

TENTH EXCURSUS

Constructing Epistemic Space

In *Meaning and Necessity* (1947), Carnap laid the foundations for much of the contemporary discussion of possible worlds and of intensional semantics. In particular, he developed a notion of ‘state-description’ that serves as the key to a linguistic construction of possible worlds. He also argued that every expression can be associated with an intension. This extensive modal and semantic project serves as the background for ‘Meaning and Synonymy in Natural Languages’, in which Carnap engages in the metasemantic project of determining what it is for a subject to use an expression with a given meaning. The modal and semantic projects deserve attention in their own right, however. In this excursus, I focus on analogs of the modal project within the current framework, while in the eleventh excursus, I focus on analogs of the semantic project.

Meaning and Necessity was published two decades after the *Aufbau*, and Carnap does not explicitly connect the projects. But as chapter 5 suggests, it is natural to draw a connection. The basic vocabulary of the *Aufbau* can be used to formulate a set of atomic sentences that characterizes the actual world. In effect, this characterization provides a state-description for the actual world. The definitional elements of the *Aufbau* provide a way to determine the truth of an arbitrary sentence, given the specification of truths in the basic vocabulary. This provides a way of evaluating the intension of a sentence given an arbitrary state-description. In the *Aufbau*, Carnap does not use his basic vocabulary to characterize state-descriptions for non-actual states of the world, but one could certainly do so in principle. In this way, one could use the materials of the *Aufbau* to construct the state-descriptions and intensions that are needed for the project of *Meaning and Necessity*.

Of course there are some differences of detail. Where the *Aufbau* uses an austere basic vocabulary (logic plus a basic relation) and a rich system of semantic rules (arbitrary definitions), *Meaning and Necessity* allows a rich basic vocabulary (atomic sentences containing arbitrary predicates and individual constants) and an austere system of semantic rules (logical relations between atomic and complex sentences). And where the *Aufbau* requires only extensional adequacy of its definitions, *Meaning and Necessity* appeals to a notion of ‘L-truth’ (truth in

virtue of the semantic rules of a language L) that Carnap says is akin to a notion of analyticity or necessity. Still, as discussed in chapter I, there is a nearby *Aufbau* project that requires something like analyticity of its definitions. This project can be used to ground a nearby *Meaning and Necessity* project that uses austere bases to define state-descriptions and rich definitional connections to define intensions.¹

Something like this construction is at the heart of contemporary linguistic constructions of possible worlds. In these constructions, metaphysical necessity plays the role that L -truth plays for Carnap, and expressions for fundamental objects and/or properties play the role of the basic vocabulary. One can thereby define possible worlds and intensions for expressions, with the key property that a sentence S is metaphysically necessary iff its intension is true in all possible worlds.

The scrutability project allows us to execute a related but quite different construction: a construction not of metaphysically possible worlds, but of epistemically possible scenarios. On this construction, apriority plays the role that L -truth plays for Carnap, and expressions in a generalized scrutability base play the role of the basic vocabulary. These scenarios can perform a crucial role in the analysis of knowledge, belief, and meaning.²

Intuitively, an epistemically possible scenario (or a *scenario* for short) is a maximally specific way the world might be, for all we know a priori. For example, we do not know a priori that gold is an element. For all we know a priori, we could be in a world in which gold is a compound. Correspondingly, there are many scenarios in which gold is a compound, as well as many scenarios in which gold is an element. This already suggests that epistemically possible scenarios are distinct from metaphysically possible worlds. On the usual understanding of metaphysical possibility, it is metaphysically necessary that gold is an element (given that it is actually an element), so there are no possible worlds in which gold is a

¹ An austere base might help with certain internal tensions in the *Meaning and Necessity* construction. Carnap says (p. 15) that where ' H ' and ' RA ' are predicates for 'human' and 'rational animal' respectively, $\forall x(Hx \equiv RAx)$ is L -true, on the grounds that ' H ' and ' RA ' are synonymous. But on Carnap's official definitions (pp. 3–4 and pp. 9–10), L -truth requires truth in all state-descriptions, state-descriptions allow arbitrary recombinations of atomic sentences or their negations, and sentences of the form Hc and RAc (for any constant c) are atomic. On these definitions, there will be state-descriptions containing both Hc and $\sim RAc$, so that $\forall x(Hx \equiv RAx)$ will not be L -true. More generally, the rich atomic language along with free recombination has the consequence that all sorts of apparently analytic sentences will be false in some state-description. An *Aufbau*-style version of the *Meaning and Necessity* project with an austere base vocabulary and a rich system of definitions corresponding to a rich notion of L -truth would help to avoid these problems. Alternatively, one could retain the rich base vocabulary but impose a constraint of L -consistency on recombinations of atomic sentences.

