

Modal Epistemology¹

Juhani Yli-Vakkuri and John Hawthorne

1. Modal epistemology

When theorizing about the a priori, philosophers typically deploy a sentential operator: ‘it is a priori that’. This operator can be combined with metaphysical modal operators, and in particular with ‘it is necessary that’ and ‘actually’ (in the standard, rigidifying sense) in a single argument or a single sentence. Arguments and theses that involve such combinations have had played a starring role in post-Kripkean metaphysics and epistemology. The phenomena the contingent a priori and the necessary a posteriori have been organizing themes in post-Kripkean discussions, and these phenomena cannot be easily discussed without using sentences and arguments that involve the interaction of the a priority, necessity, and actuality operators. However, there has been surprisingly little discussion of the logic of the interaction of these operators.² In this paper we shall attempt to make some progress on that topic.

Our main starting point is the idea that apriority has something to do with—in particular, at least entails—the metaphysical possibility of a priori knowledge. We will also take at face value some paradigm cases of the contingent a priori. The central upshot is that, given these starting points, the logic of apriority turns out to be very weak indeed. In particular, the logic of apriority turns out not to be a *normal* modal logic, and it turns that it does not even satisfy the two jointly necessary and sufficient conditions of normality: the so-called K axiom, which says that anything that follows by *modus ponens* from what is a priori is also a priori, and the ‘necessitation’ principle according to which all logical truths are a priori. On one very widely used definition—the ‘standard’ one according to Chalmers and Rabern (2014: 214)—apriority turns out not even to be factive: that is to say, it turns out that something can be both false and a priori.

These facts ought to be of philosophical interest in that they show that many forms of inference that philosophers habitually make using the apriority operator are not valid, even if in some cases they happen to preserve truth. Thus any use of those forms of inference must be justified by appeal to something other than their validity or—significantly—even to their necessary exceptionlessness. There are, of course, non-valid forms of inference that never lead from truth to falsehood, even in the scope of counterfactual suppositions. For example, since it is necessary that Hesperus =

¹ We would like to thank David Chalmers and Peter Fritz for helpful comments and discussions.

² The only examples we know of are Rabinowicz and Segerberg (1994), Restall (2012), Fritz (2013, 2014), and Chalmers and Rabern (2014).

Phosphorus, the practice of inferring that Hesperus = Phosphorus from anything whatsoever, even under a counterfactual supposition, will never lead from truth to falsehood. But the patterns of inference that will be shown to be invalid here are not like that: one cannot respond by claiming that, even though they are not logically valid, they are good enough to use in philosophical arguments, including in counterfactual suppositional reasoning (which is, of course, ubiquitous in philosophy). They are not. The very same arguments that show those forms of inference to be invalid show that they will also, in some cases, lead from truth to falsehood under the scope of counterfactual suppositions. Indeed, in some cases the arguments establish the existence of actual counterexamples: cases in which the relevant form of inference leads from an actual truth to an actual falsehood.

The notions of apriority that we are working with combine the notion of knowledge with that of metaphysical modality, via the assumption that apriority entails the metaphysical possibility of a priori knowledge, and therefore of knowledge. There are other epistemological notions that are expressed by some sentential operator Ω that involve some combination of knowledge and metaphysical modality (or some restriction thereof), and that, in particular, are such that $\Omega\phi$ entails $\Diamond K\phi$ ('It is metaphysically possible that it is known that ϕ '). We call such notions *modal-epistemological notions* and their study *modal epistemology*. For example, the notion of *being in a position to know*, which is a familiar component of the analytic epistemologist's toolkit, is a modal-epistemological notion. Might our reflections on apriority carry over to certain of these other composite notions? Taking being in a position to know as a test case, we argue that many of the themes of the earlier discussion can be adapted to show that the logic being in a position to know is also extraordinarily weak—in fact, that it fails to be normal in just the same ways as the logic of apriority. Thus great caution is in order whenever one makes certain seemingly unexceptionable inferences using the notion of being in a position to know. Just as in the case of apriority, not only are those inferences not valid, but they may well lead from truth to falsehood.

Although we do not investigate modal-epistemological notions other than apriority and being in a position to know here, it will be fairly straightforward to extrapolate from our discussion conclusions about the weakness of their logics as well. One might have thought or hoped that one could tame and normalize a species of knowledge by modalizing it: 'Even if this kind of knowledge is not closed under *modus ponens* (etc.), some related kind of *knowability* is!' But this, we suggest, is a forlorn hope: the mere entailment from the modal-epistemological notion to the possibility of knowledge is liable to block normality in the ways exhibited in this paper.

2. Informal validity

In this section we will briefly introduce the theoretical framework we will be working with throughout this paper. For now we will merely try to be explicit about our conceptions of logical validity and consequence, and to sketch the Tarski-inspired philosophical picture that motivates it. (We will later introduce further assumptions concerning the logic of the notions whose interaction we are interested in.)

A *logic* in the sense of this paper is simply a set of sentences. *The logic* of a language L is the set of precisely those sentences of L that are true on every interpretation of L 's non-logical constants—for short, *true on every interpretation*. We will call an interpretation on which a sentence ϕ is true (false) a *true (false) interpretation of ϕ* , and we will say that a sentence that is true on every interpretation is *valid* or, when a noun is needed, is a *logical truth*.

With the exception of a brief departure from this assumption in §5.1, we will assume that the language we are discussing is one whose only logical constants are sentential operators: specifically, at least the standard truth-functional connectives \neg , \wedge , \vee , \rightarrow , \leftrightarrow , the metaphysical modal operators \Box ('necessarily') and $@$ ('actually'), with $\Diamond\phi$ ('Possibly ϕ ') defined as $\neg\Box\neg\phi$, and one or more knowledge operators K_1 , K_2 , ... As usual, ' \perp ' will designate some truth-functional contradiction (it doesn't matter which). We will assume that all classical truth-functional tautologies are in the logic of the language, and that the logic is closed under *modus ponens*.

Given these assumptions, it is natural to define a finitary notion of logical consequence in terms of truth-functional connectives and validity. We will say that ψ logically follows from, or is entailed by, ϕ_1, \dots, ϕ_n (in symbols: $\phi_1, \dots, \phi_n \models \psi$) just in case $\phi_1 \wedge \dots \wedge \phi_n \rightarrow \psi$ is valid. (We will sometimes abbreviate ' ϕ is valid' as ' $\models \phi$ ') We'll say that ϕ_1, \dots, ϕ_n are *inconsistent* just in case they entail \perp .

(Of course, the notion of validity as truth on every interpretation is relative to a choice of logical constants. We don't think that there is anything especially natural about the logical constants we have chosen; they simply happen to represent the features of the world that we, as students of modal epistemology, are interested in: necessity, actuality, and various kinds of knowledge.)

Since all of the logical constants we are dealing with are sentential operators, we will, as is customary, assume that the non-logical constants of the language are its atomic sentences, i.e., those sentences that are not formed out of other sentences by the application of any sentential operators to them. A sentence ϕ , then, is valid (or is a logical truth) just in case ϕ is true on every interpretation of all of the atomic sentences.

The notion of validity we are working with is the standard *informal* one due to Tarski (1936). It is a semantic notion of validity, but we will not be giving a 'semantics' in the *formal* sense of a model-theoretic truth definition for the language we are interested in. Instead, we will rely on our own and the reader's informal understanding of the meanings of the modal and epistemic operators in our arguments.

We need not say much at this point about our conception of interpretations, except that we are thinking of them as assignments of propositions to all of the atomic sentences. (In §7 we will see that this may be something of an idealization, but it is not an idealization that makes any difference to our arguments.) We do not have any very specific conception of a proposition in mind, but we will assume that propositions are not modally coarse-grained, meaning that it is possible for necessarily equivalent sentences to express distinct propositions. We are aware that some philosophers subscribe to theories of propositions that are inconsistent with this assumption, and we will have something to say to them in §7.

The reason we will not do any model theory in this paper is that we think it would be premature. When even the most basic principles of the logic of a certain

notion are up for grabs, and especially when one is only arguing that certain principles are *not* valid, as we will be, it would be inappropriate to give a semantics and to expect one's readers to accept that what a sentence is valid only if it is valid on the semantics. It is a trivial exercise to give a semantics that invalidates any given principle. For example, one can associate \neg with any truth function or none; one can treat it like a modal operator, as in Kripke semantics for intuitionistic logic; and so on. That these things can be done by itself tells us nothing about the logic of negation. Of course, one can sometimes motivate a particular semantics by an informal argument, but even then the argument must proceed from some pretty substantive assumptions about what is informally valid. Because of our dialectical position as defenders of the claim that hardly any interesting principles concerning the notions of interest are informally valid, we do not have enough such assumptions to motivate any particular semantics.

