Epistemic Rigidity and Super-Rigidity

In *Naming and Necessity*, Saul Kripke introduced the idea of a rigid designator: an expression that picks out the same thing in every possible world. He held that names are rigid designators: for example, 'Hesperus' picks out the same planet (Venus) in all possible worlds. Natural kind terms are also rigid designators: 'water' picks out the same kind, H_2O, in all possible worlds. By contrast, many descriptions are nonrigid: 'The greatest cricket player' picks out Bradman in our world, but it picks out someone else in a world where Bradman died in his youth.

Kripke's notion of rigidity is sometimes called *metaphysical rigidity*: an expression is rigid iff it picks out the same entity in all metaphysically possible worlds. It might more aptly be called counterfactual or subjunctive rigidity, as what really matters here is stability of reference in counterfactual or subjunctive contexts: if a Twin Earth situation had obtained, it still would have been the case that water is H_2O, and the liquid in the oceans and lakes would not have been water. But I will use the more standard term here.

In the tenth excursus, I introduced the parallel notion of *epistemic rigidity*. Epistemic rigidity and the related notion of super-rigidity play an important role in some parts of this book. These notions are related to but importantly distinct from the notion of non-Twin-Earthability discussed in chapter 7. In this excursus I discuss the notions in detail.

To a first approximation, an epistemically rigid expression is one that picks out the same thing in every epistemically possible scenario. As we saw earlier, 'water' is not epistemically rigid: it picks out H_2O in an Earth scenario and XYZ in a Twin Earth scenario. Similarly, names such as 'Hesperus' are not epistemically rigid: 'Hesperus' picks out Venus in the actual scenario, but in a scenario where a star (rather than a planet) is visible in the evening sky at the relevant location, 'Hesperus' will pick out that star. If we construe predicates as picking out properties, then a predicate such as 'hot' is not epistemically rigid: it picks out a property involving molecular motion in our scenario, but in a scenario where a different property X plays the role of heat (in causing experiences,

expanding metals, and so on), it will pick out X. The same goes for general terms such as 'tiger', if these are construed as picking out properties.

By contrast, numerical expressions such as 'zero' are epistemically rigid: 'zero' picks out 0 in every scenario. The same plausibly goes for various property terms, predicates, and relations: perhaps 'consciousness', 'wise', 'part', and 'cause', for example. And the same goes for various general terms: perhaps 'philosopher', 'friend', and 'action', for example.

Of course there are scenarios within which a word pronounced 'zero' picks out other things, but those scenarios are irrelevant to epistemic rigidity. What matters here is the intension of the actual word 'zero' and the way that this intension is evaluated at other scenarios. This intension is defined in terms of a priori entailments involving the actual word, not in terms of of the way that the word or others that sound like it behave when uttered in other scenarios.

The definition of epistemic rigidity given above is intuitively useful, but as a formal definition it has a couple of problems. First, it presupposes the notion of what an expression picks out in a scenario. On some approaches to epistemic space (as in E9), the notion of epistemic rigidity is used to help characterize evaluation in scenarios, with an ensuing danger of circularity. Second, it invokes the notion of trans-scenario identity: the relation whereby an object in one scenario is the same object as an object in another. But it is not entirely clear how to make sense of trans-scenario identity.[1] One could invoke an intuitive conception of these things, but there is another approach.

As an alternative, one can define an epistemically rigid expression as one whose extension can be known a priori. For example, there is an intuitive sense in which one can know the referent of 'zero' a priori, and in which one cannot know the referent of 'Hesperus' a priori (although there is some delicacy in understanding this sense, as we will see). It is natural to expect that if one can know an expression's extension a priori, it will pick out the same extension in all epistemically possible scenarios and vice versa. Likewise, if one cannot know an expression's extension a priori, it will pick out different extensions in different scenarios, and vice versa.