² Many of the issues discussed below are elaborated at much greater length in 'The Nature of Epistemic Space'.

compound. We have to understand epistemically possible scenarios in different terms.

To construct scenarios, we can start by stipulating that S is *deeply epistemically necessary* when S is conclusively a priori (see the eighth excursus), and that S is *deeply epistemically possible* when $\sim S$ is not deeply epistemically necessary. Note that this stipulative notion of deep epistemic possibility differs from the ordinary notion of what is epistemically possible for a subject, both because it is idealized (all mathematical theorems are epistemically necessary and their negations are epistemically impossible, even if no one has proved them) and because it does not depend on what a given subject knows ('I am not conscious' is deeply epistemically possible, even though I am introspectively certain that I am conscious). For convenience I will abbreviate 'deeply epistemically possible' as 'e-possible' and sometimes as 'epistemically possible' in what follows, but this should not be confused with the non-idealized subject-relative notion. For now we can restrict attention to context-independent sentences S (specifically, epistemically invariant sentences, as discussed in E₃), although the definition can naturally be extended to define the e-possibility of a context-dependent sentence S in a context.

We can also stipulate as before that a sentence G is *epistemically complete* when G is e-possible and there is no H such that $G \& H$ and $G \& \sim H$ are e-possible. For example, if a true sentence conjoining all sentences in *PQTI* a priori entails all truths, then that sentence will be epistemically complete.³ We can also say that two epistemically complete sentences G_1 and G_2 are *equivalent* when $G_1 \rightarrow G_2$ and $G_2 \rightarrow G_1$ are both e-necessary.

Given Generalized Scrutability, there will be a compact vocabulary that can be used to specify epistemically complete sentences corresponding not just to the actual world but to arbitrary epistemic possibilities. The thesis says that there is a compact class C of sentences such that for all e-possible S , S is a priori scrutable from some e-possible subclass C' of C (where a class of sentences is e-possible iff its conjunction is e-possible). From here, one can argue that for all e-possible S , S is e-necessitated by some epistemically complete sentence in C . In effect, C provides an array of epistemically complete sentences akin to *PQTI*, each of which corresponds to a highly specific epistemic possibility.

³ Indeterminacy raises a few complications. Its treatment will depend on the issue, discussed in the first excursus, of whether the vagueness of epistemic necessity goes along with the vagueness of truth or with the vagueness of determinate truth. If we take the latter route, one should say that S is e-possible when $\sim \text{det}(S)$ (rather than $\sim S$) is not e-necessary. Then when G is the conjunction of sentences in an a priori scrutability base and H is indeterminate, $G \rightarrow \text{indet}(H)$ will be e-necessary and neither $G \& H$ nor $G \& \sim H$ will be e-possible. If we take the former route, then under plausible assumptions, each of the two latter sentences will be indeterminately e-possible, but it will be determinately false that both are e-possible. So either way, G will be epistemically complete.

We can then identify scenarios with equivalence classes of epistemically complete sentences in the vocabulary of a generalized scrutability base. Given a scenario w , any sentence D in the corresponding equivalence class is a *canonical specification* of w . For a context-independent sentence S , a scenario w *verifies* S when S is a priori scrutable from a canonical specification of w : that is, when $D \rightarrow S$ is a priori, where D is an epistemically complete sentence corresponding to w . Given a Generalized Scrutability thesis for context-dependent sentences, one can likewise say that w verifies an arbitrary sentence S in a context when S is a priori scrutable from a canonical specification of w in that context.⁴

