

## HALE ON THE ARCHITECTURE OF MODAL KNOWLEDGE

BOB FISCHER

*Texas State University*

There are many modal epistemologies available to us. Which should we endorse? According to Bob Hale, we can start to answer this question by examining the architecture of modal knowledge. That is, we can try to decide between the following claims: (1) knowing that  $p$  is possible is essentially a matter of having a well-founded belief that there are no conflicting necessities—a necessity-based approach—and (2) knowing that  $p$  is necessary is essentially a matter of having a well-founded belief that there are no conflicting possibilities—a possibility-based approach. Hale argues for the first of these claims, but I contend that his arguments fail. However, it doesn't follow that we should endorse a possibility-based approach. I repurpose Hale's arguments to show those who would endorse possibility-based approaches need to say more about our modal knowledge concerning logic and mathematics; if they can't, then they ought to endorse a hybrid modal epistemology that doesn't give priority to one modality across the board. Additionally, those who endorse possibility-based views may be committed to Peter van Inwagen's modal skepticism.

### 1. Introduction

There are many modal epistemologies available to us. Which should we endorse? According to Bob Hale (2002, 2013), we can start to answer this question by examining the architecture of modal knowledge. He reasons as follows. Since the modalities are interdefinable (it's possible that  $p$  iff it isn't necessary that  $\sim p$ ; it's necessary that  $p$  iff it isn't possible that  $\sim p$ ), we can think about each modality as the 'absence' of the other: possibility is the 'absence' of conflicting necessities, while necessity is the 'absence' of conflicting possibilities. Then, we can try to decide between the following claims:

- (1) Knowing that  $p$  is possible is essentially a matter of having a well-founded belief that there are no conflicting necessities.<sup>1</sup>
- (2) Knowing that  $p$  is necessary is essentially a matter of having a well-founded belief that there are no conflicting possibilities.
- (3) Neither (1) nor (2).

The first claim represents a necessity-based approach to modal knowledge; the second, a possibility-based approach. Hale calls these

<sup>1</sup> I'm adapting this language from (2002, 6).

“asymmetrical,” since they afford a privileged status to knowledge of one modality. To reject them is to endorse symmetry—i.e., the last claim.

Hale’s goal is to motivate a necessity-based approach.<sup>2</sup> If his arguments work, then we shouldn’t accept views like Stephen Yablo’s (1993). In short, Yablo argues that if it seems to me that I can objectively imagine a world that I take to verify  $p$  (i.e., imagine such a world *de re*), then I am justified in believing that there is such a world, and hence that  $p$  is possible. This is a possibility-based approach: I acquire a well-founded belief that there are no conflicting necessities *by objectively imagining the relevant world*; I don’t need to *begin* with a well-founded belief that there aren’t any conflicting necessities.<sup>3</sup> So if Hale is correct, Yablo isn’t. The same point probably applies to most conceivability- or imaginability-based views—including, e.g., Hart (1988), van Woudenberg (2006), Hanrahan (2007), and Kung (2010).

Granted, the architecture of modal knowledge doesn’t settle which necessity-based view to adopt. Hale, of course, plumps for his essence-based approach, but there are alternative essence-based approaches available, such as Lowe (2012) and Tahko (forthcoming).<sup>4</sup> Moreover, the two most influential views in recent years—namely, the ones due to Williamson (2007) and Chalmers (2010)—are also on the table. Williamson relies on knowledge of logic and certain constitutive principles, which appear to express the essential properties of objects (for more on this, see Roca-Royes 2011). Chalmers, of course, relies on knowledge of semantic necessities. Still, it’s hard to know where to begin in the epistemology of modality, so it’s worth exploring any way to whittle down our options.

Unfortunately, I don’t think that Hale’s arguments manage to do that. In the next two sections, I explain why Hale hasn’t made his case. However, we can repurpose Hale’s arguments to clarify the prospects for possibility-based approaches, and I take up that project in the final section.