What it is to know an expression's extension is not entirely clear, as we saw in the second excursus. For current purposes, we should understand it in much the same way that we understood knowing that a sentence is true (in 2.2), so that metalinguistic knowledge is not required. To know what 'zero' refers to is just to know what zero is, where zero is presented under the guise of 'zero'. Intuitively, one can know a priori what zero is (where zero is presented under the guise of

---

[1] See 'The Nature of Epistemic Space' for a discussion of trans-scenario identity.

'zero'). By contrast, one cannot know a priori what Hesperus is (where Hesperus is presented under the guise of 'Hesperus').

In addition, to say that one can know the extension of 'zero' priori is not simply to say that there is a truth 'Zero is such-and-such' that one can know a priori. The most obviously relevant truth around here is just 'Zero is zero', but the existence of an a priori truth of that form does not suffice for epistemic rigidity. It is true that one can know 'Zero is zero' a priori while one cannot know 'Hesperus is Hesperus' a priori, because one cannot know 'Hesperus exists' a priori. But one can also know 'The number of stars is the number of stars' a priori, and 'The number of stars' is not epistemically rigid.

A more promising suggestion is that 'zero' is epistemically rigid iff one can know a priori (*de re*) of zero that it is zero (where zero in its predicative role is presented under the guise of 'zero'). Generalizing this pattern plausibly excludes both 'Hesperus' and 'The number of stars'. Tricky issues still arise, though. Someone might suggest that there is a name 'Starnum' whose reference is fixed to be the number of stars, and that by knowing 'Starnum is Starnum' we thereby know of Starnum that it is Starnum. To exclude this case, one requires a strong reading of *de re* knowledge in which one does not know *de re* of Starnum that it is Starnum simply in virtue of knowing that Starnum is Starnum.[2] I think that there is a natural way of reading *de re* attributions so that the definition gets the right results. But for present purposes I will leave the idea of knowing an extension a priori as intuitive.

All this can be extended naturally to the key case of properties and relations. A predicate is epistemically rigid when one can know a priori what it is for something (or some things) to satisfy the predicate. Intuitively, we know a priori what it is for one thing to be part of another thing. Arguably, we know a priori what it is for something to be conscious. But we do not know a priori what it is for something to be human, or what it is for one object to be more acidic than another. Correspondingly, we might say that we can know a priori (*de re*) of the parthood relation that it is the parthood relation, but we cannot know a priori of the more-acidic-than relation that it is the more-acidic-than relation. The idea of knowing a priori what it is for something to have a certain property is perhaps the most intuitive understanding, though.

There are numerous epistemically rigid expressions for abstract objects such as numbers ('zero'), properties ('conscious'), and relations ('part'). By contrast, it is

---

[2] For example, Scott Soames (2004) suggests an exportation principle that allows names (but not descriptions) to be exported from *de dicto* knowledge attributions to yield *de re* attributions. Then if one can know a priori that Starnum is Starnum, one can know a priori of Starnum that it is Starnum. I argue against exportation principles of this sort toward the end of 'Propositions and Attitude Ascriptions: A Fregean Account'.

arguable that there are no epistemically rigid expressions for concrete objects. For any expression E for a concrete object e, it is hard to see how we could know the extension of E a priori. On the face of it, for all we know a priori, E refers to e or to some quite distinct object f. Correspondingly, any ordinary expression for a concrete object picks out what seem to be different entities in different scenarios.

Another way to bring this out: any true identity statement in which both sides are epistemically rigid, such as '2 + 2 = 4', is a priori. This is a consequence of both the intuitive definition above and the definition in terms of scenarios: such an identity statement will be true at all scenarios and will therefore be a priori. By contrast, most true identity statements involving ordinary proper names for concrete objects, such as 'Mark Twain is Samuel Clemens', are not a priori. It follows from this that at least one of the names is epistemically nonrigid. Furthermore, the two names seem on a par, so that if one is epistemically nonrigid, both are epistemically nonrigid. In the case of abstract objects, there will also be a posteriori identity statements such as '0 is the number of phlogiston atoms'. But here the two expressions are plausibly not on a par. Numerical representations have a special status as designators for numbers, so that the left side is epistemically rigid while the right side is not.[3]

When an expression is epistemically rigid and also metaphysically rigid *de jure* (roughly, one can know a priori that it is metaphysically rigid), we can say that it is *super-rigid*.[4] A super-rigid expression has the same extension in all scenarios and in all possible worlds. We can know its extension a priori, and we can even know its extension in all possible worlds a priori.

