}\varphi)$	T_1,T_{13} $\mathbf{IP}$,

where φ is an arbitrary sentence falling under the scope of the restricted version of (EC), $\mathbf{IKT}_{\blacksquare}$ an intuitionistically acceptable version of a $\mathbf{KT}$ logic for $\blacksquare$ and $\mathbf{IKT4}_{\blacksquare}$ an intuitionistically acceptable version of a $\mathbf{KT4}$ logic for $\blacksquare$. (T_{14}) is inconsistent with (EM)'s holding over the range to which the restricted version of (EC) is still supposed to apply.

The operator 'necessarily' occurring in both (EC) and (D) expresses metaphysical rather than feasible necessity. Nevertheless, (EC) and (D) jointly justify (T_2) since, as already noted, feasible necessities are (properly) included in metaphysical necessities. Furthermore, the closure step at line (T_{10}) seems to be justified as well.[18] The problem with (T_1)–(T_{14}) may then reasonably be thought to lie in the step from (T_{11}) to (T_{12}), which makes use of the $\blacklozenge$-version of the characteristic $\mathbf{KT4}$-axiom ($\blacksquare\varphi \supset \blacksquare\blacksquare\varphi$), the *transitivity* of the relevant accessibility relation being the necessary and sufficient condition for its semantic justification. The question then is whether the correct semantics for $\blacklozenge$ should employ a transitive accessibility relation. Brogaard and Salerno too seem to identify this as the *crux* of (T_1)–(T_{14}): "One might object that we simply need to treat $\Diamond$ as non-transitive. But this has not yet been argued for. And it would not be very interesting simply to suppose the non-transitivity of $\Diamond$, having no reason other than the threat of the revised Fitch paradox to motivate the supposition. Pending further discussion, the supposition of non-transitivity is ad hoc" (Brogaard, Salerno, 2002, p. 145).

I accept Brogaard and Salerno's request for a *principled* rejection of transitivity. However, a moment's reflection does suffice to show that the intuitive notion of feasible knowability, as spelled out at the outset of this paper, determines that the semantics of ⧫ cannot allow for an unrestrictedly transitive accessibility relation. We start by observing that the *factivity* constraint determines that the only worlds[19] which are relevant in the evaluation of a claim of feasible knowability (that is, the worlds which are *feasibly accessible*) are those worlds in which we hold *everything constant* (with respect to the world at which the claim of feasible knowability is to be evaluated) *except for the relevant epistemic facts* (and their presuppositions and consequences).[20] This is meant to ensure the factivity of the complex modal epistemic operator ⧫$\mathcal{K}$, and thereby to ensure its adequacy as a formalization of the operator 'it is feasibly knowable that'. For what may be regarded as merely possible and non-actual[21] in a claim that it is feasibly knowable that P is *not* the fact that P (which is, on the contrary, assumed actually to obtain), but simply the fact that *it is known that* P. What a claim that it is feasibly knowable that P entails is that it is actually the case that P and that, although it may not actually be known that P, the present state of the actual world determines an evidential situation such that a being endowed with our actual cognitive powers—or with some finite extensions thereof—can—under the constraints of the actual laws of nature—come to know that P.

With such restrictions on feasible accessibility in place, principled counterexamples to its unrestricted transitivity are forthcoming. Consider, for instance, worlds w_0, w_1 and w_2. In w_0, there are 1,963 houses in Nancy, but it is not known that there are 1,963 houses in Nancy. w_1 holds constant every fact of w_0 with the exception that in w_1 it is known that there are 1,963 houses in Nancy. w_2 holds constant every fact of w_1 possibly with the exception that in w_2 the relevant epistemic subjects reflect on their cognitive achievements and thereby come to know that it is known that there are 1,963 houses in Nancy. As the counterexample we are constructing is neutral with respect to the validity of the so-called "$\mathcal{KK}$-*thesis*":

(KK) If it is known that P, then it is known that it is known that P,

w_2 is not assumed to be distinct from w_1 (if the reader thinks that (KK) does not hold, then she is free to assume that w_2 is distinct from w_1; if, on the contrary, she thinks that (KK) does hold, then she is free to assume that w_2 is identical with w_1). Assume also that w_0, w_1 and w_2 all satisfy (i) and (ii) with respect to one another. *Relative to* there being 1,963 houses in Nancy, w_1 is thus feasibly accessible from w_0, and *relative to* its being known that there are 1,963 houses in Nancy, w_2 is feasibly accessible from w_1. Assigning 'It is known that there are 1,963 houses in Nancy' to 'φ', $\mathcal{K}\varphi$ is true at w_2, from which it follows that $\blacklozenge\mathcal{K}\varphi$ is true at w_1 and that $\blacklozenge\blacklozenge\mathcal{K}\varphi$ is true at w_0 (see fn 24 for a justification of this very last claim). Does it also follow that $\blacklozenge\mathcal{K}\varphi$ is true at w_0? Emphatically no, because that would entail, given the factivity constraint we placed on $\blacklozenge\mathcal{K}$, that φ is true at w_0, which is false, as, in w_0, it is not known that there are 1,963 houses in Nancy!