## 2. Clarifying the Options

Hale argues for the inadequacy of possibility-based approaches by contending that possibility-based approaches can’t satisfy two necessary

<sup>2</sup> Surprisingly, Hale’s argument has received very little attention, despite being the Presidential Address at the 2002 meeting of the Aristotelian Society. There are a few passing references to it in the literature—see, e.g., Vaidya (2010) and Roca-Royes (2011)—but no critical discussion.

<sup>3</sup> Two points. First, Hale is discussing knowledge; Yablo, justified belief. But the difference doesn’t matter here. Second, it may be possible to construe Yablo’s view as an instance of a necessity-based approach, in which case it wouldn’t be vulnerable to Hale’s arguments. But I take it that this would run counter to the spirit of that view.

<sup>4</sup> It’s unclear whether the ground-clearing maneuver works even if the arguments succeed, since the link between essential and modal knowledge is difficult to make out. On this, see Horvath (2014).

conditions on an asymmetrical modal epistemology. (In a possibility-based approach, possibility is the “dominant” modality and necessity is the “recessive.” Things are reversed in a necessity-based approach.) And if possibility-based approaches can’t satisfy the second of those conditions, then that’s a problem for *any* view that allows some basic knowledge of possibility. Hence, Hale’s argument creates trouble for symmetrical approaches as well, leaving necessity-based views as the last ones standing.

The two necessary conditions are as follows:

- (1) A belief about a recessive modal claim is to be justified by appeal to the fact that, after a well-directed and thorough search for relevant dominant modal truths, we found none that rules out the recessive modality claim in question.<sup>5</sup>
- (2) “[There is] a base class of dominant modal truths which meet two conditions:
  - a they can be known without reliance upon any recessive modality claims and
  - b they are collectively strong enough to support the superstructure of modal knowledge to be erected over them” (2002, 8).

Condition 1 is driven by two concerns: first, that an approach be properly asymmetrical, where one modality has epistemic priority over the other; second, that the approach not be guilty of making arguments from ignorance. Focusing on the second point, if we want to infer that *p* is possible, it isn’t enough to say that we aren’t aware of any necessity that rules it out, as that would be like inferring that the butler did it simply because there is no evidence that exonerates him (not because of any positive reason to think that he’s guilty). In addition, we need to be able to say that we have tolerably good reason to think that there is no such necessity. Hence, the crucial part of Condition 1 is the expression, “after a well-directed and thorough search.” It’s this qualifier that justifies the shift from absence of evidence to evidence of absence.

Together, Conditions 1 and 2a state what it is for an approach to be properly asymmetrical: it must both posit an epistemic priority

<sup>5</sup> This formulation is based on two statements that Hale makes. First: “According to [an asymmetrical] strategy, beliefs about the recessive modality are to be justified by appeal to the fact that we have found no dominant modal truths which rule out the recessive modal claim up for assessment” (2002, 6). Second: “the asymmetric theorist [...] must make out that in following the route she proposes, we need not be simply passing, gratuitously, from mere lack of knowledge of any relevant countervailing necessities, but may have looked responsibly for them and failed to find any. This requires that we can give decent operational sense to the idea of a well-directed and thorough search for necessities relevant to the assessment of a given possibility claim (or, of course, for possibilities relevant to a given necessity claim, in case of a possibility-based asymmetric approach)” (2002, 7).

(Condition 1) and avoid there being an in-principle objection to that priority (Condition 2a). Condition 2b amounts to a statement of conservatism: whatever we take ourselves to know about modal matters, the base class needs to be able to explain it. However, Condition 2b does *not* require that the base class be able to explain the superstructure all by itself. I can learn that  $p$  is possible, infer that  $q$  is not necessary, and infer that  $r$  is therefore possible. The base class just needs to be able to explain the superstructure via a series of such inferences. Moreover, Condition 2b is compatible with denying that we have much modal knowledge: someone might maintain that if there *were*, say, modal knowledge in metaphysics, then it would be acquired by asymmetrical methods; but since there isn't, it isn't. On this view, the base class of dominant modal truths is indeed strong enough to support the superstructure of modal knowledge to be erected over it, since you need nothing to support nothing.