In practice, most epistemically rigid expressions in natural language are also super-rigid. There are some fairly artificial expressions that are epistemically rigid but not super-rigid. Consider 'Whether (P iff actually P)', where P is any contingent sentence, and 'whether' is an operator that serves to pick the truth-value of the embedded sentence. Then this expression picks out *true* in all scenarios, but it picks out *false* in some non-actual worlds. One can even devise expressions that are epistemically rigid and metaphysically rigid *de facto* without being super-rigid. Still, any epistemically rigid expression E can easily be turned into a super-

rigid expression E' by rigidifying it *de jure*. For example, one can simply take E' to be 'the actual E''. Then $E = E'$ will be a priori, although it will not be necessary unless E is metaphysically rigid. So where epistemic (although not modal) matters are concerned, one can move easily between epistemic rigidity and super-rigidity.

These distinctions can naturally be represented in the two-dimensional semantic framework (discussed in E10 and E11), according to which expressions have primary intensions (functions from scenarios to extensions), secondary intensions (functions from worlds to extensions), and two-dimensional intensions (functions from scenario–world pairs to extensions). An epistemically rigid expression is an expression with a constant primary intension. A metaphysically rigid expression is one with a constant secondary intension. A super-rigid expression is one with a constant two-dimensional intension.

A closely related notion is that of semantic neutrality. An expression is semantically neutral roughly when its extension in any given possible world is independent of which scenario is actual.[5] Every super-rigid expression is semantically neutral, but the reverse is not the case. For example, 'the only conscious being in the world' is semantically neutral but not super-rigid (it picks out different entities in different worlds, but in a way that can be known without knowing which world is actual). Still, any semantically neutral expression is equivalent to a compound of super-rigid expressions. For example, the semantically neutral description just mentioned can be decomposed as 'the F', where the predicate F super-rigidly expresses the property of being the only conscious being in the world, and where 'the' contributes logical expressions that can be regarded as super-rigid. So there is little difference for our purposes between the class of sentences containing only super-rigid expressions and the class containing only semantically neutral expressions.

Epistemic rigidity is highly reminiscent of non-Twin-Earthability. The epistemically rigid expressions I have discussed here are roughly the same expressions as the non-Twin-Earthable expressions discussed at the start of chapter 7. The epistemically nonrigid expressions correspond to the Twin-Earthable expressions. Still, the two notions are not quite the same. A non-Twin-Earthable expression is roughly one whose extension does not depend metaphysically on the environment (all possible duplicates use corresponding expressions with the same extension). An epistemically rigid expression is roughly one whose extension

---

[5] Nida-Rümelin (2007) calls semantic neutrality 'actuality-independence'. In other work on these topics I have given semantic neutrality a larger role. Here I put more weight on super-rigidity as I think the notion is both more fundamental and easier to grasp.

does not depend epistemologically on empirical evidence. The application of these two notions coincides in many cases, but they can come apart.[6]

Consider the expression, 'Fred', stipulated to pick out 1 if there are any thinkers and 0 if not. Then 'Fred' is non-Twin-Earthable: any token of 'Fred' picks out 1. The same applies at the level of thought: any user of a *Fred* concept picks out 1. Still, 'Fred' is not epistemically rigid: it picks out 1 in scenarios containing thinkers, and 0 in scenarios not containing thinkers. Likewise, subjects are not in a position to know its referent a priori. 'Fred = 1' is true but not a priori: to know it, subjects need either introspective evidence that they are thinking or non-introspective knowledge that others are thinking. So non-Twin-Earthability and epistemic rigidity come apart here.[7]

Something similar goes for concepts of other intrinsic properties (for a notion of Twin-Earthability tied to intrinsic duplicates) or functional and phenomenal properties (for a notion tied to functional and phenomenal duplicates). For example, if 'Bill' is stipulated to pick out the phenomenal color in the center of my visual field, then 'Bill' will not be Twin-Earthable (in every duplicate the corresponding token will pick out phenomenal blueness), but 'Bill = phenomenal blueness' is still not a priori.