The lesson to be learnt from this kind of example can be elaborated as follows. A world w_j is feasibly accessible or not from a world w_i only *relative to* a particular fact candidate for being feasibly knowable at w_i. In particular, the factivity constraint determines that w_j is feasibly accessible from w_i relative to a particular fact *only if* that fact obtains both at w_i and w_j (call the facts which meet this condition $\ulcorner w_i$-w_j-constant$\urcorner$). $\blacklozenge$ in $\blacklozenge\mathcal{K}\varphi$ can then be interpreted as selecting as relevant feasible-accessibility relation the specific one relative to the fact expressed by φ. Thus, for instance, in our counterexample w_2 is feasibly accessible from w_0 relative to the fact that there 1,963 houses in Nancy, as this fact is indeed w_0-w_2-constant, but it is not feasibly accessible from w_0 relative to the fact that it is known that there 1,963 houses in Nancy, as this latter fact is not w_0-w_2-constant. Consequently, whilst w_2 can count as a witness for the feasible possibility, in w_0, that it is known that there 1,963 houses in Nancy, it cannot count as a witness for the feasible possibility, in w_0, that it is known that it is known that there 1,963 houses in Nancy. Indeed, since for no w_j is the fact that it is known that there 1,963 houses in Nancy w_0-w_j-constant (for that fact does not obtain at w_0 in the first place!), no world can count as such a witness and so, at w_0, it is not the case that it is feasibly possible to know that it is known that that there 1,963 houses in Nancy, as desired.

The facts which are not w_i-w_j-constant will include epistemic facts of the relevant order n, which obviously have to be allowed to change from not obtaining at w_i to obtaining at w_j if the notion of feasible possibility (and, in particular, of feasible knowability) is to have any interest at all and not to collapse on that of actuality (and, in particular, of actual knowledge). However, the very same nth-order epistemic facts which so vary from w_i to w_j may very well, with respect to a world w_k which meets the other conditions for being feasibly accessible from w_j relative to any nth-order epistemic fact, be w_j-w_k-constant.[22] Importantly, any such w_j-w_k-constant fact will not be w_i-w_k-constant. Thus, for instance, in our counterexample the first-order epistemic fact that it is known that there 1,963 houses in Nancy is not w_0-w_1-constant, and yet is w_1-w_2-constant.[23] Importantly, such w_1-w_2-constant fact is not w_0-w_2-constant. It is then clear that, under such circumstances, one cannot conclude without further qualification, from the holding of these two *relativized* feasible-accessibility relations between w_j and w_i and between w_k and w_j, that w_k is also feasibly accessible from w_i relative to *every* fact whatsoever. For, whereas w_k may very well be assumed to be, in turn, feasibly accessible from w_i *relative to* every w_i-w_j-constant fact, it cannot be assumed to be feasibly accessible from w_i *relative to* every w_j-w_k-constant fact, for some such fact (namely, the holding of some epistemic relation to some w_i-w_j-constant fact) may not be w_i-w_k-constant.

Crucially, this peculiar kind of failure of transitivity suffices to determine failures of the characteristic $\mathbf{KT4}$-axiom. For collapse at a world w_i of $\blacklozenge\blacklozenge\varphi$ on $\blacklozenge\varphi$ (where $\blacklozenge\blacklozenge\varphi$ is true at w_i because $\blacklozenge\varphi$ is true at w_j[24] and this in turn because φ is true at w_k) is ensured to be truth preserving only if it is not the case that φ describes the holding of an epistemic relation to a fact which, while w_j-w_k-constant, is not w_i-w_j-constant. As the structure of our counterexample highlights, *only if* this requirement is met can the operator $\blacklozenge\mathcal{K}$ be legitimately taken to be factive. Thus, under this requirement, one cannot, in the context of (T_1)–(T_{14}), collapse the sentence $\blacklozenge\blacklozenge\mathcal{K}\mathcal{K}\varphi$ on $\blacklozenge\mathcal{K}\mathcal{K}\varphi$ (as at line (T_{12})), for, in the context of (T_1)–(T_{14}), $\mathcal{K}\mathcal{K}\varphi$ describes exactly the holding of an epistemic relation to a fact (expressed by $\mathcal{K}\varphi$) which, while w_j-w_k-constant, is not w_i-w_j-constant.[25]