A note about terminology: we will often speak somewhat loosely of taking attitudes like acceptance and rejection towards 'principles' or 'axioms' like $A\phi \rightarrow \phi$ or rules of inference like $\phi/A\phi$, or of an operator like A 'obeying' such 'axioms', 'principles' or rules. Strictly speaking, what is at issue in such cases is the *schema* $A\phi \rightarrow \phi$ and whether all of its instances are valid: to 'accept' or 'reject' it, or to say that the relevant operator 'obeys' or doesn't 'obey' it, is to accept or reject the claim that all of its instances are valid. Similarly, what's at issue in the case of a rule of inference is whether the logic is closed under it.

3. Chalmers and Rabern's observation

Chalmers and Rabern (2014) have made a start on the project we are interested in. Like us, they are interested, *inter alia*, in observing the consequences of a definition of a priority in terms of what it is possible to know a priori. With some justification, they call the following definition the 'the standard way' of understanding the apriority operator A .

Definition 1. $A\phi =_{df} \Diamond K^A\phi$

Here ' $K^A\phi$ ' is to be read as 'It is known a priori that ϕ '. (Like us, Chalmers and Rabern do not take Definition 1 to be set in stone. They consider semantics on which its two sides are not logically equivalent.) What they show is that, given a very minimal logical assumption, namely the validity of the axiom **K** ($\Box(\phi \rightarrow \psi) \rightarrow (\Box\phi \rightarrow \Box\psi)$), the following trio is inconsistent.

$$(A1) \quad A\phi \wedge \neg\Box\phi$$

$$(A2) \quad A\phi \rightarrow \Box A\phi$$

$$(A3) \quad \Box(A\phi \rightarrow \phi)$$

They also claim, in effect,³ that it is *prima facie* plausible that there is an interpretation on which (A1), (A2), and (A3) are all true. Clearly, as they observe, there is at least one true interpretation of (A1) provided that there is at least one contingent a priori truth—a piece of post-Kripkean orthodoxy that neither they nor we have any interest in challenging. (A2), on the other hand, is plausibly valid ‘if $[\Box\phi]$ is understood in the standard way’ (i.e., according to Definition 1), because, so understood, (A2) is an instance of axiom **5** ($\Diamond\phi \rightarrow \Box\Diamond\phi$). Regarding (A3), Chalmers and Rabern only say that it ‘is an instance of the claim that apriority is modally factive’ (p. 214). Later in the paper they cite ‘our clear intuitions of modal factivity’ (p. 222). But intuitions aside, it does seem that the assumption that (A2) is valid underwrites many discussions of apriority. After all, many philosophers write as if it would be a limitation of rationality to fail to know a proposition that is a priori. But if there could be an a priori proposition that is false, one could hardly find an agent who fails to know it when it is false to be less than perfectly rational on account of that.

Thus, it seems, there is fairly compelling *prima facie* case that there is an interpretation on which all of (A1)-(A3) are true. But the combination of (A1)-(A3) is inconsistent with **K**, which is surely valid, so there is no interpretation on which (A1)-(A3) are all true. This, Chalmers and Rabern say, is a ‘general problem about the interaction of modal and epistemic operators’ (p. 213). They go on (in §5) to develop three semantics that invalidate one or another of (A2) and (A3).

The observations that the trio (A1)-(A3) is inconsistent with **K** and that (A2) is an instance of **5** on Definition 1 are beyond question and not very surprising. It is, after all, pretty obvious that the phenomenon of the contingent a priori is in tension with the theses that anything a priori is necessarily a priori and that, necessarily, anything a priori is true. Given that we are not questioning the existence of contingent a priori truths, the following alternative definition is a natural retreat.

Definition 2. $A\phi =_{df} \phi \wedge \Diamond K_A\phi$

With Definition 2 in place, we keep the necessary factivity of apriority ($\Box(A\phi \rightarrow \phi)$) while giving up the principle that whatever is a priori is necessarily a priori ($A\phi \rightarrow \Box A\phi$). Alternatively, of course, one might learn to live with the idea that something could be both a priori and false. One might, as a back-up, try claiming that the principle that anything that is a priori is true is a contingent logical truth. Since post-Kripkean orthodoxy already embraces the idea of contingent logical truths anyway,⁴ it would not be an enormous further step to think that $A\phi \rightarrow \phi$ is one of them. (We will return to this back-up plan in §5.)

This is not the end of the story, however. Chalmers and Rabern’s recent reflections on the contingent a priori manage to scratch the surface of some quite serious problems for logicians of apriority—indeed, arguably, problems for anyone

³ We are perhaps not being entirely true to Chalmers and Rabern’s paper here, since they are not explicitly presenting the problem in terms of validity as truth on all interpretations. They simply say, of (A1)-(A3), that ‘[a]ll of these claims are initially plausible’ (p. 214). But presumably they have in mind something like: *at least one trio of this form is such that all of its members are initially plausible* or *it is initially plausible that all of these sentences are true on at least one interpretation*.

⁴ We consider Kaplan’s ‘Demonstratives’ to be one of the gospels. See Kaplan 1977: §XV, esp. p. 539, n. 65.

who both wishes to make serious philosophical use of apriority conceived as a modal-epistemological notion.

4. Background logical assumptions

As advertised in §2, we will now make our main logical assumptions explicit. As the logic of apriority and its interaction with other operators is up for grabs, we are not going to lay down any ground rules in that area. Instead we will be content here to make some uncontentious assumptions about the behavior of the other operators in the language.

First, we will make the standard assumption that the modal sublogic of the logic of the language is **S5**. This means that the sublogic is the smallest set of sentences that includes, in addition to all tautologies, axioms **5**, **T**, and **K**:

$$\mathbf{5}: \quad \diamond\phi \rightarrow \diamond\Box\phi$$

$$\mathbf{T}: \quad \Box\phi \rightarrow \phi$$

$$\mathbf{K}: \quad \Box(\phi \rightarrow \psi) \rightarrow (\Box\phi \rightarrow \Box\psi)$$

and is closed under *modus ponens* and the rule of *necessitation*,⁵

$$\mathbf{NEC}: \quad \frac{\phi}{\Box\phi.}$$

S5 encodes the standard ideas that (i) what is necessary is a non-contingent matter (axiom **5**), that (ii) necessity implies truth (**T**), that (iii) whatever necessarily follows from a necessary truth is necessary (**K**), that (iv) all logical truths (in the sublogic) are necessary (**NEC**), and that (v) anything that logically follows (in the sublogic) from a necessary truth is necessary (**NEC** and **K**).

Following established usage, we will say that an operator Ω has *normal* logic iff the logic includes (or Ω ‘obeys’) the *K axiom for Ω* ($\Omega(\phi \rightarrow \psi) \rightarrow (\Omega\phi \rightarrow \Omega\psi)$) and is closed under (‘obeys’) *NEC for Ω* ($\phi/\Omega\phi$). We will use ‘ \mathbf{K}_Ω ’ to designate the axiom $\Omega(\phi \rightarrow \psi) \rightarrow (\Omega\phi \rightarrow \Omega\psi)$, and similarly for ‘ \mathbf{T}_Ω ’, and we will use ‘ \mathbf{NEC}_Ω ’ to designate the rule $\phi/\Omega\phi$ (thus $\mathbf{K}_\Box = \mathbf{K}$ and $\mathbf{NEC}_\Box = \mathbf{NEC}$).

Note that, while we have assumed that the modal sublogic of the language is normal, we have not assumed (closure under) **NEC** for the logic of the whole language. We have two reasons for not assuming **NEC**. The first reason is that there are other operators in the language whose interaction with \Box we should not prejudge. The second reason is that there is one operator in the language whose interaction with \Box we *should* prejudge, because the combination of that interaction with **NEC** would make a complete hash of things.

The operator in question is $@$. We will make the standard assumptions that $@$ obeys the axiom

⁵ Note that closure under **NEC** does not imply that $\phi \rightarrow \Box\phi$ is in the modal sublogic, but only that $\phi \rightarrow \Box\phi$ is in it when ϕ is in it.

T_@: $@\phi \rightarrow \phi$,

which says that whatever is actually so is so, and that it also obeys the standard rigidity axiom

RIG: $\phi \rightarrow \Box @\phi$.