In the other direction: it is plausible that in our world, any epistemically rigid expression is non-Twin-Earthable. But in some possible worlds, this might not be so. For example, if there are Edenic worlds (see 7.4 and 'Perception and the Fall from Eden') in which subjects are directly acquainted with instances of primitive redness in their environments, then their expression 'redness' or 'primitive redness' may be Twin-Earthable (for reasons discussed in 8.4), at least in a sense where Twin-Earthability is tied to intrinsic duplication. But our expression 'Edenic redness' is plausibly epistemically rigid, and the same goes for the corresponding expressions in the Edenic world: Edenic subjects are in a position to know just what property they are talking about, simply by possessing the concept of Edenic redness. So this is at least a potential case of epistemic rigidity without non-Twin-Earthability.

---

[6] Epistemically nonrigid but semantically neutral expressions such as 'The only conscious being' will be extensionally Twin-Earthable but not intensionally Twin-Earthable, in the sense defined in the additional excursus on Twin-Earthability. Roughly, epistemic rigidity stands to extensional non-Twin-Earthability as semantic neutrality stands to intensional non-Twin-Earthability. Another approximate parallel is that epistemic rigidity stands to Twin-Earthability roughly as scenarios stand to contexts of utterance (although see footnote 8).

[7] I discuss cases like this in 'Does Conceivability Entail Possibility?' and 'The Two-Dimensional Argument against Materialism' as counterexamples to George Bealer's thesis that there are no a posteriori necessities involving semantically stable expressions. Semantic stability is a sort of non-Twin-Earthability (closest to the intensional Twin-Earthability discussed in the additional excursus). This putative role for semantic stability is better played by semantic neutrality or super-rigidity.

I think that epistemic rigidity is clearly the more fundamental of the two concepts here. At least where epistemological and modal matters are concerned, epistemic rigidity and super-rigidity cut things closer to the joints. Non-Twin-Earthability is interesting for these purposes to the extent that it approximates epistemic rigidity, and is independently interesting for its connections to internalism and externalism about content, but epistemic rigidity runs deeper.

Epistemic rigidity should also be distinguished from context-independence.[8] It is arguable that ordinary proper names such as 'Gödel' are extensionally context-independent: they pick out the same referent in every context. They are not epistemically rigid, however: they do not pick out the same referent in every scenario. In reverse, a term such as 'small' (construed as a predicate of numbers) may be context-dependent while being epistemically rigid in every context. Some terms ('small' as a more general predicate, perhaps) may even be epistemically rigid in some contexts but not others

We can think of epistemically rigid expressions as *referentially transparent* expressions, and epistemically nonrigid expressions as *referentially opaque* expressions. As defined earlier, an epistemically rigid expression is one whose extension is available on ideal a priori reflection, while the extension of an epistemically nonrigid expression is not knowable a priori. Correspondingly, we can think of epistemically rigid expressions as expressing referentially transparent concepts—concepts whose extension is knowable a priori—while epistemically nonrigid expressions express referentially opaque concepts. Referentially transparent concepts come with an especially direct grip on the corresponding entities in the world.

What sorts of expressions are epistemically rigid? We can approach the question by first examining epistemically nonrigid expressions. Analyzed from within the scrutability framework, the obvious examples fall into two classes. The first class includes indexicals: primitive indexicals, such as 'I', 'now', and phenomenal demonstratives, along with other less primitive indexicals that derive from these, such as 'today', 'here', and ordinary demonstratives. These are epistemically nonrigid because they function indexically to pick out a certain ostended entity—the current time, the present subject, and so on—and the subject is not in a position to know a priori what the ostended entity is.