4. Epistemic Possibility of Knowledge and Feasible Knowability

I would like to close by considering a proposal made by Brogaard and Salerno for an identification of the modality expressed by $\blacklozenge$ which satisfies the constraints on it unveiled by our foregoing discussion. Since they (correctly, as we have seen in our discussion of (T_1)–(T_{14})) think that a restricted version of (EC) entails that the accessibility relation featuring in the semantics of $\blacklozenge$ is non-transitive, and (incorrectly, as I don't have space to expand on here) think that a restricted version of (EC) is also incompatible, under (EM), with (KK), Brogaard and Salerno eventually cast about for an interpretation of $\blacklozenge$ which, under the rejection of (KK), verifies the non-transitivity of $\blacklozenge$. They rightly note that *epistemic possibility*, defined as "*consistency*" with what is known,[26] behaves non-transitively under the rejection of (KK).

Indeed, it is easily shown that, given the above definition of epistemic possibility, failure of (KK) is both necessary and sufficient for the accessibility relation featuring in the semantics of the epistemic-possibility operator (that is, *epistemic accessibility*) to be non-transitive. As for necessity, observe that, if (KK) holds, then every world w_0 at which $\mathcal{K}\varphi$ is true is a world at which $\mathcal{K}\mathcal{K}\varphi$ (and $\mathcal{K}\mathcal{K}\mathcal{K}\varphi$, and $\mathcal{K}\mathcal{K}\mathcal{K}\mathcal{K}\varphi\ldots$) is true, so that every world w_1 epistemically accessible from w_0 is a world at which (both φ and) $\mathcal{K}\varphi$ (and $\mathcal{K}\mathcal{K}\varphi$, and $\mathcal{K}\mathcal{K}\mathcal{K}\varphi\ldots$) are true, so that every world w_2 epistemically accessible from w_1 is a world at which φ is true (because of the truth at w_1 of $\mathcal{K}\varphi$), and $\mathcal{K}\varphi$ is true (because of the truth at w_1 of $\mathcal{K}\mathcal{K}\varphi$), and $\mathcal{K}\mathcal{K}\varphi$ is true (because of the truth at w_1 of $\mathcal{KKK}\varphi$)... Therefore, the necessary and sufficient condition for the epistemic accessibility of w_1 from w_0 (truth at w_1 of what is known at w_0) is satisfied by every world w_2 which satisfies the necessary and sufficient condition for being epistemically accessible from w_1. It then follows that, if w_1 is epistemically accessible from w_0 and w_2 from w_1, so is w_2 from w_0—in other words, if (KK) holds, epistemic accessibility is transitive. Contraposing, if epistemic accessibility is not transitive, (KK) does not hold.[27] As for sufficiency, observe that, if (KK) does not hold, then the following model invalidates the transitivity of epistemic accessibility (as Brogaard, Salerno, 2002, pp. 146–147, fn 7 correctly point out): $\mathcal{K}\varphi$ and $\neg\mathcal{KK}\varphi$ are true at w_0, φ and $\neg\mathcal{K}\varphi$ are true at w_1, $\neg\varphi$ and $\mathcal{K}\neg\varphi$ are true at w_2 (w_1 is epistemically accessible from w_0, and w_2 is epistemically accessible from w_1, but w_2 is not epistemically accessible from w_0).[28]

However, one may very well wonder whether the contribution made by 'it is feasibly knowable that' to the truth conditions of a sentence is the same as the contribution made by 'it is epistemically possible that it is known that'. Given the above definition of epistemic possibility, if it is not known that it is not known that P, then it is epistemically possible that it is known that P. Since its being feasibly knowable that P entails that it is the case that P, its not being known that it is not known that P should then likewise entail that it is the case that P. But that is clearly incorrect: consider any case where it is false that P, but it is not known that it is not known that P (just because, say, everyone mistakenly believes, on the contrary, that it is known that P). In any such case, it would lead to a straightforward contradiction to require that it be the case that P. Suppose, for instance, that everyone confidently but mistakenly believes that whales are fish and that, as a consequence, everyone also believes that it is known that whales are fish. Assuming a *modicum* of rationality, no one will also believe that it is not known that whales are fish, and hence, assuming that knowledge requires belief, no one will know that it is not known that whales are fish. Of course it still doesn't follow that, in such a situation, whales are fish, even though it is epistemically possible to know that they are. Epistemic possibility of knowledge does not sustain factivity.