Given the above, this construction ensures the crucial principle of Plenitude (along with a number of other principles discussed in ‘The Nature of Epistemic Space’): S is e-possible iff there exists a scenario that verifies S . Likewise, for context-dependent sentences S , S is e-possible in a context iff there exists a scenario that verifies S in that context. For example, the e-possible sentence ‘Gold is a compound’ will be verified by many scenarios: intuitively, these are scenarios in which a compound gives rise to the appearances that we associate with gold. Likewise, the e-possible sentence ‘Hesperus is not Phosphorus’ will be verified by many scenarios: intuitively, these are scenarios in which the relevant bright objects in the evening and morning skies are distinct.

For any subject s (at time t in world w), there will be one scenario that is *actualized* for s (at t in w). This will be a scenario corresponding to an epistemically complete sentence (such as $PQTI$) that is true of w for s at t : specifying the objective character of w (perhaps using P , Q , and T) and the position of s and t within it (using I). The scrutability thesis tells us that a sentence S will be true for s (at t in w) if and only if it is verified by the scenario that is actualized for s (at t in w). Different scenarios will be actualized for different subjects (even within the same world), as reflected in the fact that the I component of $PQTI$ will be different for different subjects.

Scenarios as defined have many applications. They can be used to help understand talk about skeptical scenarios in epistemology, and more generally to serve as ‘epistemically possible worlds’ in the analysis of knowledge and belief. They can be used to help understand the objects of subjective probability: it is arguable that idealized subjective probabilities are in effect distributed over the space of scenarios. And perhaps most importantly, they can play a key role in the analysis of meaning and content, helping to analyze Fregean notions of meaning and internalist notions of mental content.

⁴ Here the relevant sort of context-dependence is epistemic context-dependence. Primitive indexicals such as ‘I’ and ‘now’ can be taken to be epistemically context-independent. The required Generalized Scrutability thesis then requires a base of epistemically context-independent sentences while allowing context-dependent sentences in the dependent class. See E11 for more on this issue.

For these purposes (which are discussed further in the eleventh excursus and chapter 8), a central role is played by intensions: functions from scenarios to truth-values. The intension of S (in a context) is true at a scenario w if w verifies S (in that context), false at w if w verifies $\sim S$ (in that context), indeterminate at w if w verifies *indet*(S), and so on. The intension of S will be true at all scenarios iff S is a priori. So the intension of ‘Gold is an element’ will be true at some scenarios and false at others.

The intension so defined is a version of the primary or epistemic intension familiar from two-dimensional semantics. A sentence’s primary intension is its epistemic profile (chapter 1), mapping epistemically possible scenarios to truth-values. A sentence’s secondary intension is its modal profile, mapping metaphysically possible worlds to truth-values. A sentence S is a priori (epistemically necessary) iff its primary intension is true in all scenarios, and metaphysically necessary iff its secondary intension is true in all worlds.

When S is an a posteriori necessity, such as ‘Hesperus is Phosphorus’, its secondary intension will be true at all worlds, but its primary intension will be false at some scenarios. There will be some specifications of scenarios— $PQTI^*$, say—describing a scenario in which the objects visible in the morning sky (around the individual designated by ‘I’) are entirely distinct from the objects visible in the evening sky. If we discovered that we were in such a scenario, we would accept ‘Hesperus is not Phosphorus’. Likewise, conditional on the hypothesis that we are in such a scenario, we should accept ‘Hesperus is not Phosphorus’. So ‘Hesperus is not Phosphorus’ is conditionally scrutable from $PQTI^*$. The arguments earlier in this chapter then suggest that it is a priori scrutable from $PQTI^*$. So the primary intension of ‘Hesperus is Phosphorus’ will be false at a scenario specified by $PQTI^*$.

More generally, to evaluate the primary intension of a sentence S at a scenario w , one considers w as actual: that is, one considers the hypothesis that w actually obtains, or equivalently, the hypothesis that D is actually the case, where D specifies w . Here we can use conditional scrutability at least as a heuristic guide to a priori scrutability, and ask: conditional on the hypothesis that w is actual, should one accept S ? For example, conditional on the hypothesis that a Twin Earth scenario in which the oceans and lakes are filled by XYZ is actual, one should accept ‘Water is not H_2O ’. So the primary intension of ‘Water is H_2O ’ is false at this scenario.