RIG encodes the idea that $@$ is a modally rigidifying operator: it takes the proposition its operand expresses to a non-contingent proposition. In particular, if the proposition that p is true, then the proposition that actually p is a necessary truth, and otherwise the proposition that actually p is a necessary falsehood. In a semantic framework like Kaplan's (1977), $@$ is treated as an indexical operator: what it does to its operand depends on the context-*cum*-world in which it is used, but we are not making use of the ideology of contexts in our official conception of validity. (Although we take that ideology seriously and we will find it useful to appeal to it in some of our informal arguments below.) The rigidifying property of $@$ is what enables it play its characteristic role in correctly formalizing various truths we express using the English 'actually'. For example, as an homage to Russell (1905), we note that the sentence

Frank's yacht could have been longer than it actually is

is true because, while it is a contingent matter how long Frank's yacht is, it is a non-contingent matter how long Frank's house actually is.

Now we are in a position to see precisely why could not assume **NEC**. Suppose for a contradiction that it does. By **K** and **RIG**, $\Box(\phi \leftrightarrow @\phi)$ is false whenever ϕ is contingent (i.e., whenever $\neg\Box\phi \wedge \neg\Box\neg\phi$ is true). By **RIG** and **T_@**, $\phi \leftrightarrow @\phi$ is valid, so assuming **NEC** would give us the incorrect result that $\Box(\phi \leftrightarrow @\phi)$, which may be false, is valid.

This is a familiar observation. Since $@$ gives rise to contingent logical truths, **NEC** fails in a language that includes $@$. But even in such a language the logic of \Box is well-behaved and close enough to normal: we still have **K**, and the exceptions to **NEC** can be neatly quarantined, enabling a complete axiomatization of the logic of necessity and actuality (see Crossley and Humberstone 1977).

We will use K^A to formalize 'it is known a priori that' and ' K ' to formalize 'it is known that'. Here we will only be assuming the principles that knowledge is necessarily factive:

NecFac: $\Box(K\phi \rightarrow \phi)$

and that it is necessary that whatever is known a priori is known ($\Box(K_A\phi \rightarrow K\phi)$). The principle that apriority is necessarily factive:

NecFac_A: $\Box(K_A\phi \rightarrow \phi)$

follows from these.

5. How weak is the logic of the apriori?

It is a standard assumption that the logic of apriority is at least as strong as the normal modal logic **KT**, in the sense that it is normal and includes the \mathbf{T}_A ($A\phi \rightarrow \phi$) and is normal. If it is normal it also includes

$$\mathbf{K}_A: \quad A(\phi \rightarrow \psi) \rightarrow (A\phi \rightarrow A\psi)$$

and is closed under \mathbf{NEC}_A ($\phi/A\phi$).

Each component of the standard assumption seems *prima facie* very compelling. \mathbf{T}_A says that whatever is a priori is true. \mathbf{K}_A says that whatever follows from what is a priori by *modus ponens* is also a priori. \mathbf{NEC}_A amounts to the assumption that all logical truths are a priori. The latter is a way of precisifying the commonplace that *logic is a priori*.

The well-known counterexamples to **NEC** do not seem to make any trouble for that commonplace. As we have already noted (note 4), it is affirmed in the gospels that anything of the form $\phi \leftrightarrow @\phi$ is valid and a priori, and in particular that $\phi \leftrightarrow @\phi$ is a contingent a priori truth whenever ϕ is contingent.⁶ Contingent instances of $\phi \leftrightarrow @\phi$ are among the paradigms of the contingent a priori.

Above we introduced two alternative definitions of the apriority operator A , namely:

Definition 1. $A\phi =_{df} \Diamond K_A\phi$,

Definition 2. $A\phi =_{df} \phi \wedge \Diamond K_A\phi$

We will begin our investigation by asking how the standard assumption that the logic of A is at least as strong as **KT** looks through the lens of Definition 1. Then we will move on to examine it through the lens of Definition 2.

5.1. Definition 1

Because we did not assume **NEC**, we cannot prove $\Box(A\phi \rightarrow \phi)$ from \mathbf{T}_A and **NEC**. For this reason, we cannot make immediate trouble for the standard assumption of that the logic of the apriority is at least as strong as **TK** by combining the presumptive validity of $\Box(A\phi \rightarrow \phi)$ with that of $A\phi \rightarrow \Box A\phi$ (which is valid on Definition 1, being an instance of **5**) and the fact that, given the phenomenon of the contingent a priori, there is a true interpretation of $A\phi \wedge \neg\Box\phi$. (Thus we cannot reinstate Chalmers and Rabern's problem even though we are working with Definition 1.) On our assumptions, there is simply no presumption of the validity of $\Box(A\phi \rightarrow \phi)$, so no immediate problem for \mathbf{T}_A . If \mathbf{T}_A is valid, some of its instances are contingent logical truths—and contingent logical truth is, again, a phenomenon whose reality is at the core post-Kripkean orthodoxy and is affirmed in the gospels (see note 4).

So let's set the issue of its necessitateability aside and directly ask: is \mathbf{T}_A valid? Given Definition 1 it is, in fact, easy to show that \mathbf{T}_A is not valid using standard

⁶ See Chalmers (2011) for a dissenting view.

examples of the contingent a priori (and contingent logical truth). Again, $\phi \leftrightarrow @\phi$ expresses a contingent proposition on any interpretation on which ϕ expresses a contingent proposition. In particular, if p is the proposition assigned to ϕ , then $\phi \leftrightarrow @\phi$ expresses a proposition that is necessarily equivalent to p if p is true, and otherwise $\phi \leftrightarrow @\phi$ expresses a proposition that is necessarily equivalent to the negation of p . Now consider a world w in which some sentence χ expresses a proposition p_w that is actually false but true in w . To make things vivid, suppose further that χ is the eternal, non-indexical sentence ‘Donald J. Trump loses the U.S. presidential election in 2016’. (Thus w may be a world very close to the actual one.) Then p_w is the proposition that DJT loses, etc., and χ expresses that proposition both actually and in w , and in w the sentence $\chi \leftrightarrow @\chi$ expresses a proposition p'_w necessarily equivalent to p_w . Suppose further that p'_w comes to be known a priori in the usual way in w : in w , it is known that the sentence $\chi \leftrightarrow @\chi$ is a logical truth, and on that basis p'_w comes to be known a priori. Now consider an interpretation that assigns p'_w to ϕ . Clearly, on that interpretation, $\Diamond K_A \phi$ is true and ϕ is false, so $\Diamond K_A \phi \rightarrow \phi$ is false on that interpretation. By Definition 1, $\Diamond K_A \phi \rightarrow \phi = A\phi \rightarrow \phi$, so $A\phi \rightarrow \phi$ is not valid.

On Definition 1, then, the logic of apriority is not at least as strong as **KT**, because it does not include **T_A**.

Note that our argument against **T_A** does not essentially depend on the presence of **@** in the language. It suffices that it is possible that in some language there is an operator that works like **@**, a material biconditional connective, and a sentence that expresses a proposition p that is true but actually false, and a contingent truth q necessarily equivalent to p is known a priori by any of the methods we use to come to know $\phi \leftrightarrow @\phi$ a priori. q is then possibly known a priori but actually false, and $A\phi \rightarrow \phi$ is false on any interpretation that assigns q to ϕ .

Nor does the argument even essentially depend on the possibility of a language with an operator that works like **@**, or any particular modal operator. It simply turns on the observation that there are contingent propositions that are actually false but can be known *a priori* by a counterfactual use of one or another of the paradigmatic methods for acquiring *a priori* knowledge of contingent matters. For example, to take a standard example, we actually fix (following Evans 1979) the reference of ‘Julius’ using the description ‘the inventor of the zip’, but there is a possible counterfactual situation in which the semantic workings of English are otherwise as they actually are, and in which the actual inventor of the zip—call him ‘ x ’—who is not actually the unique inventor of sliced bread, is the unique inventor of sliced bread, and we fix the reference of ‘Julius’ instead using ‘the inventor of sliced bread’. In that counterfactual situation we come to know *a priori* through one of the usual methods the actually false proposition that x , if x exists, invented sliced bread. Any interpretation that assigns that proposition to ϕ is a counterexample to **T_A**.

Note that we have, in effect, noticed that a backup plan anticipated in §3 is hopeless on Definition 1. The backup plan was to deny $\Box(A\phi \rightarrow \phi)$ while accepting **T_A**.