---

[8] Context-dependence should also be distinguished from Twin-Earthability. 'Gödel' is Twin-Earthable: a corresponding expression could be used by a twin with a different referent. It is also arguably context-independent: the English word 'Gödel' picks out the same referent in every context. The difference arises from the fact that the corresponding expression on Twin Earth need not be the English word. In reverse, 'small' might be context-dependent without being Twin-Earthable, as long as any pairs of contexts in which it is uttered with different referents are not contexts involving twins.

The second class includes *role-scrutable* expressions: roughly, expressions whose extension is a priori scrutable from more basic truths by determining what plays a certain role (typically although not necessarily a causal role). For example, the extension of 'water' is a priori scrutable by determining what plays the role of (roughly) being the clear drinkable liquid that we have seen in our environment. The extension of 'Gödel' is a priori scrutable by determining what plays the role of being called 'Gödel' by others and of being at the other end of a causal chain.

Within a definitional framework, epistemically nonrigid expressions will include primitive indexicals (such as 'I' and 'now') and definite descriptions 'the D' (for example, 'the watery stuff around here'), where D is made up of primitives and it is not a priori what is the D. They will also include certain descriptive predicates (for example, 'has my favorite property') and general terms. In some cases D may include only epistemically rigid expressions: for example, 'The most friendly being in the universe'. These cases will turn on the fact that even if we know a priori what property a predicate F refers to, we often are not in a position to know which entities satisfy F. In other cases D will also include primitive indexicals: 'today' corresponds roughly to 'the day including now'.

In all of these cases, the definitional framework provides a clear explanation of epistemic nonrigidity: we would expect primitive indexicals to be epistemically nonrigid, and we would expect the relevant descriptions to be epistemically nonrigid also. We might say that in the definitional framework, these expressions are *conceptually opaque*: their referential opacity is apparent through conceptual reflection, so that their conceptual structure guarantees that they are referentially opaque.

Within the scrutability framework, something similar applies. We will still have primitive indexicals in the first class. The expressions in the second class will no longer be precisely equivalent to descriptions, and there may not be a simple specification of the relevant role. But if we follow the approximate definition model of chapters 1 and 8, a role-scrutable expression will be at least approximately a priori equivalent (in a given context) to certain descriptions of the form 'the thing that plays such-and-such complex role'. And even if we eschew approximate definitions, the relevant expressions will still be scrutable from truths about the various roles that the extension plays, just as truths such as 'water is H_2O' are scrutable from truths about the various properties of H_2O.

As with the definitional framework, the scrutability framework provides a natural explanation of epistemic nonrigidity. As before we would expect primitive indexicals to be epistemically nonrigid, and we would expect the relevant role-scrutable expressions to be epistemically nonrigid also. These expressions can also reasonably count as conceptually opaque: their referential opacity is apparent through conceptual reflection, so that their conceptual structure and role guarantees that they are referentially opaque.

If these are the epistemically nonrigid expressions, which expressions are epistemically rigid? If we buy into the version of the scrutability framework that involves conceptual priority, the most obvious candidates are the non-indexical primitive expressions in a scrutability base. For example, 'and', 'zero', 'law', 'fundamental', and 'consciousness' are all plausible candidates to be in a scrutability base. Other candidates are expressions that derive from these non-indexical primitives, either through definition or scrutability, as long as we avoid role-scrutability. For example, other logical expressions ('some'), mathematical expressions ('plus'), mental expressions ('believe'), and nomic expressions ('cause') are also plausible candidates to be epistemically rigid, as are various expressions that derive from a combination of these (e.g., 'friend' or 'philosopher', at least on certain readings).[9]

Unlike the epistemically nonrigid expressions considered above, none of these expressions appear to be conceptually opaque (at least granted views on which they are conceptually primitive). They are *conceptually transparent*, in that conceptually they appear to be transparent: no referential opacity is revealed by conceptual reflection. On the face of it, 'zero' transparently picks out zero, and 'consciousness' transparently picks out consciousness. In effect, these expressions at least *seem* to be epistemically rigid. That is, they seem to give a direct grip on their referent in the world (phenomenal properties, fundamentality, lawhood, addition, and so on), whether or not they really do.