By contrast, to evaluate the secondary intension of a sentence S at a world w , one considers w as counterfactual: that is, one considers counterfactually what would have been the case if w had obtained, or equivalently, if D had been the case, where D specifies w . Here as a heuristic we can ask: if w had been obtained (that is, if D had been the case), would S have been the case? In the case of a Twin Earth world, we can ask: if the oceans and lakes had been filled by XYZ, would

water have been H_2O ? Following Kripke and Putnam, the standard judgment about this subjunctive conditional is ‘yes’. (By contrast, the intuitive judgment about the indicative conditional ‘If the oceans and lakes are filled by XYZ, is water H_2O ’ is ‘no’.) If so, the secondary intension of ‘Water is H_2O ’ is false at the Twin Earth world.

We can also associate subsentential expressions with primary and secondary intensions. In general, a primary intension maps an epistemically possible scenario to extensions—objects for singular terms, properties for predicates, and so on—while a secondary intension maps metaphysically possible worlds to extensions. For example, intuitively the primary intension of ‘water’ picks out H_2O in the actual scenario and picks out XYZ in a Twin Earth scenario. Making this precise requires us first to make sense of the notion of objects within scenarios, which takes some work (as scenarios have so far just been constructed from sentences), but we can work with an intuitive understanding for present purposes.

The secondary intension of ‘Hesperus is Phosphorus’ reflects the fact that ‘Hesperus’ and ‘Phosphorus’ are *metaphysically rigid*, picking out the same entity—the planet Venus, their referent in the actual world—in all metaphysically possible worlds. By contrast, the primary intension of ‘Hesperus is Phosphorus’ (in paradigmatic contexts) suggests that ‘Hesperus’ and ‘Phosphorus’ are not *epistemically rigid*: they do not pick out the same entity in all epistemically possible scenarios (see E13 for much more here). If they did pick out the same entity in all scenarios, then ‘Hesperus is Phosphorus’ would be true in all scenarios and therefore a priori, which it is not. Rather, to a rough first approximation, ‘Hesperus’ picks out a bright object in a certain location in the evening sky in a given scenario, while ‘Phosphorus’ picks out a bright object in a certain location in the morning sky. In many scenarios, these two objects will be distinct.

This point generalizes to arbitrary names of concrete objects, suggesting that no name for a concrete object is epistemically rigid. By contrast, it is arguable that some names for abstract objects are epistemically rigid: for example, ‘0’ arguably picks out zero in all scenarios, and ‘identity’ arguably picks out the relation of identity in all scenarios. One can similarly hold that numerous predicates—perhaps ‘conscious’, ‘causes’, and ‘omniscient’, among many others—are epistemically rigid, having the same property as extension in all scenarios.

Making the notion of epistemic rigidity precise is tricky, in part because we have not yet formally populated scenarios with objects, and in part because it is unclear that there is a coherent general notion of trans-scenario identity: that is, of what it is for entities in two different scenarios to be the same entity. But a useful intuitive gloss on the notion is that an epistemically rigid expression is one that expresses an epistemically rigid concept, and that an epistemically rigid concept is one whose extension we can know a priori. For example, there is an intuitive sense in which we cannot know what water is a priori, and in which we

cannot know what Hesperus is a priori, but in which we might be able to know what zero is a priori or what identity is a priori. That is roughly the sense at play here. I discuss this and many other issues about epistemic rigidity at greater length in the fourteenth excursus.

When an expression is epistemically rigid and also metaphysically rigid (metaphysically rigid *de jure* rather than *de facto*, in the terminology of Kripke 1980), it is *super-rigid*. In this case, the expression will pick out the same entity in all scenarios and all worlds. For example, it is plausible that ‘o’ picks out zero in all scenarios and all worlds. As with epistemically rigid expressions, there are plausibly no super-rigid expressions for concrete objects, but there are plausibly super-rigid expressions for some abstract objects and properties.