Having given up \mathbf{T}_A , one might still hope that Definition 1 allows A to have a normal logic, but this is a vain hope. Consider the principle of agglomeration, which is valid if the logic is normal:⁷

$$\mathbf{AGL}_A: \quad (A\phi \wedge A\psi) \rightarrow A(\phi \wedge \psi)$$

To find a false interpretation of \mathbf{AGL}_A , consider again the case of the sentence $\chi \leftrightarrow @\chi$, as described above. Because DJT did not lose, $\chi \leftrightarrow @\chi$ expresses a proposition p necessarily equivalent to the proposition that DJT did not lose (because he actually won). Moreover, because DJT did not lose in w , in w $\chi \leftrightarrow @\chi$ expresses a proposition q necessarily equivalent to the proposition that DJT did not lose. p is known a priori in w , and q is actually known a priori, so p is possibly known a priori and so is q . Consider an interpretation that assigns p to ϕ and q to ψ . On that interpretation, clearly, $\Diamond K_A \phi \wedge \Diamond K_A \psi$ is true; but, just as clearly, $\Diamond(\phi \wedge \psi)$ is false on that interpretation (it is impossible for DJT to both lose and not lose), so, by \mathbf{NecFac}_A , $\Diamond K_A(\phi \wedge \psi)$ is false on that interpretation. Therefore

$$(\Diamond K_A \phi \wedge \Diamond K_A \psi) \rightarrow \Diamond K_A(\phi \wedge \psi)$$

is false on that interpretation. By Definition 1, the above is none other than $(A\phi \wedge A\psi) \rightarrow A(\phi \wedge \psi)$, so \mathbf{AGL}_A is not valid, and the logic is not normal.

Having given up any hope of a normal modal logic of apriority, one might still hope that Definition 1 might allow one to at least have one of the two ingredients of normality: \mathbf{K}_A and \mathbf{NEC}_A . And indeed nothing we have shown so far rules this out: \mathbf{AGL}_A follows from the combination of \mathbf{K}_A and \mathbf{NEC}_A , but it does not follow from either without the other. Perhaps, one might hope, we can still have at least one of the two.

This is yet another vain hope. Let us begin by considering \mathbf{K}_A .

To find a false interpretation of \mathbf{K}_A , suppose that we actually fix the reference of ‘Julius’ by the description ‘the number of ants in the universe’, and that we then come to know a priori, in the usual way, that Julius $\neq 1 +$ the number of ants in the universe, and suppose that, by deduction, we also come to know a priori that, if Julius $= 1 +$ the number of ants in the universe, then \perp . Suppose further that the speakers in some world w use ‘the successor of the number of ants in the universe’ to fix the reference of ‘Julius’ (but that the meanings-*cum*-characters of all other expressions are the same in w and in the actual world), and that they accordingly come to know a priori what they express by ‘Julius $= 1 +$ the number of ants in the universe’—call that proposition ‘ j ’. Finally, suppose that there is exactly one more ant in the actual world than there is in w . Thus j is the proposition that $n + 1$ is the number of ants in the universe, where n is the actual number of ants in the universe, and the proposition that Julius $= 1 +$ the number of ants in the universe is the proposition that $n + 1$ is the number of ants in the universe. Now consider an interpretation that assigns j to ϕ . Clearly $\Diamond K_A(\phi \rightarrow \perp)$ is true on that interpretation, since we in fact do know a priori that, if $n = 1 +$ the number of ants in the universe, then \perp . $\Diamond K_A \phi$ is also true on the

⁷ Since $\phi \rightarrow (\psi \rightarrow (\phi \wedge \psi))$ is a tautology, $A(\phi \rightarrow (\psi \rightarrow (\phi \wedge \psi)))$ is valid by closure under A -necessitation, and by the \mathbf{K} axiom for A , so is $(A\phi \wedge A\psi) \rightarrow A(\phi \wedge \psi)$.

interpretation, because in w it is known a priori that $n = 1 +$ the number of ants in the universe. By **NecFac_A**, $\Diamond K_A \perp$ is false on every interpretation, so follows that

$$\Diamond K_A(\phi \rightarrow \perp) \rightarrow (\Diamond K_A \phi \rightarrow \Diamond K_A \perp),$$

which, by Definition 1, is an instance of **K_A**, is false on the interpretation under consideration.

Finally, consider **NEC_A**. Counterexamples to **NEC_A** are easy to come by if we enrich the language with propositional quantification, so let us do so for the moment. Since $\phi \leftrightarrow @\phi$ is valid, so, by universal generalization, is $\forall p(p \leftrightarrow @p)$.⁸ $\forall p(p \leftrightarrow @p)$ has one very interesting feature: the proposition it expresses is true in exactly one world: the actual one. It says that every proposition is such that it is true if and only if it is actually true, i.e., true in the actual world. (Used in any other world w , $\forall p(p \leftrightarrow @p)$ expresses a proposition that is true in w and in no other world.) By **NecFac_A**, it follows that $\forall p(p \leftrightarrow @p)$ can be known if and only if it is actually known. Now it is plausible that $\forall p(p \leftrightarrow @p)$ is actually known (since it is plausible that, even if you didn't know it before you read this paragraph, you do now). But of course not every proposition expressed by a valid sentence is known. This is particularly clear if we allow—as we do in the present setting⁹—truths of second-order logic to count as valid. There is a second-order sentence κ such that κ is valid if the Continuum Hypothesis (CH) is true, and $\neg\kappa$ is valid if CH is not true,¹⁰ and it is not known whether κ is true. Now consider the sentences:

$$(\kappa) \quad \kappa \wedge \forall p(p \leftrightarrow @p)$$

$$(\neg\kappa) \quad \neg\kappa \wedge \forall p(p \leftrightarrow @p)$$

If CH is true, then κ and therefore (κ) is valid, and the proposition expressed by (κ) is true only in the actual world. If CH is not true, then $\neg\kappa$ and therefore $(\neg\kappa)$ is valid and the proposition expressed by $(\neg\kappa)$ is true only in the actual world. By **NecFac_A** and Definition 1, then, either (κ) or $(\neg\kappa)$ is a counterexample to **NEC_A**. After all, we have seen that one of the following sentences is true

$$\neg\Diamond K_A(\kappa \wedge \forall p(p \leftrightarrow @p))$$

$$\neg\Diamond K_A(\neg\kappa \wedge \forall p(p \leftrightarrow @p)),$$

and these sentences are, by Definition 1, none other than the following.

$$\neg A(\kappa \wedge \forall p(p \leftrightarrow @p))$$

⁸ It can also be shown to be valid on Kaplan's (1977) standard semantics for $@$ supplemented by Kaplan's (1970) equally standard semantics for propositionally quantified modal logic.

⁹ Propositional quantification is second-order quantification (into the position of a 0-place predicate).

¹⁰ See Etchemendy (1989: 123).

$$\neg A(\neg\kappa \wedge \forall p(p \leftrightarrow @p))$$

Depending on whether CH is true, either (κ) or $(\neg\kappa)$ expresses a proposition that, so to speak, *you only get one shot at knowing*: if it isn't actually known, then it's impossible for it to be known, and, *a fortiori*, it's impossible for it to be known a priori.

Of course, in §2 we said we would be concerned with propositional logic, so the above argument doesn't directly speak to the narrow question we posed about the strength of the logic, in the sense of propositional logic, of apriority. Nevertheless the result is significant: it shows NEC_A fails in the second-order logic of apriority, independently of any definitions, as long as apriority implies the possibility of a priori knowledge. In fact, the result only requires extremely weak second-order logic: namely, second-order universal generalization, by which the validity of $\phi \leftrightarrow @\phi$ implies the validity of $\forall p(p \leftrightarrow @p)$. Given this much second-order logic, κ could be any unknown logical truth, and would follow that (κ) is a counterexample to NEC_A . In light of this observation, it would not be a very exciting fact, if it were a fact at all, that there are no counterexamples to NEC_A in the propositional logic of apriority.