These considerations are already sufficient to block the third argument Brogaard and Salerno give against any restricted version of (EC). The argument assumes that the feasible-possibility operator behaves like an epistemic-possibility operator at least in the sense of satisfying:

(EP) If it is known that it is not the case that P, then it is not feasibly possible that P,[29]

and assumes $\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi$ and $\neg\varphi$ to be subject to (EC) if φ is. It runs as follows (adapted from Brogaard, Salerno, 2002, pp. 147–148):

E_1	(E_1)	$\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi$	A
E_2	(E_2)	$\mathcal{K}(\varphi \equiv \blacklozenge\mathcal{K}\varphi)$	$\mathcal{K}$((EC) $\wedge$ (D)) $\mathbf{IKT}_{\Box,\mathcal{K}}$
E_3	(E_3)	$\varphi \equiv \blacklozenge\mathcal{K}\varphi$	A
E_2	(E_4)	$\varphi \equiv \blacklozenge\mathcal{K}\varphi$	E_2 $\mathbf{IKT}_{\mathcal{K}}$
E_5	(E_5)	$\mathcal{K}(\mathcal{K}\neg\varphi \supset \neg\blacklozenge\varphi)$	$\mathcal{K}$(EP)
E_6	(E_6)	$\mathcal{K}\neg\varphi \supset \neg\blacklozenge\varphi$	A
E_1,E_2	(E_7)	$\blacklozenge\mathcal{K}(\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi)$	E_4 $[\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi/\varphi]$, E_1 **IP**
E_8	(E_8)	$\mathcal{K}(\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi)$	A
E_8	(E_9)	$\mathcal{K}\neg\mathcal{K}\varphi$	E_8 $\mathbf{IKT}_{\mathcal{K}}$
E_6,E_8	(E_{10})	$\neg\blacklozenge\mathcal{K}\varphi$	E_6 $[\mathcal{K}\varphi/\varphi]$, E_9 **IP**
E_3,E_6,E_8	(E_{11})	$\neg\varphi$	E_3,E_{10} **IP**
E_8	(E_{12})	$\mathcal{K}\neg\mathcal{K}\neg\varphi$	E_8 $\mathbf{IKT}_{\mathcal{K}}$
E_6,E_8	(E_{13})	$\neg\blacklozenge\mathcal{K}\neg\varphi$	E_6 $[\mathcal{K}\neg\varphi/\varphi]$,$E_{12}$ **IP**
E_3,E_6,E_8	(E_{14})	$\neg\neg\varphi$	E_3 $[\neg\varphi/\varphi]$,E_{13} **IP**
E_3,E_6,E_8	(E_{15})	$\neg\varphi \wedge \neg\neg\varphi$	E_{11},E_{14} **IP**
E_3,E_6	(E_{16})	$\neg\mathcal{K}(\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi)$	E_8,E_{15} **IP**
E_2,E_5	(E_{17})	$\mathcal{K}\neg\mathcal{K}(\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi)$	E_2,E_3,E_5,E_6,E_{16} $\mathbf{IKT}_{\mathcal{K}}$
E_2,E_5	(E_{18})	$\neg\blacklozenge\mathcal{K}(\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi)$	E_6 $[\mathcal{K}(\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi)/\varphi]$,$E_{17}$ **IP**
E_1,E_2,E_5	(E_{19})	$\blacklozenge\mathcal{K}(\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi) \wedge \neg\blacklozenge\mathcal{K}(\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi)$	E_7,E_{18} **IP**
E_2,E_5	(E_{20})	$\neg(\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi)$	E_1,E_{19} **IP**

where φ is an arbitrary sentence falling under the scope of the restricted version of (EC) and $\mathbf{IKT}_{\Box,\mathcal{K}}$ an intuitionistically acceptable combination of $\mathbf{IKT}_{\Box}$ and $\mathbf{IKT}_{\mathcal{K}}$. What (E_1)–(E_{20}) would show is that (even) a restricted (known) version of (EC) is committed to the rejection that neither φ nor $\neg\varphi$ are known. Although this needn't be a straightforward *reductio* of a restricted (known) version of (EC), it is worth noting that the conclusion (E_{20}) intuitionistically entails $\neg\mathcal{K}\varphi \supset \neg\varphi$ and classically entails $\varphi \supset \mathcal{K}\varphi$ (which are the main results of the original paradox of knowability), as the following simple proof demonstrates:

1	(1)	$\neg(\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi)$	A
2	(2)	$\neg\mathcal{K}\varphi$	A
3	(3)	φ	A
4	(4)	$\mathcal{K}\neg\varphi$	A
4	(5)	$\neg\varphi$	4 $\mathbf{IKT}_{\mathcal{K}}$
3,4	(6)	$\varphi \wedge \neg\varphi$	3,5 **IP**
3	(7)	$\neg\mathcal{K}\neg\varphi$	4,6 **IP**
2,3	(8)	$\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi$	2,7 **IP**
1,2,3	(9)	$\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi \wedge \neg(\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi)$	8,1 **IP**
1,2	(10)	$\neg\varphi$	3,9 **IP**
1	(11)	$\neg\mathcal{K}\varphi \supset \neg\varphi$	2,10 **IP**

(contraposition on (11), $\varphi \supset \neg\neg\varphi$ and the distinctively classical $\neg\neg\mathcal{K}\varphi \supset \mathcal{K}\varphi$ would then yield $\varphi \supset \mathcal{K}\varphi$). ($E_1$)–($E_{20}$) would entail that a restricted (known) version of (EC) is no better off than the unrestricted one, and therefore has no point. However, (E_1)–(E_{20}) crucially relies on (known) (EP), and this has been shown not to hold generally by the counterexample just offered.

We then have that the two modal epistemic operators 'it is feasibly knowable that' and 'it is epistemically possible to know that' are very plausibly not even *extensionally* equivalent (and certainly, in any event, not *intensionally* equivalent), for it is very plausibly *actually* the case that some instance of the schema 'It is epistemically possible that it is known that P' is true whereas the corresponding instance of the schema 'It is feasibly knowable that P' is false (and, in any event, it is certainly *possible* that some instance of the former schema is true whereas the corresponding instance of the latter schema is false).[30]

The previous counterexample assumed that consistency with what is known (on some interpretation or other) is the necessary *and sufficient* condition for epistemic possibility. Yet, the sufficiency direction of the consistency condition may reasonably be rejected as being too weak, at least in some context (cf Hacking, 1967, p. 148). Suppose that a computer collects the results of a complex experiment with subatomic particles which conclusively refute a theory T. Badly misinterpreting the results, I utter ⌜It might be that T is true⌝: my utterance is intuitively false, even though, in that situation, no one knows that T is not true. However, the counterexample can easily be modified to accommodate any plausible stronger condition placed on epistemic possibility, as long as its not being the case that P does not require that it is known that it is not known that P—the vast implausibility of such a requirement can be appreciated by noting that it is equivalent to the characteristic $\mathbf{B}$-axiom for the knowledge operator ($\varphi \supset \mathcal{K}\neg\mathcal{K}\neg\varphi$), which is notoriously highly problematic (see Williamson, 2000a, pp. 166–167, 226–227).

Epistemic possibility of knowledge is thus not sufficient for feasible knowability. Nor is it necessary. For it may very well be the case that it is feasibly knowable that P but that it is known that it is not known that P (just because, say, it is known that no one will ever bother to check whether P), and therefore, contraposing on the uncontroversial necessity direction of the consistency condition, it is not epistemically possible that it is known that P. Suppose, for instance, that the number of books in my office at noon 28/10/2006 is 92, and that I come to know at noon 28/10/2006 that, alas, the Big Crunch is going to happen in 10 minutes. Then it is certainly feasibly knowable that the number of books in my office at noon 28/10/2006 is 92 (it would take less than 10 minutes to count them, and we would know that no book has either left or come in since noon), boring as this truth may be. However, it is also known, in the circumstances, that, in actuality, no one has ever counted, is counting or will ever count how many books there are in my office at noon 28/10/2006. Therefore, it is known that it is not known that the number of books in my office at noon 28/10/2006 is 92, and so it is not epistemically possible that it is known that the number of books in my office at noon 28/10/2006 is 92. The previous considerations concerning extensional and intensional non-equivalence still apply to this second kind of counterexample.[31]

Epistemic possibility of knowledge is thus neither necessary nor sufficient for feasible knowability. Therefore, the peculiar failures of transitivity we observed for the latter in the previous section cannot be explained by the non-transitivity (under failure of (KK)) of the former.