### 3. Hale's Two Arguments

With Hale's two conditions in view, we can now consider his criticisms of the possibility-based approach.

#### 3.1. Hale's First Argument

Hale's first argument aims to show that it's unclear whether a possibility-based approach can satisfy Condition 2a—namely, that there is a base class of possibilities that can be known without reliance upon any necessities. Hale assumes that any possibility-based approach will involve the view that conceivability is a guide to possibility. He doesn't argue for the connection between conceivability and possibility-based views, and I doubt that it's genuine (for views that appear to be possibility-based and yet don't employ conceivability, see, e.g., Roca-Royes 2007; Leon 2009). Nevertheless, I grant the assumption here.

On Hale's view, conceivability "strictly implies" possibility, which is to say that you haven't conceived that  $p$  unless  $p$  is metaphysically possible (2002, 11, n. 19). Given this assumption, the real question concerns when we're justified in believing that we've conceived that  $p$ . If we can be justified in believing that we've conceived that  $p$  without knowing any necessities, then possibility-based approaches have some chance of satisfying Condition 2a. If not, then not.

Why should we doubt that we can be justified in believing that we've conceived that  $p$  without knowing any necessities? Hale writes:

I am not sure that one can be justified in taking it to be conceivable that  $p$  without some assurance that potentially relevant broadly logical necessities—those relating to concepts involved in  $p$ —do not

ensure that  $\sim p$ . And to have any such assurance, it seems, one needs to know what the potentially relevant necessities are—those necessities which are directly reflected in the requirements for the application of the concepts involved. (2002, 14–15)

Hale's thought seems to be as follows. Suppose that you've never seen a hairless cat, and you're considering whether such a cat is possible. You think you can conceive of one, but are you justified in believing that you have? Hale's suggestion is that you can't be unless you know what has to be true of something for it to count as a cat, as well as what has to be true of something for it to count as hairless. Absent this knowledge, you can't tell whether being hairless is precluded by cathood—i.e., you can't tell whether you're running afoul of a conceptual necessity.

The crucial premise here is that we can't be justified in thinking that it's conceivable that  $p$  unless we're justified in thinking that there are no conceptual necessities that entail  $\sim p$ . So, we could make short work of the argument by denying that there are conceptual necessities to know.<sup>6</sup> But this move aside, how does the argument fare?

Two objections remain. The first concerns what's involved in being competent with a concept. The second concerns the epistemic standard to which Hale is holding the possibility-based approach.

The first objection takes issue with Hale's view that you're a competent user of a concept only if you know the concept's application conditions—i.e., the *necessary* conditions for a concept's application. But that seems too strong: it's enough to know conditions that are *necessary-in-a-context*.<sup>7</sup>

Very roughly, the argument for this is based on the point of having concepts: namely, to divvy up the things in the world. It would be startling if this need led us to generate concepts precise enough to allow us to decide what they require in all circumstances, since our circumstances rarely require such precision. (Linguistic evolution, like evolution generally, is satisficing, not optimizing.) Most people's circumstances, for example, have never required the conceptual precision needed to answer questions like these: How much hair can you have on your head while still being bald? Is an El Camino a car? If you breed foxes until they become very docile and look like dogs (a feat that you can accomplish in about forty generations), are they dogs? If a transgender man (i.e., someone who was identified female at birth, but now identifies as male) is unmarried, is he a bachelor? What about Giaquinto's (2008) example of a male who marries to help someone get citizenship, but otherwise lives a

<sup>6</sup> For three very different routes to this conclusion, see Quine (1951), Williamson (2006), and Giaquinto (2008).