But is it a fact at all? The answer turns out to be 'No'. Recall that, while we made various assumptions about which operators are present in the propositional language, we made no assumptions at all about which operators are *not* present in it. Here is one that could be easily introduced into Kaplan's Logic of Demonstratives, on which we have been freely drawing, since it is, once again, one of the gospels. Our 0-place operator $@_0$ is a close relative of the standard actuality operator $@$. It is, in fact, a generalization of it—a more powerful operator in terms of which we can *define* the standard actuality operator. (For what it's worth, our view is that the correct approach to the Logic of Demonstratives is to treat $@_0$ as primitive and omit $@$.) In any world (context) w , $@_0$ expresses the *world proposition* that is true exactly in w (the *world proposition of* w). Without getting into the formal details,¹¹ one can think of the world-proposition of w as the conjunction of all of the propositions that are true in w . On a coarse-grained conception, on which propositions are sets of worlds, this is simply $\{w\}$. Theorists of fine-grained propositions will allow that many propositions may be true in exactly one world. Exactly which proposition $@_0$ expresses on fine-grained views will depend on both the details of those views and decisions about the semantics of $@_0$ —we will not worry about these details, although we are confident

¹¹ Here is how $@_0$ should be introduced into Kaplan's Logic of Demonstratives. In order to be completely unambiguous, we use exactly his notation. First, add to §XVIII, p. 542, the syntactic clause 8.1, which says ' $@_0$ is a 0-place modal operator', and modify syntactic clause 10 (*ibid.*) accordingly, deleting its third line. Then, on p. 545, replace semantic clause 10(ii) with:

$$(ii) \models_{c/w} @_0 \text{ iff } w = c_w$$

Finally, on p. 543, after syntactic clause 11, add:

$$\textbf{Definition: } @_0\phi = \Box(@_0 \rightarrow \phi)$$

With these changes in place, semantics validates every formula that Kaplan's semantics validates, but there is no primitive actuality operator in the language, and the semantics also validates certain formulae that Kaplan's doesn't. Most importantly, we get the result:

$$\models @_0,$$

but not:

$$\models \Box @_0,$$

that they can be worked out provided that those views are consistent.¹² Since we could define $@\phi$ as $\Box(@_0 \rightarrow \phi)$, we could replace the assumption that $@$ is in the language with the assumption that $@_0$ is. (We only made the former assumption because $@$ is a more familiar beast.)

$@_0$ is valid. Therefore so is $@_0 \wedge \lambda$ if λ is valid. But of course not every logical truth is known. (Logical omniscience is an idealization that serves many investigations in epistemic logic well, but it is false. We will return to this theme in §7.) Let λ then be an unknown logical truth, and suppose that the proposition expressed by $@_0 \wedge \lambda$ is not known. (We would not need to make this additional supposition if we could help ourselves to the principle that knowledge distributes over conjunction in the sense of $K(\phi \wedge \psi) \rightarrow (K\phi \wedge K\psi)$, but this begins to look questionable when the idealization of logical omniscience is given up.) Since the proposition expressed by $@_0 \wedge \lambda$ is not known, it is not known a priori, and, by **NecFac_A**, since that proposition is only true in the actual world, it cannot be known a priori, and $\neg\Diamond K_A(@_0 \wedge \lambda)$ is true. Validity entails truth, so $\Diamond K_A(@_0 \wedge \lambda)$ is not valid, and now we have a counterexample to **NEC_A**: $@_0 \wedge \lambda$ is valid and, by Definition 1, $\neg\Diamond K_A(@_0 \wedge \lambda) = \neg A(@_0 \wedge \lambda)$. As before, the counterexample is independent of any definition: it is a counterexample as long as apriority implies the possibility of a priori knowledge.

In summary, Definition 1 has no hope of securing even the most basic putative principles of the logic of apriority.

5.2. Definition 2

Let us now turn to Definition 2. Here **T_A** is secured by brute force. But, sadly, it turns out, nothing else of interest is secured.

We already know that the logic of apriority is not normal on Definition 2, because we know that **NEC_A** fails on Definition 2: exactly the same counterexamples that worked on Definition 1 also work on Definition 2.

All that is left, of normality, then, is **K_A**. But even this last vestige of normality turns out to be a mirage, for familiar reasons. Suppose now that λ is a non-contingent valid sentence such that $K\lambda$ is false but possibly true, and further, for that reason, $K(\lambda \wedge @_0)$ is false too. Suppose further that $A(@_0 \rightarrow (\lambda \wedge @_0))$ is true. (Note that the proposition expressed by $@_0 \rightarrow (\lambda \wedge @_0)$ is *not* one of those that ‘you only get one shot at knowing’; it is necessarily true.) The first supposition is completely

¹² Perhaps the most difficult case is that of the theorist (such as Salmon 1986, Soames 1989, 2010, and King 2007) who holds that propositions have sentence-like structure. We would not be terribly disturbed if we encountered difficulties there, since the inconsistency of such theories seem to be demonstrable by the Russell-Myhill argument (see Goodman 2017 for a higher-order reconstruction of the argument). But, if we set aside worries about their consistency, there is no problem at all about using their resources for finding a suitable proposition to associate with $@_0$ in each world in which it is used. Of course, theorists of structured propositions cannot accept the view that there is any such thing as the—or even *a*—conjunction of all true propositions (being a true proposition, it would have to have itself as a constituent), but they provide us with plentiful other resources for specifying the proposition $@_0$ expresses when used in a world. Soames, for example, deploys both Kaplan’s (1977) ‘dthat’ operator and first-order quantification over worlds with abandon, and he is happy to talk about worlds obtaining. In his framework, then, we could say that $@_0$ has the same character as the sentence ‘Dthat(the world that obtains) obtains’.

unremarkable: again, there is no logical omniscience. The second supposition is no more remarkable: it follows from the first supposition that $@_0 \rightarrow (\lambda \wedge @_0)$ is valid, and even if not every validity is a priori, it would be utterly bizarre to lay down a ground rule according to which *no* validity of the form $@_0 \rightarrow (\lambda \wedge @_0)$ is. Finally, clearly $A@_0$ is true—among other reasons, because it is actually known a priori that *actuality obtains*, and that is what $@$ says, and whatever is actually known a priori is a priori. Now consider the following instance of \mathbf{K}_A .

$$A(@_0 \rightarrow (\lambda \wedge @_0)) \rightarrow (A@_0 \rightarrow A(\lambda \wedge @_0))$$

It is false, because it and our assumptions imply $A(\lambda \rightarrow @_0)$, which, by Definition 2, implies $\Diamond K_A(\lambda \wedge @_0)$, and $\Diamond K_A(\lambda \wedge @_0)$ is false. After all, the proposition expressed by $\lambda \wedge @_0$ is one of those that ‘you only get one shot at knowing’. It is not actually known a priori, and it is only true in the actual world, so, by \mathbf{NecFac}_A , it is not possibly known a priori.

One objection to the foregoing might proceed by attacking the assumption that we don’t ‘get only one shot at knowing’ the proposition expressed by $@_0 \rightarrow (\lambda \wedge @_0)$. While that proposition is true in every world, one might think that (e.g., on account of its being ‘directly about’ the actual world) it can be grasped, and therefore known, only in the actual world. Given that apriority entails the possibility of knowledge one might thus balk at the supposition that $A(@_0 \rightarrow (\lambda \wedge @_0))$ is true. Notice, however, that, given that logical omniscience is false, we could instead have made trouble for \mathbf{K}_A by replacing $@_0 \rightarrow (\lambda \wedge @_0)$ with the tautology $(@_0 \wedge \chi) \rightarrow (\lambda \rightarrow (@_0 \wedge \chi))$, where χ is some complicated logical truth. It is extremely plausible that there is some complicated logical truth χ such that both $(@_0 \wedge \chi) \rightarrow (\lambda \rightarrow (@_0 \wedge \chi))$ and $(@_0 \wedge \chi)$ are known but $(\lambda \rightarrow (@_0 \wedge \chi))$ is not known for the simple reason that no one knows both $(@_0 \wedge \chi) \rightarrow (\lambda \rightarrow (@_0 \wedge \chi))$ and $(@_0 \wedge \chi)$ at the same time and then goes on to perform a *modus ponens*. (We don’t care much whether the counterexample will involve different agents who fail to combine their knowledge or a single agent who fails to combine knowledge had at different times, or knows both at one time but dies or gets distracted before performing a *modus ponens*.)

In conclusion, there is no hope for a normal logic of apriority given Definition 2, nor even for the individual components of normality, \mathbf{K}_A and \mathbf{NEC}_A .

5.3. The logic of a priority in the absence of definitions

We have observed some surprising consequences of two natural definitions of a priority, Definition 1 and Definition 2. But we have given no reasons (nor will we give any reasons) for thinking that the correct definition of a priority must be one of the two, or that there must be any definition at all. Suppose we reject both definitions, and we further give up on definitions altogether, and simply treat apriority as primitive. And suppose we only have the constraint that a priority is a modal-epistemological notion expressed by a sentential operator A . Does the logic of apriority look any better from that perspective?