5. Conclusion

We have seen how the factivity constraint on feasible knowability is crucial in generating peculiar failures of transitivity for the feasible-accessibility relation, and how such a failure cannot be explained, tempting as that may be, by an analysis of feasible knowability in terms of epistemic possibility of knowledge. I conclude that, quite unsurprisingly, the arguments we have been reviewing against a restricted version of (EC), when carefully scrutinized, prove unsound. They do so because they obliterate the distinctive semantics and logic—induced by the factivity constraint—which govern the modal epistemic operator 'it is feasibly knowable that' and simply take for granted that it is analysable in terms of knowledge and some usual kind of metaphysical (or epistemic) possibility. Although this paper hasn't been concerned with the fine details of such semantics and logic, it has indicated some important overall features that these should have, and has traced their source back to the the strain generated by the opposite requirements of *convergence* with actuality (as far as *what is known* is concerned) and of *divergence* with it (as far as the relevant *state of knowing* is concerned).

Notes

1. Throughout, I use simple quotation marks for standard quotation, corner quotes and display for autonymous quasi-quotation (where non-metalinguistic quantification into such environment has to be understood substitutionally). Formal vocabulary refers to itself.

2. Call the unqualified version of anti-realism, quantifying over every sentence belonging to *whichever* discourse, 'global anti-realism'. Call a qualified version of anti-realism, quantifying just over every sentence belonging to a *particular* discourse D, ⌜local anti-realism with respect to D⌝. The discussion in the paper is intended to be insensitive to this distinction.

3. Throughout, 'P' will be used as a substitutional sentential variable (if free, amounting in effect to a sentence schema). 'φ' will instead be used as a metalinguistic variable over sentences.

4. Possible restrictions on the schema in order to deal with context sensitivity and the semantic paradoxes will be ignored here, as not relevant to the issues under discussion (see Zardini, 2008 for some discussion).

5. I should stress that neither (i) nor (ii) are non-negotiable for a broadly anti-realist perspective. I only report them as conditions typically occurring in anti-realist discussions (as can be variously found e.g. in Dummett, 1975) and they will actually play no role in the rest of my discussion, which focusses instead on the factivity feature. Thanks to an anonymous referee for raising this point.

6. 'Knowledge' and its relatives are in turn henceforth understood as implicitly existentially generalizing over epistemic subjects and times: knowledge by *someone* at *some time*.

7. Note that this distinction between unrestricted and restricted anti-realism is orthogonal to the distinction between global and local anti-realism introduced in fn 2.

8. But most likely due to Alonzo Church, see Salerno, 2009, pp. 34–37.

9. Unfortunately, there is no unique such version. See Bozic, Dosen, 1984; Dosen, 1985 for some plausible proposals.

10. To see that this modal move from the feasible necessity of a negation to the negation of a feasible possibility should count as intuitionistically valid, think of operators of feasible modality as first-order quantifiers over the range of feasibly possible worlds. Then the *modal* move in question amounts to the intuitionistically kosher *first-order* move from $\forall\xi\neg\varphi$ to $\neg\exists\xi\varphi$.

11. Henceforth, I mostly use the word 'truth' merely as a handy short for the clumsy substitutional locution 'thing which is the case'. Thus, for example, 'There are no unknown truths' is understood to mean 'For no P, P and it is unknown that P'.

12. For example, as Mackie, 1980, p. 91 first noted, it seems just as conceptually impossible to be justified in believing that [P and no one is ever justified in believing that P] as it is to know that [P and no one ever knows that P].

13. For the record, Tennant, 2009; 2010, pp. 21–23 proposes a related but interestingly different restriction.

14. Strictly speaking, that requires a qualification, since some not completely trivial assumptions are made about what the restriction is (in the argument of section 3, it is assumed that $\mathcal{K}\varphi$ is subject to (EC) if φ is; in the argument of section 4, it is assumed that $\neg\mathcal{K}\varphi \wedge \neg\mathcal{K}\neg\varphi$ and $\neg\varphi$ are subject to (EC) if φ is). These assumptions are however so minimal that, in what follows, I will mostly leave the qualification implicit.

15. Think especially of the limit-case in which the class of sentences is the singleton of an arbitrarily chosen true but unknown sentence.

16. Of course, relative to the class of sentences for which they believe that (EC) holds good, the mandated attitude for anti-realists would be epistemic arrogance.

17. I should like to stress that the kernel of the argument goes actually back to Williamson, 1992, p. 67.

18. See fn 25 for a more precise statement concerning (T_{10}).

19. Henceforth, I will set aside complications arising from time.

20. This distinction between the relevant epistemic facts with their presuppositions and consequences on the one side and the rest of all other facts on the other side presupposes a mild form of non-holism. Although such issues clearly go beyond the scope of this paper and I won't try to defend this claim here, it is arguable that it is under that presupposition that the concept of feasible knowability finds its usefulness in ordinary and philosophical thinking. Thanks to Jonathan Lowe, Sven Rosenkranz and anonymous referee for discussions on this point.