<sup>7</sup> Alternately, we may need to know the necessary conditions for the application of a concept *in close worlds* (i.e., in close worlds, if  $x$  is a bachelor, then  $x$  is an unmarried male).

single lifestyle? Is water  $H_2O$ , or just one *species* of  $H_2O$ —namely, the one *without* extra deuterium (which makes for heavy water; or without protium, for semiheavy water; or without tritium, for tritiated water)?<sup>8</sup>

Granted, we could refine the relevant concepts so that there would be answers to the above questions. But precisification is *not* uncovering existing detail; rather, it is *adding* detail where none existed previously. Moreover, it isn't clear how and when precisification happens. Just because some scientists begin to talk a certain way about water, and we follow suit, that doesn't show that the concept has changed. It's equally likely that we're just deferring to scientists concerning *truths* about water. On such a view, it needn't be conceptually necessary that water is  $H_2O$ ; rather, water's being  $H_2O$  might just happen to be among the truths that we all take for granted about water.

The upshot of the above is that Hale bears the burden of proof. If someone wants to maintain that our concepts are more precise than would be required for both their original purpose and their normal use, and that competent use involves knowing these finer details, then she owes us some arguments. However, Hale doesn't provide them here or elsewhere.

So let's suppose that competence involves knowledge of conditions that are necessary-in-a-context for the application of a concept, not knowledge of necessary conditions *simpliciter*. If so, then being competent with a concept will require knowing some *local* necessary conditions—i.e., knowing, in a context, some of what the concept requires and precludes. However, it won't require knowing anything interesting about what those concepts require in all circumstances—i.e., the (unrestrictedly) necessary conditions for their application. Hence, when our attempts at conceiving concern perfectly ordinary objects in ordinary circumstances—e.g., my mug being several inches to the right of where it is on my desk—we may be justified in believing that we've conceived, since there's no reason to doubt that we can tell whether our concepts apply. Granted, we may not be justified in weirder and wilder circumstances. But that isn't a problem for conceiving generally; rather, it's a problem for *those specific conceptions*. So, absent further arguments, it looks as though defenders of the possibility-based approach are entitled to the view that there are *some* circumstances in which we can—*contra* Hale—conceive that  $p$  without knowledge of the relevant broadly logical necessities.

So much for the first objection. What about the second? Recall the crucial premise in Hale's argument: we can't be justified in thinking that it's conceivable that  $p$  unless we're justified in thinking that there are no conceptual necessities that entail  $\sim p$ . Defenders of the possibility-based approach should concede that there are cases in which I'm mistaken about whether I've conceived that  $p$ . But they're fallibilists,

<sup>8</sup> For more on this point, see Weisberg (2006).

so they'll take such cases in stride. To make an *objection* out of these errors requires saying something about their *frequency* and *distribution*. In particular, the errors need to be *fairly common* and *distributed in an unpredictable way*. For suppose they aren't fairly common: then, conceiving is reliable. And suppose, for example, that they cluster in some particular domain of inquiry: then, we can avoid them by refusing to trust our conceivings in that domain. However, Hale offers no reason to believe either that the errors are frequent or that they're distributed unpredictably. We have only the allegation that we *might* be making too many errors, which isn't much of a reason to distrust conceivings generally.<sup>9</sup>

What's more, defenders of conceivability can explain away the intuitive appeal of the thesis that conceivings often lead us astray. That intuition probably exists because of an availability bias: as philosophers, we tend to focus on cases that have captured the attention of philosophers: e.g., disembodied minds, zombies, swamp people, maximally perfect beings, etc. When we reflect on the countless philosophically *boring* conceivings, it's implausible that the philosophically interesting ones are representative: e.g., the shelves in my garage being about a yard closer to the door (making more room to hang tools on the wall), the chain on my bike not being covered in rust, etc.

So there appear to be two problems with Hale's first argument: first, it relies on an overly-demanding view of conceptual competence; second, Hale hasn't given us a reason to think that conceiving's errors are so frequent or so distributed that we shouldn't trust it at all. It seems safe to conclude, therefore, that Hale's first argument fails.