In fact, it doesn’t, and it may even look worse: now we have no guarantee that even \mathbf{T}_A is valid. (Of course, we have been given so little guidance on how to think about a priority that we know of no way of constructing a counterexample to its

validity.) As for \mathbf{K}_A and \mathbf{NEC}_A , the reader may have noticed that every one of our arguments against these in §5.2 made use of only one direction of the equivalence Definition 2 would give us (namely $A\phi \rightarrow (\phi \wedge \Diamond K_A\phi)$), and in fact they only required a principle even weaker than that direction: $A\phi \rightarrow \Diamond K_A\phi$. And this principle we have as long as we assume that apriority is a modal-epistemological notion. As long as being a priori entails the possibility of being known a priori, the logic of a priori can be shown not to be normal, and not to even have either of the individual components of normality, \mathbf{K}_A and \mathbf{NEC}_A .

6. Being in a position to know

Analytic epistemologists these days make frequent use of the ideology of ‘being in a position to know’.¹³ Being in a position to know is thought to be factive, and, further, it is thought that one can be in a position to know things one does not in fact know. This last thought makes attractive the thought that being in a position to know might have a normal logic. Even granting that the idealization of logical omniscience is false (so \mathbf{NEC}_K fails) and that knowledge is not closed under logical consequence, and in particular not under *modus ponens* (so \mathbf{K}_K fails), it may nevertheless seem *prima facie* plausible—and indeed various standard applications of the notion require it to be true—that one is in a position to know any logical truth, and that one is in a position to know whatever follows by *modus ponens* from what one is in a position to know. Thus one might expect, or at least hope, that the logic of being in a position to know is at least as strong as \mathbf{KT} .

It is furthermore natural to think that the notion has connections to metaphysical modality, even though being in a position to know may not be straightforwardly definable in terms of \Box and K . Purveyors of this ideology are certainly not working with definitions analogous to either Definition 1 or Definition 2. There may be a cow in the distance that you have no hope of seeing given your actual visual apparatus and actual supply of visual aids but that you could have seen with a supercharged visual apparatus or, more mundanely, with very powerful binoculars. In such a case it may be that you are not in a position to know that there is a cow in the distance even though it is true, and it is metaphysically possible for you to know that there is a cow in the distance. Nevertheless, it is plausible that an entailment runs in the opposite direction: just as someone who cannot speak is not in a position to speak, someone who cannot know is not in a position to know. Contraposing, it is plausible that if one is in a position to know, then it is possible for one to know:

$$K^P\phi \rightarrow \Diamond K\phi,$$

where K^P is to be read as ‘one is in a position to know that’. This principle just amounts to the claim that being in a position to know being a modal-epistemological notion.

Supposing that being in a position to know is a modal-epistemological notion, how much of the hypothesis that its logic is at least as strong as \mathbf{KT} can we save?

¹³ For example, Williamson (2000: 95) defines his central notion of *luminosity* in terms of that of being in a position to know.

As in the case of apriority, it not easy to think of counterexamples to axiom **T** K_P , and of course one could secure **T** K^P by brute force—for example, by introducing a definition analogous to Definition 2 that uses a restricted possibility operator. We shall leave **T** K^P alone.

However, it should come as no great surprise that the logic of being in a position to know is not normal, and that there are counterexamples to both **K** K^P and **NEC** K^P . In fact, exactly the same counterexamples work (*mutatis mutandis*) here as in the case of apriority. One may be in a position to know both $(@_0 \rightarrow (\lambda \wedge @_0))$ and $@_0$, while not knowing $(\lambda \wedge @_0)$. Then, by **NecFac** and our assumption that K^P expresses a modal-epistemological notion, familiar reasoning will show that

$$K_P(@_0 \rightarrow (\lambda \wedge @_0)) \rightarrow (K_P@_0 \rightarrow K_P(\lambda \wedge @_0))$$

is a false instance of **K** K^P . As for **NEC** K^P , by a familiar form of argument, some sentence of the form $\neg \diamond K_P(@_0 \wedge \lambda)$ will deliver a counterexample.

Here and elsewhere there are plenty of additional examples that we could use to argue that the logics of the modal-epistemological notions are not normal if we could help ourselves to the assumption that knowing a conjunction necessarily implies knowing each conjunct, in the sense of the necessitated distribution principle

DIST: $\square((K(\phi \wedge \psi) \rightarrow (K\phi \wedge K\psi)))$.

Although we ourselves do not find **DIST** secure enough to assume, there is some plausibility to it even when the assumption of logical omniscience is given up: perhaps, as Williamson (2000: 283) conjectures, ‘knowing a conjunction *constitutes* knowledge of its conjuncts’. Suppose for the moment, then, that **DIST** is valid. Under this supposition, a variety of fairly pedestrian counterexamples to the normality of the logic of K^P come into view. For example, there are many cases in which one is in a position to know that p while also being in a position to know that one does not know that p . For example, suppose that you are wondering about whether there is a prime number between j and k , that you don’t believe that there is such a prime (nor do you believe that there is not), and that, reflecting on the situation, you come to know that you don’t believe that there is a prime between j and k , and, since you know that knowledge requires belief, you also come to know by deduction that you don’t know that there is such a prime ($K\neg K\pi$). Suppose further that there is easy proof of the existence of such a prime just within your reach, and you would hit upon it if you gave the issue just a little more thought. You could easily know that there is such a prime—so easily, in fact, that you are in a position to know so ($K^P\pi$). Since knowing entails being in a position to know, $K^P\pi \wedge K^P\neg K\pi$ is true. Now suppose for a contradiction that the logic of K^P is normal. It follows that

$$(K^P\pi \wedge K^P\neg K\pi) \rightarrow K^P(\pi \wedge \neg K\pi),$$

being an instance of \mathbf{AGL}_{K^P} , is valid, and so $K^P(\pi \wedge \neg K\pi)$ is true. By the modal-epistemological property of K^P , $\Diamond K(\pi \wedge \neg K\pi)$ is true. By **DIST**, $\Diamond(K\pi \wedge K\neg K\pi)$ is true, but $\Diamond(K\pi \wedge K\neg K\pi)$ is inconsistent with **NecFac**.

7. Objections and replies

Some of our arguments relied on a combination of a conception of knowledge as a relation to propositions and a conception of propositions as non-coarse, so that, for example, necessarily equivalent valid sentences do not all express the same proposition. We also, in effect, made the assumption that it is fairly easy for a single proposition to be known in a variety of possible situations. In talking about knowledge in this way, we were not being oblivious to the phenomenon of ‘guise’ or ‘mode of presentation’ sensitivity that has figured prominently in discussions of knowledge and belief ascriptions in the philosophy of language since Kaplan’s classic paper ‘Quantifying In’ (1968) (albeit less prominently in work on epistemic logic). For example, everything we have said is consistent with the view that the most natural relation in the vicinity of our ‘knowledge’ talk is a three-place relation R between an agent, a proposition, and a guise. This view is consistent with the further view that knowledge is the two-place relation that results from existentially generalizing over the guise argument of R (as in Salmon’s [1986] theory of belief ascriptions).¹⁴ Nevertheless, some may react to our discussion by claiming that it presupposes either an incorrect conception of propositions or an incorrect conception of epistemic operators, according to which they operate on propositions. In this section, we shall indicate why we doubt that such maneuvers will help to reinstate interesting—in particular, normal—logics for the modal-epistemological notions we have discussed.

The least subtle way to try to undermine our arguments is to opt for a coarse-grained conception of propositions, according to which propositions that are necessarily equivalent are identical, and to insist that to know that p is just to know, under some guise or other, the proposition that p . We will call the advocate of this view the ‘coarse-grainer’. This would threaten some of our anti-normality arguments concerning apriority and being in a position to know, which assumed the distinctness of certain necessarily equivalent propositions:¹⁵ for the coarse-grainer, $@_0$ and $@_0 \wedge \lambda$ express the same proposition, so the coarse-grainer could reply that, in assuming that the proposition expressed by $@_0$ is known, we, by Leibniz’s law, *ipso facto* assumed that the proposition expressed by $@_0 \wedge \lambda$ is known as well. But the coarse-grainer’s objection is both ineffective and self-defeating.

¹⁴ That is to say, $K(x, p)$ is defined as $\exists gR(x, g, p)$, where g is a guise variable. There are contextualist versions of this idea, according to which the range of guises being generalized over varies with context: see Chalmers (2011) for a recent example.

¹⁵ In the anti- $\mathbf{K}_K/\mathbf{K}K^P$ arguments λ could be, for all we assumed, both contingent and valid, but then, of course, the proposition expressed by λ would have to be true in at least one world other than the actual one. In fact, we do think there are, or at least could be, such valid sentences, but we do not want to lean too heavily on that possibility, since we have far more powerful ammunition against the coarse-grainer. (Just to give a flavor of the kind of thing we have in mind, consider $@_1$, a 0-place modal operator that expresses, in each world w in which it is used, the disjunction of the world proposition of w and the proposition that there are talking donkeys.)