What is the relationship between epistemically possible scenarios and metaphysically possible worlds? As I have discussed them so far, these are independent sorts of entities. But it is common to see a close relationship between them, modeling epistemically possible scenarios as centered metaphysically possible worlds. The existence of super-rigid expressions, which function to pick out the same entities in scenarios and in worlds, helps us to explore the connections between these entities. They can also help us to construct scenarios nonlinguistically using worldly entities such as properties and propositions.

To analyze the correspondence between scenarios and worlds it is helpful to highlight two theses about super-rigid expressions that I discuss in the fourteenth excursus and chapter 8. The first is Super-Rigid Scrutability: all epistemically possible sentences S are scrutable from sentences including only super-rigid expressions and primitive indexicals (such as ‘I’ and ‘now’).⁵ The second is an Apriority/Necessity thesis: when a sentence S contains only super-rigid expressions, S is a priori iff S is necessary. The first thesis is supported by the character of the scrutability bases we arrive at and by general considerations concerning the scrutability of sentences containing epistemically nonrigid expressions. The second thesis is supported by the observation that paradigmatic Kripkean a posteriori necessities all appear to involve epistemically nonrigid expressions.

Super-Rigid Scrutability (and the considerations that support it) suggests that a generalized scrutability base need contain only certain basic super-rigid sentences (that is, sentences containing only super-rigid expressions) and certain indexical sentences such as ‘I am F_1 ’ and ‘Now is F_2 ’, where F_1 and F_2 are predicates containing only super-rigid expressions. Given this, scenarios can be identified with epistemically complete sentences of the form $D \& I$, where D is a complex super-rigid sentence (conjoining basic sentences) and I is a conjunction of indexical sentences as above. D will say roughly that there exist objects bear-

⁵ This has roughly the strength of a generalized scrutability thesis, but I omit ‘Generalized’ here and later for ease of discussion.

ing certain specific properties and relations, and I will attribute certain specific properties and relations of oneself and the current time.

D will express a complex Russellian proposition p , containing properties and relations (perhaps along with other abstract objects) as constituents, connected by logical structure. This proposition is quite reminiscent of a possible world. It is common to regard possible worlds as complex Russellian propositions, specifying the distribution of certain basic properties and relations over objects. If the Apriority/Necessity thesis is correct, then D will be metaphysically possible (as it is super-rigid and epistemically possible), so p will be metaphysically possible. In this case, p will correspond to a metaphysically possible world.⁶ If the Apriority/Necessity thesis is incorrect, then D may be metaphysically impossible, in which case p will correspond to a metaphysically impossible world. But either way p will correspond to a world-like entity involving the distribution of properties and relations over objects.

If a scenario w were specified by D alone, we could then identify the scenario with the world specified by p .⁷ Given the additional role of I , which says something like ‘I am F_1 and now is F_2 ’, we can instead identify the scenario with a *centered world*. Centered worlds are usually taken to be ordered triples of worlds, individuals, and times. For present purposes, we can take them to be ordered triples of a Russellian proposition p and properties ϕ_1 and ϕ_2 (corresponding to maximally specific properties possessed by the individual and the time respectively). For the scenario in question, we can take ϕ_1 and ϕ_2 to be the properties expressed by the predicates F_1 and F_2 respectively.

⁶ Strictly speaking, p will correspond either to a complete or an incomplete metaphysically possible world. Which is correct depends on the thesis of Super-Rigid Necessitation (an analog of Super-Rigid Scrutability): for any metaphysically possible sentence S , S is metaphysically necessitated by some sentence T including only super-rigid expressions. If this thesis (along with Apriority/Necessity) is true, the Russellian propositions in the text will be metaphysically complete (by analogy with epistemic completeness) and will specify full metaphysically possible worlds. If this thesis is false, these Russellian propositions may be metaphysically incomplete and will specify incomplete worlds (worlds without all details filled in), which in effect correspond to equivalence classes of metaphysically possible worlds.