On the face of it, it is most plausible to hold that these expressions are epistemically rigid. That is, it is plausible to endorse a Conceptual/Referential Transparency thesis: all conceptually transparent expressions are referentially transparent. This thesis has the consequence (given the above) that all non-indexical primitive expressions are epistemically rigid. I will stipulate that conceptually transparent expressions must also be metaphysically rigid *de jure*. Then the thesis allows us to conclude that conceptually transparent expressions are super-rigid.

There will be philosophical views that deny this thesis. For example, some type-B materialists may hold both that 'consciousness' is conceptually primitive and that it is epistemically nonrigid: it refers to a certain physical property, even though one could not know that a priori.[10] Some 'type-B color physicalists' who

---

[9] In some of these cases one may need to disambiguate, precisify, or fix a context first.

[10] Not all type-B materialists will deny the Conceptual/Referential Transparency thesis. One sort of type-B materialist holds that phenomenal properties are necessitated by physical properties but are not identical to them; this version can accept the thesis and will probably instead deny the Apriority/Necessity thesis below. A second sort holds that phenomenal properties are identical to physical properties while holding that all physical expressions for those properties are epistemically nonrigid. That view can accept both of these theses, although doing so probably leads to a version of Russellian monism. A third sort holds that 'consciousness' is a primitive indexical or derives from primitive indexicals such as demonstratives. A fourth sort holds that 'consciousness' is role-scrutable from more basic primitives, although here the issue will recur for the more basic primitives.

are primitivists about color concepts but not about color properties may hold a similar view of color expressions.[11] An analogous view about spatial expressions, nomic expressions, and others is possible at least in principle. These views can still agree that the relevant expressions are conceptually transparent in the sense above, while holding that conceptual transparency does not entail referential transparency.

We might usefully divide the corresponding concepts into three groups. Expressions that are both conceptually and referentially transparent express *transparent* concepts: concepts that reveal their referents. Expressions that are both conceptually and referentially opaque express *opaque* concepts: concepts that obscure their referents, at least in the sense that they do not reveal their referent. Expressions that are both conceptually transparent but referentially opaque express *pseudo-transparent* concepts: concepts that appear to reveal their referents but in fact obscure them. The type-B theorists in the previous paragraph are naturally allied with the view that phenomenal concepts or color concepts are pseudo-transparent.[12]

These views are varieties of *primitive externalism*: externalism about reference for (non-indexical) primitive concepts. These will often be externalist in the sense that the referent of a primitive concept is determined by factors outside the skin, but they need not be: a type-B theorist might hold that 'consciousness' refers to an internal neurophysiological property, for example. But they will be externalist in at least the sense that their referent lies outside our immediate cognitive grasp: even when full a priori mastery of the relevant primitive concept does not yield knowledge of what it picks out. Primitive externalism is naturally allied with an externalist account of what grounds reference for primitive expressions: perhaps a causal, teleological, or reference-magnet account.

A full assessment of primitive externalism is a substantial project in its own right. For now, I note that the major arguments for externalism are not arguments for primitive externalism. Putnam-style arguments apply best to role-scrutable expressions, and Burge-style arguments apply to expressions used deferentially. A quite new sort of argument would be needed to establish primitive externalism. So following the methodology laid out in the introduction, I take the default view to be primitive internalism.

Someone might argue for primitive externalism by appealing to the causal theory of reference or some other externalist theory. As in the previous paragraph, though, I think the arguments for the causal theory are grounded in the

---

[11] Here I have in mind Byrne and Hilbert 2007, who seem to treat color concepts as primitive and hold that at least some color truths are inscrutable from underlying physical truths while holding that colors are physical properties. By contrast Jackson (1998) holds what we might think of as type-A color physicalism, involving functionalism about color concepts.