First, it is obviously ineffective against our arguments against \mathbf{T}_A , \mathbf{AGL}_A , and \mathbf{K}_A on Definition 1, but moreover it is ineffective against close variants of the arguments that it would undermine if it had anything going for it. For consider variants of the arguments against \mathbf{K}_A and \mathbf{NEC}_A that replace $@_0 \wedge \lambda$ with $(@_0 \vee \phi) \wedge \lambda$, where λ is as before and ϕ is interpreted as the proposition that no one knows anything. Clearly the proposition expressed by $@_0$ is not necessarily equivalent to the proposition expressed by $(@_0 \vee \phi) \wedge \lambda$. Just as clearly, the latter is unknowable unless actually known, so it will do the job.

Second, to recapitulate Yli-Vakkuri (2016: 827), the objection is self-defeating because the coarse-grainer can only secure a normal logic for A and K^P by accepting the absurd conclusions that every truth is a priori and that one is always in a position to know every truth. Recall that the proposition expressed by $\phi \leftrightarrow @\phi$ is necessarily equivalent either to the proposition expressed by ϕ or to the proposition expressed by $\neg\phi$, depending on which is true. According to the coarse-grainer, then, the proposition expressed by $\phi \leftrightarrow @\phi$ is either the proposition expressed by ϕ or the proposition expressed by $\neg\phi$, depending on which is true. By \mathbf{NEC}_A , $A(\phi \leftrightarrow @\phi)$ is valid, and therefore true. Since $\phi \leftrightarrow @\phi$ expresses either the same proposition as ϕ or $\neg\phi$ (whichever is true), it follows that $A\phi \leftrightarrow \phi$ is true. By a similar argument, it follows that $K^P\phi \leftrightarrow \phi$ is true given \mathbf{NEC}_{K^P} . And, as if this wasn't bad enough, the coarse-grainer is, entirely independently of \mathbf{NEC}_A and \mathbf{NEC}_{K^P} , committed something dangerously close to omniscience. If he accepts logical omniscience, then he accepts the validity of $K(\phi \leftrightarrow @\phi)$ and therefore of $K\phi \vee K\neg\phi$. (On his view, logical omniscience implies omniscience!) If he does not accept logical omniscience, he nevertheless accepts

$$K(\phi \leftrightarrow @\phi) \rightarrow (K\phi \vee K\neg\phi),$$

and by that principle, one knows any truth whatsoever simply by knowing a platitude of the form $\phi \leftrightarrow @\phi$. So much for the coarse-grainer.

Of course, actual lovers of coarse-grained propositions are typically not coarse-grainers in the sense defined above: typically they do not say that for it to be known that p is for the proposition that p to be known under some guise or other. Typically they acknowledge, for example, that one can know that p while failing to know that q even when p and q are necessarily equivalent, and therefore identical by their lights, and they have some complicated story about the semantics of epistemic operators that they hope will make this *prima facie* violation of Leibniz's law look less offensive to reason. Here we only wanted to make vivid how disastrous a coarse-grained view of propositions would be if offered as a way out of some of the logical difficulties we have pointed out without any compensating adjustments in one's semantics of epistemic operators. (We will consider such adjustments below.)

Alternatively, one might consider going in precisely the opposite direction. Consider an approach that makes propositions extremely fine-grained, perhaps by somehow incorporating guises or modes of presentation into them. To give a flavor of how this would work, suppose that, as Leibniz thought, this is the best possible world. Suppose further that, with complete logical security, one of us writes down $@_0$,

thereby coming to know the proposition expressed by $@_0$. Consider, meanwhile, a nearby world w in which someone driven by Leibnizian optimism writes down:

(δ) Dthat(the best possible world) obtains.

On various fairly coarse-grained views, (δ) in w expresses the proposition that $@_0$ actually expresses, but on the fine-grainer's view, these are distinct propositions, because, for example, one is about betterness and the other is not, or one presents itself via a guise that incorporates betterness (or a representation of betterness) and the other does not, or something along such lines.¹⁶

Once one is in the grip of this picture, one might well begin to wonder whether the proposition actually expressed, for example, by $@_0 \rightarrow (\lambda \wedge @_0)$ could be known in any world other than the actual one. Knowing it might require acquaintance with the very special mode of presentation by which the proposition expressed by $@_0$ presents itself only in the actual world. And if the proposition expressed by $@_0 \rightarrow (\lambda \wedge @_0)$ is not known in any other world, that will make trouble for some of our arguments.

Recall, however, that we have already presented arguments against both \mathbf{K}_A and \mathbf{K}_{K^P} that do not require the assumption that any proposition is known in any non-actual world. So the fine-grained approach is not a promising strategy for defending either \mathbf{K}_A or \mathbf{K}_{K^P} . Furthermore, our arguments against both \mathbf{NEC}_A and \mathbf{NEC}_{K^P} are entirely untouched by a move to a fine-grained conception of propositions. After all, what made trouble for \mathbf{NEC}_A and \mathbf{NEC}_{K^P} , was not knowledge of propositions expressed in the actual world in other worlds, but the *lack* of knowledge, in the actual world, of propositions expressed in the actual world (which, in turn, implied that those propositions are also not known in any other world). Finally, our argument against \mathbf{AGL}_A in the context of Definition 1 are also entirely untouched by fine-graining.

We will now turn to a third kind of theorist, who endorses *hypersensitivism*, the view that epistemic operators are sensitive not only to the propositions expressed by their operands but also to some other feature of their semantic profiles—perhaps Kaplanian character or Chalmers' primary intension. One natural view of this kind is that the objects of knowledge are pairs $\langle I_1, I_2 \rangle$ of a primary intension I_1 and a secondary intension (i.e., a proposition) I_2 .¹⁷ On such a view, the semantic value of a sentence is the ordered pair of its primary intension and its secondary intension, and interpretations are reconceived as assignments of ordered pairs of primary and secondary intensions to the atomic sentences.¹⁸ Of course, the details of

¹⁶ For those who, like us, prefer to think of possible worlds as propositions or states of affairs, to be generalized over using second-order quantifiers, there are generalizations of Kaplan's 'dthat' to other types that could be used here. A 'dthat' of type $(t \rightarrow t) \rightarrow t$, applied to $\lambda p\phi(p)$, where $\phi(p)$ is

$$\begin{aligned} & \diamond p \wedge \forall q(\Box(p \rightarrow q) \vee \Box(p \rightarrow \neg q)) \\ & \wedge \forall r((\diamond r \wedge \forall q(\Box(r \rightarrow q) \vee \Box(r \rightarrow \neg q))) \wedge r \neq p) \rightarrow p > r), \end{aligned}$$

where $>$ expresses betterness, works.

¹⁷ See Chalmers (2006, 2011).

¹⁸ If you can quantify in with propositional quantifiers, you need a story about how propositional quantification interacts with epistemic operators, and this will not be straightforward. Again, Kaplan (1968) might be a good place to start.

hypersensitivism could be worked out in ever so many different ways, but it is good to have a rough picture of at least one natural way of filling them in before we press on.

There are two varieties of hypersensitivism that are of no interest to us. The views of theorists who think of epistemic operators as covert metalinguistic predicates—who, in effect, draw a line in the sand at Quine’s first grade of modal involvement for epistemic discourse¹⁹—are of no interest to us. On such views one should not expect these operators to have normal modal logics, and their defenders may well take our arguments to support their views. Hypersensitivist views according to which apriority and being in a position to know are not modal-epistemological notions are also of no interest to us, since the topic here is modal epistemology.

The forms of hypersensitivism that do fall within our purview might seem ideally suited to resist our arguments. After all, our entire discussion so far has been framed in terms of assignments of *propositions* to atomic sentences. But, for the hypersensitivist, validity is truth under all assignments of semantic values that are *not* propositions—for example, all assignments of pairs of primary and secondary intensions—to atomic sentences. It might seem, then, that none of our arguments will carry over into a hypersensitivist framework for theorizing about validity.