21. Note that, as has just happened in the text, I sometimes use 'actually' and its relatives not just to refer to the actual world, but, more generally, to refer back to whichever world is the world at which a certain claim is to be evaluated.

22. Notice though that the epistemic facts which do change from not obtaining at w_j to obtaining at w_k may still include the $n + 1$th-order epistemic facts constituted by the holding of epistemic relations to the nth-order epistemic facts in turn constituted by the holding of those epistemic relations which have varied from w_i to w_j but which are held constant from w_j to w_k.

23. Notice though that what does change from not obtaining at w_1 to obtaining at w_2 still includes the second-order epistemic fact that it is known that it is known that there 1,963 houses in Nancy.

24. In our counterexample, we have tacitly assumed that w_1 can count as a witness for the feasible possibility, in w_0, not only that it is known there are 1,963 houses in Nancy, but also that it is feasibly knowable that it is known that there are 1,963 houses in Nancy. The assumption is warranted by the consideration that those two facts would seem to belong to the same order, since they would seem to concern the same level of actual knowledge (namely, the first one) and by the consideration that, if two facts belong to the same order, a world is feasibly accessible from a world relative to one fact iff it is so accessible relative to the other fact.

25. These considerations concerning the relativity to particular facts of the feasible-accessibility relation can be developed in a formal framework whose main characteristic is the use of *multiple* feasible-accessibility relations. The facts to which feasible accessibility is relative can in effect—with some simplification—be merged together in the following sets: the set of non-epistemic facts, the set of first-order epistemic facts, the set of second-order epistemic facts etc. It is then natural to mirror this hierarchy by classifying formulae according to their degree of complexity with respect to occurrences of $\mathcal{K}$, and to assign then a specific feasible-accessibility relation to each degree of complexity. The resulting logics for the feasible-possibility operator are investigated in detail in Zardini, 2010b. Here, I should like to mention that the logics are sufficient to invalidate also the second argument considered by Brogaard and Salerno, which I don't have space to address in this paper. I should also report that, even though the logics are non-normal in the sense that the $\mathbf{K}$-axiom $(\blacksquare(\varphi \supset \psi) \supset (\blacksquare\varphi \supset \blacksquare\psi))$ does not hold unrestrictedly, they still validate the particular instance used at line (T_{10}). I should finally add that, while transitivity fails in the peculiar sense explained in the text (which suffices to determine failures of the characteristic $\mathbf{KT4}$-axiom), this is compatible with each of the specific feasible-accessibility relations still being transitive.

26. Given our previous stipulations about 'knowledge' and its relatives in fn 6, the notion of epistemic possibility defined in the text is neither *group*- nor *time*-relative. Group-relativity is certainly a crucial feature of the notion expressed by ordinary uses of 'might' and its like, but it is inappropriate when trying to analyze a notion that is not group-relative, such as the notion of feasible possibility studied in this paper. That notion may however be time-relative (if condition (ii) is accepted), in which case one should also focus on a time-relative notion of epistemic possibility when discussing the relation between epistemic possibility of knowledge and feasible knowability. The discussion in this section can easily be adapted to the time-relative case.

27. The gloss in the text 'truth at w_1 of what is known at w_0' clarifies the way in which the rather ambiguous consistency condition for epistemic possibility is to be understood: what is known at a world must be *true* at every world epistemically accessible from it (and not just: what is known at a world must *not* be *false* at every world epistemically accessible from it). Note that the weaker interpretation of the consistency condition would belie failure of (KK) as a necessary condition for the non-transitivity of epistemic accessibility. The following (non-classical) model shows how, under the weaker interpretation of the consistency condition, transitivity fails even under (KK) conditions: $\mathcal{K}\varphi$ and $\mathcal{K}\mathcal{K}\varphi$ are true at w_0, neither $\mathcal{K}\varphi$ nor $\neg\mathcal{K}\varphi$ are true at w_1, $\neg\varphi$ is true at w_2 (w_1 is epistemically accessible from w_0, and w_2 is epistemically accessible from w_1, but w_2 is not epistemically accessible from w_0). Of course, if the semantics is classical, the weaker interpretation of the consistency condition collapses on the stronger. I will ignore such niceties in what follows.

28. Notice that, in the model described in the text, w_1 and w_2 also invalidate the *symmetry* of epistemic accessibility, and this independently of the question whether (KK) is valid or not.