### 3.2. Hale's Second Argument

Hale's second criticism aims to show that it's unclear whether a possibility-based approach can satisfy Condition 1—i.e., the belief that *p* is necessary is to be justified by appeal to the fact that, after a well-directed and thorough search for relevant possibilities, we found none that

<sup>9</sup> We can make the same point by thinking about its perceptually seeming to you that *p* (where if veridical, we'd say that you perceive that *p*). Plainly, its seeming to you that *p* doesn't entail that *p*. This, however, is not a general undercutting defeater for seemings—namely, a reason to think that seemings are generally unreliable. To get a general undercutting defeater for seemings, we would need an argument for the claim that seemings lead us astray often enough that we would be unjustified in believing that things are as they appear to be. So suppose we had a compelling argument for the view that we're more likely to be in a skeptical scenario than one in which our seemings are veridical. Then, perhaps, we would have a general undercutting defeater for seemings, and it would be appropriate to reject beliefs as unjustified if they were formed on the basis of seemings. Absent that, we might have the more-familiar *local* undercutting defeaters: if I've had enough alcohol to impair my vision, and I know this, then I have a local undercutting defeater for the way that things (visually) seem to me. But acquiring such a defeater doesn't give me a *general* reason to doubt my seemings; rather, it is context-specific, applying only until I sober up.

rules out  $p$ 's necessity.<sup>10</sup> Essentially, his contention is that it isn't clear you can engage in such a search for relevant possibilities, since it isn't clear that there's any way to assess what the relevant possibilities are. He puts the point as follows:

Whereas there is a reasonably well-circumscribed class of broadly logical necessities relevant to the assessment of a claim about conceivability—those necessities, appreciation of which goes along with mastery of the concepts involved—there is no similarly well-circumscribed class of possibilities relevant to the assessment of a claim about broadly logical necessity. Of course, there will be indefinitely many *prima facie* broadly logical possibilities, involving some or all of the concepts featured in the given necessity-claim, which competence with those concepts does not enable us to rule out. But these will not be germane to the assessment of that necessity-claim. And of course, to justify the claim that it is broadly logically necessary that  $p$ , we must justify the claim that it is *not* (perhaps despite some appearance to the contrary) broadly logically possible that  $p$ —for these are the same claim. My point is that there is no well-defined class of *genuine* (broadly logical) possibilities we need to review. (2002, 15)

We should pause to observe that you can make a quick *tu quoque* response if you criticize Hale's view about the relationship between broadly logical and absolute necessity—a view he defends in his (1996), (2012), and (2013):

*Broadly logical* necessities (which comprise, in my usage, just narrowly or strictly logical together with analytic or conceptual necessities) are *properly included* within the class of *absolute* necessities. Strictly, of course, there are two assumptions here: first, that broadly logical necessities are indeed absolute, and second, that they do not exhaust the class of absolute necessities. (2002, 8)

Suppose this weren't true, and absolute necessities were properly included in the set of broadly logical necessities.<sup>11</sup> Then, you could know that  $p$  is broadly logically necessary without being in a position to know whether  $p$  is absolutely necessary. And if that were the case, then there wouldn't be a well-defined class of broadly logical necessities on which to draw to assess possibility claims.

However, I don't want to pursue that objection here, as I'm sympathetic to Hale's view about the relationship between absolute and broadly logical necessity. So let's proceed to note that there are two facets to Hale's argument: one about finding relevant possibilities; the

<sup>10</sup> Hale presents this as a reply to an objection to his first argument. However, the point he makes is independent of that argument.

<sup>11</sup> For arguments to this effect, see Shalkowski (1997) and Lowe (1998).

other about those possibilities being genuine (i.e., not *merely* broadly logical).

It's important to see that the concern about relevance is toothless on its own. Consider, for example, the thesis that our origins are essential: no person could have been produced by a different sperm and ovum. What possibilities are relevant to the truth of this thesis? Obviously enough, ones like *Bob could have had different parents*. (Indeed, we can present them as a reasonably well-defined class: i.e., the set of all propositions such that if  $p$ , then it isn't the case that our origins are essential.) Granted, conceivability's defenders don't know the truth values of all these possibility claims, nor even what all the relevant claims are (except in the general way just mentioned). But why should they need to?