In fact, every one of our arguments carries over. We will begin with what may seem to be the most difficult case: our argument against the **T** axiom for apriority on Definition 1. The key observation there was that there was an assignment, call it ‘*a*’, of propositions to atomic sentences that assigns a proposition, call it ‘*p**’, that is actually false but possibly known a priori to ϕ . $\Diamond K_A \phi \rightarrow \phi$ would then be false on *a*, and, by Definition 1, **T**_{*A*} would be false on *a*, so—provided that *a* is an interpretation—**T**_{*A*} would not be not valid. According to the hypersensitivist, *a* is not an interpretation. But no matter. To see why the hypersensitivist too must give up **T**_{*A*} on Definition 1, consider, for concreteness, the Chalmers-inspired proposal described above. (In fact, nothing here turns on the details of any hypersensitivist proposal, but the general problem that the hypersensitivist faces is easier to make vivid by considering a proposal than by reasoning in complete abstraction from the details of any hypersensitivist views.) And consider again the world *w* in which, from our non-hypersensitivist point of view, *p*, which is actually false, is known a priori. From the hypersensitivist’s point of view the thing that is known a priori in *w* is not *p** but $\langle I_1, p^* \rangle$, where *I*₁ some primary intension. Now consider a hypersensitivist interpretation *H* that assigns $\langle I_1, p^* \rangle$ to ϕ . Clearly ϕ is false on *H* because *p** is false, and $\Diamond K_A \phi$ is true on *H* because $\langle I_1, p^* \rangle$ is known a priori in *w*, so is possibly known. It follows that $\Diamond K_A \phi \rightarrow \phi$ is false on *H*. The key features of hypersensitivism on which this argument depends are just these: first, that a sentence is true iff the proposition incorporated into its hypersensitive semantic value is true and, second, that $\Box \phi$ is true iff the proposition incorporated into the hypersensitive semantic value of ϕ is necessarily true.

We are not going to provide a translation manual, but it should be clear that all of our arguments can be translated into hypersensitivist arguments that work provided that the original arguments work within our own theoretical framework. As long as (i) those semantic values incorporate (or determine) propositions—as they must, because otherwise \Box would have nothing to be sensitive to—and (ii) a sentence is true as

¹⁹ See Quine (1955).

evaluated at a world w on an assignment of a semantic value V to it iff the proposition contained in (or determined by) V is true in w , the hypersensitivist translations of our arguments will work without a hitch.

But it is also worth noting that many of our arguments do not even require any such translation scheme to be acceptable to the hypersensitivist. After all, many of them simply turn on assumptions about which sentences are true. None the anti-normality arguments that relied only on the entailment of the possibility of knowledge by the relevant modal-epistemological notion essentially relied on any assumptions about the nature of interpretations or about the semantics of epistemic operators. At best the hypersensitivist could hope to resist our argument against \mathbf{T}_A on Definition 1, but even there, as we have just seen, things end badly for him.

One should not be very surprised that hypersensitivism did nothing to restore a substantial logic of apriority or of being in a position to know. After all, little else separates the fine-grainer from the hypersensitivist other than that the latter claims that the objects of knowledge are not propositions. But one should not expect such taxonomic decisions to have much bearing on the logic of the modal-epistemological notions, and the impotence of fine-graining in the face of our arguments has already been demonstrated.

The final objection we will consider runs along the following lines: it is a harmless idealization to suppose that we have logical omniscience, and it's not a big step from there to suppose that every logical truth is known a priori. On that supposition, there are no counterexamples to \mathbf{NEC}_A or to \mathbf{NEC}_{KP} . (Of course this only gives a contingent status to the principles that all logical truths are a priori and that we are in a position to know all logical truths, but we already signed up for the contingency of some logical truths in admitting either $@$ or $@_0$ as a logical constant.) Our response to this objection is that the slogan that logical omniscience is a harmless idealization must be used with some care, and this objection does not use it with appropriate care. After all, assuming that there is no God, it is almost certainly false that any thinking being is logically omniscient. And while a falsehood might be harmless to assume in some theoretical context, it would be absurd to suppose that a falsehood that is harmless to assume in some theoretical context is harmless *tout court*. If one is trying to figure out whether anyone is logically omniscient, the assumption of logical omniscience is not harmless. Relatedly, if one is trying to figure out how the contours of apriority relate to the contours of what is a priori known and what it is possible to know a priori, a logical omniscience assumption is completely misplaced.

8. Concluding remarks

What are the take-home lessons of these logical reflections? One such lesson is *not* that the notions of apriority and being in a position to know must be discarded. Even if one or both of these notions is or too vague or inchoate to be theoretically useful, the logical reflections above will only constitute a part of the case for that conclusion.²⁰ That said, the discovery that the logic of a theoretically central epistemic operator is extremely weak is far from insignificant. It shows that certain common

²⁰ In fact, we are inclined to think that one of them (a priority) deserves to be discarded, but we make that case elsewhere. See Yli-Vakkuri and Hawthorne (2018: ch. 6, n. 46).

forms of inference do not have the backing of logic. For example, philosophers tend to assume without argument that the conjunction of two a priori truths will itself be a priori. We have discovered, in effect, that not only is this principle is not logically valid, but, provided that a priority entails the possibility of a priori knowledge, it is in fact false. If we are to continue to theorize about apriority by means of a sentential operator, we should self-consciously disavow a wide range of seemingly innocent assumptions and examine things case by case. And the same goes for being in a position to know. One may of course prefer to give up the idea that either notion entails the possibility of knowledge, and try to justify a stronger logic in a rather different theoretical framework, but, as we said at the outset, that is a topic we postpone for another occasion.

References

- Almog, Joseph, John Perry, and Howard Wettstein (1989). *Themes from Kaplan*. Oxford: OUP
- Bacon, Andrew, John Hawthorne, and Gabriel Uzquiano (2017). ‘Higher-Order Logic and the Prior-Kaplan Paradox’. In Yli-Vakkuri and McCullagh 2017.
- Chalmers, David (2011). ‘Propositions and Attitude Ascriptions: A Fregean Account’, *Noûs*, 45: 595–639
- Chalmers, David (2006). ‘The Foundations of Two-Dimensional Semantics’, in Manuel Garcia-Carpintero and Josep Macià, eds., *Two-Dimensional Semantics*. Oxford: OUP.
- Chalmers, David (2011). ‘Actuality and Knowability’, *Analysis*, 71: 411–19,
- Chalmers, David and Brian Rabern (2014). ‘Two-Dimensional Semantics and the Nesting Problem’, *Analysis*, 74: 210–224
- Crossley, J. N., and Humberstone, L. (1977). ‘The logic of ‘actually’, *Reports on Mathematical Logic*, 8: 11–29.
- Evans, Gareth (1979). ‘Reference and Contingency’, *The Monist*, 62: 161–89.
- Fritz, Peter (2014). ‘What is the correct logic of necessity, actuality and apriority?’ *Review of Symbolic Logic*, 7: 385–414.
- Fritz, P. (2013). ‘A logic for epistemic two-dimensional semantics’, *Synthese*, 190: 1753–70.
- Goodman, J. (2017). ‘Reality is not structured’, *Analysis*, 77: 43–53,
- Kaplan, D. (1968). ‘Quantifying In’, *Synthese*, 19: 178–214
- Kaplan, D. (1987). ‘Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals’, in J. Almog, J. Perry and H. Wettstein 1989.
- King, J. (2007). *The Structure and Nature of Content*. Oxford: OUP.
- Quine, W. V. (1955). ‘Three Grades of Modal Involvement’, *Journal of Symbolic Logic*, 20:168–169.
- Rabinowicz, W. and K Segerberg (1994). ‘Actual Truth, Possible Knowledge’, *Topoi*, 13: 101–115.
- Restall, G. (2012). ‘A cut-free sequent system for two-dimensional modal logic, and why it matters’, *Annals of Pure and Applied Logic*, 163, 1611–1623.
- Salmon, N. (1986). *Frege’s Puzzle*. Cambridge, MA: MIT Press.
- Stalnaker, Robert C. (2012). *Mere Possibilities*. Oxford: OUP.
- Soames, S, (1989). ‘Direct Reference and Propositional Attitudes’. In Almog, Perry, and Wettstein 1989.
- _____. (2010). *What Is Meaning?* Princeton, NJ: Princeton University Press.
- Tarski, A. (1936). ‘On the Concept of Logical Consequence’. In *Logic, Semantics, Metamathematics*, ed. by J. Corcoran. Indianapolis, IN: Hackett, 1983.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: OUP.
- Williamson, T. (2013). *Modal Logic as Metaphysics*. Oxford: OUP.
- Yli-Vakkuri, J. (2017). ‘Epistemicism and Modality’, in Yli-Vakkuri and McCullagh 2017.
- Yli-Vakkuri, J. and J. Hawthorne (2018). *Narrow Content*. Oxford: OUP.
- Yli-Vakkuri, J. and M. McCullagh, eds., (2017). *Williamson on Modality*. London: Routledge.