[12] Philip Goff (2011) uses the terminology of transparent and opaque concepts for a similar distinction.

cases of role-scrutable expressions and expressions used deferentially and do not have much purchase on the case of primitive expressions. There is little reason to accept a causal theory of reference for expressions such as 'zero' or 'part', so there is little reason to think these theories are universal. In fact, one can argue (as I do on the additional excursus on reference magnets and the grounds of intentionality) that the role of causation and other apparently externalist factors in reference is grounded in certain features that are internal to a subject's grasp, suggesting that the purely external role needed for primitive externalism would require a distinct mechanism that there is not much positive reason to believe in.

Most fundamentally, I think that primitive externalism is to be rejected because it gives us too little grip on what we are thinking and saying. We have a substantial grasp of what we are talking about when we talk about laws of nature or parthood or consciousness, and primitive externalism is not in a position to explain that substantial grasp. That issue requires a sustained investigation in its own right, though. In the meantime, I flag the issue and I register my own view, which rejects primitive externalism and accepts the Conceptual/Referential Transparency thesis.

Epistemic rigidity can help us to analyze Kripke's examples of the necessary a posteriori. We have already seen that true identity statements involving epistemically rigid expressions are a priori. Correspondingly, any a posteriori identity statement must involve at least one epistemically nonrigid expression. This is just what we find in Kripke's examples of necessary *a posteriori* identity sentences: 'Hesperus is Phosphorus', 'heat is the motion of molecules', 'water is H_2O', and so on. In each of these cases, a key term is metaphysically rigid but not epistemically rigid.

We can put the point by saying that identity sentences involving super-rigid expressions on each side are a priori iff they are necessary. This thesis follows immediately from the definition of super-rigidity. Similarly, any necessary a posteriori identity sentence must involve at least one expression that is not super-rigid. Assuming that both expressions are metaphysically rigid, as in the paradigm cases, then at least one expression must be epistemically nonrigid. In a sense, the combination of epistemic nonrigidity and metaphysical rigidity can be seen as the source of the necessary a posteriori.

This trivial thesis can be strengthened in a couple of ways. First, if we accept the Conceptual/Referential Transparency thesis, then we can derive the thesis that all identity statements involving conceptually transparent expressions on each side are a priori iff they are necessary. In effect, on this view conceptual opacity is the source of epistemic nonrigidity, and the combination of conceptual opacity and metaphysical rigidity is the source of necessary a posteriori identity statements. This model certainly fits all the a posteriori identity statements that Kripke discusses.

Second, one could strengthen the thesis in a different direction by generalizing from a posteriori identities to all a posteriori necessities, as follows:

Apriority/Necessity Thesis: If a sentence S contains only super-rigid expressions, S is a priori iff S is necessary.

This thesis is trivially true when S is an identity statement, and nontrivial but plausible when S is not. Certainly, all of Kripke's examples of the necessary a posteriori involve epistemically nonrigid expressions. Still, some philosophical positions will deny the thesis. For example, 'An omniscient being exists' plausibly involves only super-rigid expressions, and some theist views entail that this sentence is necessary but not a priori. Likewise, some views of mathematics (as discussed in chapter 6) may allow that there are mathematical truths that are necessary but not a priori. Once again, however, the Apriority/Necessity thesis fits the a posteriori necessities that Kripke discusses, all of which involve epistemically nonrigid expressions.

Finally, one can make both strengthenings at once, holding that if a sentence S contains only conceptually transparent expressions, S is a priori iff S is necessary. We might call this the Strong Apriority/Necessity thesis: it follows from the original Apriority/Necessity thesis and the Conceptual/Referential Transparency thesis. The Strong Apriority/Necessity thesis in effect says that all necessary a posteriori sentences and all contingent a priori sentences involve conceptually opaque expressions.