For the argument to have any force, the concern about relevance needs to be combined with a demand that the possibilities in question be genuine. This is a high standard, and it's hard to make sense of it apart from the assumption that, while it's fine to be a fallibilist, it *isn't* fine to have your basic method be fallible. On this view, there's nothing wrong with the view that we sometimes make mistakes in the way we handle our evidence. Among other things, weariness, laziness, and false background beliefs can ruin our inferences. However, unless our evidence entails that for which it's evidence, it doesn't justify belief. For suppose that weren't Hale's view. Then, conceivability's defenders could insist there are some possibilities that they justifiably believe to be genuine—on the basis of seemingly successful conceivings—and they could rifle through *those* as they conduct a well-directed and thorough search. Why can't they take this line? Not because we have reason to think that some of their justified beliefs are false; Hale gives no argument to that effect. Rather, the view has to be that those beliefs aren't justified in the first place. But the problem can't just be that the source of justification is fallible, for reasons already discussed. Moreover, attempts to conceive are vulnerable to almost all the same problems as are attempts to assess broadly logical necessities.

All, that is, save one: broadly logical necessity entails absolute necessity, whereas broadly logical possibility doesn't entail absolute possibility. So it looks as though Hale's view is that evidence must entail that for which it's evidence.

Let's suppose that this is indeed Hale's view. It isn't an unreasonable one. Nevertheless, it also isn't one that he can expect proponents of conceivability to share. After all, the same consideration tells against induction, inference to the best explanation, and any other form of ampliative inference. Surely conceivability's defenders can't be blamed for siding with these standard modes of reasoning.

Hale might reply that his argument doesn't rest on a *general* principle about fallibilism, but only a *specific* principle about conceivability. That's a fair move, but it's worth noting that conceivability's defend-

ers certainly won't accept the specific principle unless the general one is plausible. And they'll insist—quite reasonably—that the general one isn't. So, it seems that Hale's second argument is unsuccessful too: he's given us no reason to think that defenders of conceivability will run afoul of Condition 1.

#### 4. The Prospects for Possibility-based Approaches

Hale's arguments don't show what they were supposed to show—namely, that we ought to endorse a necessity-based approach. It certainly doesn't follow, though, that possibility-based approaches are in the clear. In fact, I think we can use Hale's arguments to make two points about possibility-based approaches that clarify their prospects. First, those who endorse possibility-based views need to say more about our modal knowledge concerning logic and mathematics; if they can't, then they ought to take a symmetrical approach. Second, those who endorse possibility-based views may be committed to Peter van Inwagen's (1998) modal skepticism.

##### 4.1. *Logic and Mathematics*

Necessarily, if a conjunction is true, then both conjuncts are true. The challenge for possibility-based views is to provide some account of how we know this. Necessity-based views have no trouble here. Suppose that there are conceptual truths concerning logic and mathematics. If there are such truths, then it's plausible that we can use them to account for our knowledge of the corresponding modal truths. According to Hale, conceptual truths are broadly logical truths, and all broadly logical truths are absolutely necessary. So if we know as much, we can deduce that a particular conceptual truth is absolutely necessary. This sort of explanation isn't available given a possibility-based approach—or, at least, it isn't available unless you can give possibility-based justifications of the modal claims involved.

Likewise, possibility-based approaches don't allow you to say that proofs are evidence that truths are necessary. Saying that presupposes knowing both that the premises are necessary and that the form is necessarily truth-preserving. So either this move pushes the problem back, or it capitulates to the necessity-first approach.