29. Notice that (EP) is intuitionistically equivalent with the necessity direction of the consistency condition (substituting 'feasibly' for 'epistemically' in 'If it is epistemically possible that P, then it is not known that it is not the case that P').

30. Of course, for any candidate instance for extensional divergence, we cannot, on purely logical grounds, know that it is a ultimately suitable one (this being just another case of the structural unknowability so nicely illustrated by the Fitch-Church paradox): if it is known that it is not feasibly knowable that P, then it is known that it is not known that P (by the entailment from knowledge to feasible knowability and closure of knowledge), and so it is not known that it is not known that it is not known that P (by factivity of knowledge)—i.e. it is not known that it is epistemically possible to know that P. Hence the suppositional character of the counterexample against the sufficiency of epistemic possibility of knowledge for feasible knowability.

31. Of course, again, for any candidate instance for extensional divergence, we cannot, on purely logical grounds, know that it is a ultimately suitable one (this being yet just another case of the structural unknowability so nicely illustrated by the Fitch-Church paradox): if it is known that it is feasibly knowable that P, then it is known that P (by factivity of feasible knowability and closure of knowledge), and so it is not known that it is not known that P (by factivity of knowledge), and so it is not known that it is known that it is not known that P (by factivity of knowledge)—i.e. it is not known that it is not epistemically possible to know that P. Hence, again, the suppositional character of the counterexample against the necessity of epistemic possibility of knowledge for feasible knowability.

References

Bozic M., Dosen K. [1984], 'Models for Normal Intuitionistic Modal Logics', *Studia Logica* **43**, pp. 217–245.

Brogaard B., Salerno J. [2002], 'Clues to paradoxes of knowability: reply to Dummett and Tennant', *Analysis* **62**, pp. 143–150.

Brogaard B., Salerno J. [2009], 'Fitch's Paradox of Knowability', *Stanford Encyclopedia of Philosophy*, http://plato.stanford.edu/entries/fitch-paradox/, accessed 03/04/2010.

Dosen K. [1985], 'Models for Stronger Normal Intuitionistic Modal Logics', *Studia Logica* **44**, pp. 39–70.

Dummett M. [1975], *Truth and Other Enigmas*, Duckworth, London.

Dummett M. [2001], 'Victor's error', *Analysis* **61**, pp. 1–2.

Fitch F. [1963], 'A Logical Analysis of Some Value Concepts', *The Journal of Symbolic Logic* **28**, pp. 135–142.

Hacking I. [1967], 'Possibility', *The Philosophical Review* **76**, pp. 143–168.

Mackie J. [1980], 'Truth and Knowability', *Analysis* **40**, pp. 90–92.

Rosenkranz S. [2004], 'Fitch back in action again?', *Analysis* **64**, pp. 67–71.

Salerno J. [2009], 'Knowability Noir: 1945–1963', in Salerno J. (ed.), *New Essays on the Knowability Paradox*, Oxford University Press, Oxford, pp. 29–48.

Tennant N. [1997], *The Taming of the True*, Clarendon Press, Oxford.

Tennant N. [2001], 'Is Every Truth Knowable? Reply to Williamson', *Ratio* **14**, pp. 263–280.

Tennant N. [2002], 'Victor vanquished', *Analysis* **62**, pp. 135–142.

Tennant N. [2009], 'Revamping the Restriction Strategy', in Salerno J. (ed.), *New Essays on the Knowability Paradox*, Oxford University Press, Oxford, pp. 223–238.

Tennant N. [2010], 'Williamson's Woes', *Synthese* **173**, pp. 9–23.

Williamson T. [1992], 'On Intuitionistic Modal Epistemic Logic', *Journal of Philosophical Logic* **21**, pp. 63–89.

Williamson T. [2000a], *Knowledge and Its Limits*, Oxford University Press, Oxford.

Williamson T. [2000b], 'Tennant on Knowable Truth', *Ratio* **13**, pp. 99–114.

Williamson T. [2009], 'Tennant's Troubles', in Salerno J. (ed.), *New Essays on the Knowability Paradox*, Oxford University Press, Oxford, pp. 183–204.

Zardini E. [2008], 'Truth and What Is Said', *Philosophical Perspectives* **22**, pp. 545–574.

Zardini E. [2010a], 'Truth, Demonstration and Knowledge. A Solution to the Paradox of Knowability', ms.

Zardini E. [2010b], 'Black Boxes. The Semantics and Logic of Obliterative Modalities', ms.