A counterexample to the Strong Apriority/Necessity thesis would be a necessary a posteriori or contingent a priori sentence involving only conceptually transparent expressions. Such a sentence would be what I have elsewhere called a *strong* a posteriori necessity (as opposed to Kripke's weak a posteriori necessities involving conceptually opaque expressions), or a strong priori contingency. I have argued at length (in, e.g., 'The Two-Dimensional Argument against Materialism') that there are no strong a posteriori necessities, and those arguments apply equally to strong a priori contingencies. So I accept the Strong Apriority/Necessity thesis.

A key thesis about super-rigidity that I discuss elsewhere in this book (tenth and sixteenth excursuses and chapter 8) is the following.

Super-Rigid Scrutability: All truths are scrutable from super-rigid truths and indexical truths.

There is also a generalized version: all epistemically possible truths are scrutable from epistemically possible super-rigid sentences and indexical sentences. Here super-rigid sentences are those containing only super-rigid expressions, and indexical sentences are those of the form 'E is D' where E is a primitive indexical and D contains only super-rigid expressions.

Generalized Super-Rigid Scrutability is a consequence of the thesis that all epistemically nonrigid sentences are scrutable from epistemically rigid sentences and indexical sentences. One can also derive a version of Super-Rigid Scrutability from Conceptual/Referential Transparency, along with the theses that all truths are scrutable from truths involving conceptual primitives and that all non-indexical conceptual primitives are conceptually transparent. In practice, the most important sort of challenge to Super-Rigid Scrutability (from those otherwise sympathetic with the scrutability framework) is likely to come from primitive externalism.[13]

As before, I think that primitive externalism is false: all epistemic nonrigidity derives from either primitive indexicality or role-scrutability. So I hold that all epistemically nonrigid sentences are scrutable from epistemically rigid sentences and indexical sentences, and therefore accept Generalized Super-Rigid Scrutability.[14] I return to the issue in chapter 8.

[13] A tempting argument for Generalized Super-Rigid Scrutability runs as follows. Even if we cannot know the extensions of our expressions a priori, we can know their primary intensions a priori. So we can refer super-rigidly to primary intensions. But then all truths will be scrutable from truths of the form 'p is true', where p specifies the primary intension of a truth in the scrutability base. Similarly, we can refer super-rigidly to scenarios, so all truths will be scrutable from 's is actualized', where s specifies a scenario super-rigidly. However, an opponent can note that if Super-Rigid Scrutability is false, primary intensions and scenarios are best understood as linguistic or Fregean entities (E10). They can then hold that 'true' and 'actualized' as predicates of these entities are not super-rigid. Because the properties picked out by the basic linguistic or Fregean entities can be known only empirically, the reference relation involving these entities can be known only empirically, and likewise for truth and actualization.

[14] One can develop a weak sense in which even pseudo-transparent concepts count as epistemically rigid. There is an intuitive sense in which 'consciousness' at least picks out the *feature* of consciousness in every scenario, whether or not it picks out the property of consciousness. Here features are roughly projections of conceptually transparent concepts: they correspond to the way that properties would be if those conceptually transparent concepts were referentially transparent, as they seem to be. They are abstract objects that are akin to properties in that they can be predicated of objects, but they are individuated by the transparent or pseudo-transparent concepts that pick them out. Then we can say that phenomenal features are distinct from neural features, even if phenomenal properties are identical to neural properties. On my view, features correspond one-to-one with properties, so these two will stand and fall together. But for someone who believe in pseudo-transparency, distinctness of features does not lead to dualism about the mind-independent world, as features are in the relevant sense mind-dependent. We can then say that phenomenal concepts are weakly epistemically rigid in that they pick out the feature of consciousness in every scenario, and color concepts are weakly epistemically rigid in a similar way. The same goes for weak super-rigidity. This would then allow even those who believe in pseudo-transparent concepts to accept Generalized (Weak) Super-Rigid Scrutability. This may be useful for allowing them to accept some applications of Generalized Super-Rigid Scrutability: for example, we could then use features instead of properties to construct scenarios in the fashion of the tenth excursus.