Super-Rigid Necessitation will be false on certain haecceitist views (discussed in E16), on which a super-rigid specification of a world may underdetermine which objects are present in a world. It will also be false on certain quidditist views (discussed in 7.9 and E16), on which a super-rigid specification of a world may underdetermine which intrinsic properties are present in that world. Still, given Apriority/Necessity and Super-Rigid Scrutability, these propositions will correspond at least to incomplete metaphysically possible worlds. Given Super-Rigid Necessitation in addition, the propositions will correspond precisely to metaphysically possible worlds.

⁷ This only works when D is super-rigid. If non-super-rigid expressions are involved in the base, identifying scenarios with Russellian propositions will give the wrong results. For example, if ‘Hesperus’ and ‘Phosphorus’ are in the base, ‘Hesperus is such-and-such’ will always specify a Russellian proposition about Venus, as will ‘Phosphorus is such-and-such’. So ‘Hesperus is Phosphorus’ will come out true in all scenarios, even though it is not a priori.

If Super-Rigid Scrutability is true, we can always model scenarios as centered worlds in this way. If Apriority/Necessity is also true, scenarios will correspond to centered worlds where the worlds in question are always metaphysically possible.⁸

Take a paradigmatic a posteriori necessity, such as ‘Water is H₂O’ which is false in a Twin Earth scenario. Super-Rigid Scrutability entails that a Twin Earth scenario can be specified by an epistemically possible sentence *TE* & *I* conjoining a super-rigid sentence *TE* with indexical claims involving ‘I’ and ‘Now’. Here *TE* might involve a super-rigid specification of the microphysical and phenomenological character of the world, without using epistemically nonrigid terms such as ‘water’. Because *TE* is super-rigid, it can be used to specify a world: that is, a Russellian complex of objects and properties. Apriority/Necessity entails that *TE* is metaphysically possible, so that this world is a metaphysically possible world. *TE* metaphysically necessitates ‘Water is H₂O’, while *TE* & *I* epistemically necessitates ‘Water is not H₂O’. So the primary intension of ‘Water is H₂O’ is false at the Twin Earth scenario while its secondary intension is true at the Twin Earth world. But there remains a close correspondence between the scenario and the world. The difference in intensions for ‘Water is H₂O’ arises not so much due to differences between them as due to the difference between epistemic and metaphysical necessitation.

I accept Apriority/Necessity and Super-Rigid Scrutability. (Relatives of these theses play crucial roles in ‘The Two-Dimensional Argument against Materialism’.) So I hold that scenarios correspond closely to centered metaphysically possible worlds. In practice, many philosophers already use centered metaphysically possible worlds to model epistemically possible scenarios (this is a standard practice in the literature on subjective probability, for example). The analysis above can be seen as providing a partial grounding for this practice.

At the same time, some philosophers will reject Apriority/Necessity or Super-Rigid Scrutability. If these theses are rejected, the relation between scenarios and possible worlds becomes more complex.

First, consider a view that rejects Apriority/Necessity while retaining Super-Rigid Scrutability: for example, a theist view on which ‘There is an omniscient being’ is necessary but not a priori. This sentence *S* plausibly contains only super-rigid expressions. So on the view in question, *S* will be a counterexample to the Apriority/Necessity thesis. Correspondingly, *S* will be false in some

⁸ As before, if Super-Rigid Necessitation is false, scenarios will correspond to centered incomplete metaphysically possible worlds, or to equivalence classes of centered metaphysically possible worlds. If Super-Rigid Necessitation (along with the other two theses) is true, scenarios will correspond near-perfectly with centered metaphysically possible worlds. The only exception involves certain symmetrical worlds, where more than one centered world (centered on symmetrical counterparts) may correspond to the same scenario.

epistemically possible scenario. Given Super-Rigid Scrutability, there will be a sentence $D \& I$ specifying such a scenario, where D is super-rigid and I is indexical. D will then specify a world (a Russellian proposition constructed from properties): in effect, a world without an omniscient being. But on the view in question, this world will not be metaphysically possible. So the scenarios here will correspond to centered metaphysically impossible worlds.

Next, consider a view on which Super-Rigid Scrutability is false: for example, a type-B materialist view (discussed in the fourteenth excursus) on which phenomenal concepts are both conceptually primitive and epistemically nonrigid. Then corresponding expressions such as ‘consciousness’ are needed in a scrutability base, but one cannot know what these expressions pick out a priori: perhaps ‘consciousness’ picks out a neurophysiological property N empirically. In this case, scenario descriptions (such as $PQTI$) will not be decomposable into a super-rigid part and an indexical part. Even once indexicals are removed, some remaining expressions (such as the phenomenal expressions in Q) will be epistemically nonrigid, so that they will not map a priori onto properties. While the sentence will still correspond empirically to a complex of properties (such as N), the resulting property-involving world will not adequately model the epistemic properties of the scenario. (For example, it becomes hard to see how ‘Consciousness is N ’ can be false at a centered world.) In this case, we can either stay with a linguistic model of scenarios without invoking properties, or perhaps better, we can see scenarios as structures of Fregean senses, where we associate the primitives here with primitive Fregean senses that determine the properties that serve as their referents only a posteriori.

To summarize: if both Apriority/Necessity and Super-Rigid Scrutability are true, as on my view, scenarios correspond closely to centered metaphysically possible worlds. If the former but not the latter is true, scenarios still correspond closely to centered worlds, but the worlds in questions may be metaphysically impossible in some cases. If the latter is false, one has to break the close connection between scenarios and centered (property-involving) worlds, instead constructing scenarios from sentences or from Fregean propositions.

What about two-dimensional evaluation? On the current model, this is tied to epistemic possibilities concerning what is metaphysically possible. For example, someone might hold that it is epistemically possible that there is exactly one metaphysically possible world, and epistemically possible that there are enormously many metaphysically possible worlds. We could model this by a two-dimensional modal structure on which every epistemically possible scenario v is associated with a space of worlds that are metaphysically possible relative to v . If one holds (as I do) that Apriority/Necessity and Super-Rigid Scrutability are both a priori, then every scenario will be associated with a space of worlds such that there is a centered world for every scenario. Under slightly stronger assump-

tions (which I discuss in ‘The Foundations of Two-Dimensional Semantics’), one can use the same space of metaphysically possible worlds to play the role of the putative worlds for each of these scenarios. Under these assumptions one gets a simple ‘rectangular’ two-dimensional structure, with every scenario corresponding to the same space of worlds. Under weaker assumptions one gets a more complex structure, on which some scenarios may be associated with smaller or larger spaces of worlds, and on which some worlds may be metaphysically impossible (relative to the actual scenario).

Either way, one can use this two-dimensional structure to define two-dimensional intensions for sentences: given a scenario v and a world w that is possible relative to v , the two-dimensional intension of S will map (v, w) to the truth-value of S at w , conditional on the assumption that v is actual. These two-dimensional intensions are a version of the two-dimensional matrices familiar from two-dimensional semantics.

An important residual issue concerns the idealized notion of epistemic possibility that we started with. This works well in modeling the epistemic states of idealized agents, but not as well in modeling the epistemic states of non-ideal agents. For example, a non-ideal agent might disbelieve certain moral or mathematical truths, even though these truths are a priori, or the agent might simply fail to believe these. The current framework does not easily model these states, as a priori truths are true in all scenarios. To model these states, it would be useful to invoke a less idealized notion of epistemic necessity, and a correspondingly more fine-grained space of non-ideal scenarios.

One idea here is that one might understand non-ideal epistemic necessity in terms of a notion of analyticity rather than apriority (discussed in chapter 8) and construct fine-grained scenarios and fine-grained intensions from there. Then insofar as the relevant moral truths (for example) are a priori but not analytic, there will be scenarios where they are false. If so, we may have groups of fine-grained scenarios with the same natural truths and different moral truths. However, it is not clear that this sort of model will help with ignorance of logical truths: these truths are often taken to be analytic, and it is not obvious how best to model scenarios in which they are false. The understanding of non-ideal epistemic space remains open as a challenge for future work.⁹

⁹ Jens Christian Bjerring (2010, forthcoming) explores a number of models of non-ideal epistemic space, and demonstrates that there are serious difficulties in developing a model in which complex logical truths are false but simple logical truths are not. As a result, one may be left with a choice between logical omniscience (logical truths are true in all worlds) and triviality (any set of sentences determines a scenario).