Given the above, what are the prospects for giving such possibility-based justifications? I'm aware of only one attempt. In his defense of the sensuous imagination as a guide to possibility, W. D. Hart writes this about our knowledge of necessity:

...you have reason to believe necessary a truth whose necessity is the best explanation of your inability to imagine that it be false. [In

this way we] confine the inference [to a truth's necessity] to those cases where you cannot imagine its denial because of something intrinsic to the nature of its subject matter and to exclude those cases where you cannot imagine a very complex denial simply because its complexity exceeds the capacities of, let us suppose, all finite imaginations. (1988, 21)

This view probably implies that we don't know many philosophically interesting necessities. Suppose that Suzy can't imagine being born to different parents. Is the best explanation of the former failure that she couldn't have had different parents? It isn't likely. It seems that our imaginations are limited by what we take to be necessary—not what *is* necessary—and there is pretty good evidence that our tendency to attribute essential properties is based on a heuristic, as Cimpiana and Salomon (2014) argue. So, the best explanation of the imaginative failure is probably that a fast and frugal cognitive strategy is limiting the imagination, not that we've run up against an absolute necessity. Or, at the very least, this explanation is at least as good as its rival, which implies that we aren't justified in maintaining that the claims in question are necessary.

Nevertheless, Hart's view might be able to get modal knowledge concerning logic and mathematics. Defenders of possibility-based approaches could understand logical and mathematical proofs as establishing particular truths, and our failures to conceive of their contraries as evidence of their necessity—at least where the best explanation of our failure concerns the relations between the various logical and mathematical notions that are involved.

Of course, whether this is plausible depends on the merits of the other hypotheses. We might instead fail to imagine because of the limits of our cognitive machinery. Indeed, conventionalists about necessity use this point to argue for their position—one which Hart certainly rejects. Edward Craig (1985, 92), for example, thinks we should see the link between unimaginability and impossibility as one we forged, and one we could equally well abandon:

...the decision to adopt a theory (and especially, because of its general applicability, an arithmetic) which presents reality as imaginable rather than one which does not, is not to be thought of as on a par with the choice between two theories both of which fall into one or the other category. Whether we can achieve imaginable or only predictive theories is the question whether we can insightfully grasp the nature of the world or only learn to 'operate' it, to anticipate and control perceivable events.

In other words, we needn't require of our best theories—even our best mathematical theories—that they make the world imaginable. We

might drop that requirement, insisting only that they enable accurate predictions. Of course, being able to imagine the structure behind a phenomenon is part of what produces the feeling of explaining that phenomenon, and we value explanation a great deal. So, the latter choice is quite a radical one. Still, if we do require that our theories make the world imaginable, then although we'll count the unimaginable as impossible, we should recognize that we needn't have. This suggests that we're imposing a structure on the world that it doesn't actually have, which is just to say that the limits of that structure—namely, the necessities to which it commits us—are conventional. So while the conversation about modal epistemology tends to ignore conventionalism, a Hart-style reply to the problem of knowing logical and mathematical necessities brings it to the fore. Those who defend possibility-based views will need to find a different reply, show that something goes wrong in Craig's argument, or argue that there are independent grounds for rejecting conventionalism about necessity.

#### 4.2. *Modal Skepticism*

Let's grant that defenders of possibility-based approaches can stave off the threat of conventionalism. Even so, they'll have to be cautious as to how much they claim for our knowledge of necessity. As I'll argue in this section, they'll also need to be cautious about how much they claim for our knowledge of possibility. Hale's conceptual competence argument suggests that possibility-based approaches are probably committed to Peter van Inwagen's (1998) modal skepticism—or something in that neighborhood.

Van Inwagen characterizes modal skepticism in a few different ways, but we can get a handle on the position by considering his examples:

...we often do know modal propositions, ones that are of use to us in everyday life and in science and even in philosophy [e.g., that there could be a full-scale papier-mâché mock-up of a barn that looked like a real barn from a distance, or that the legs and top of this table might never have been joined to one another], but do not and cannot know [...] modal propositions like [it's possible for there to be a perfect being, it's possible that I exist and nothing material exist, and it's possible that there exist vast amounts of suffering for which there is no explanation]. (1998, p. 67–69, 81 n. 3)

In addition to these cases, he also denies that we know whether there could be transparent iron, or an iron bar that floats in water, or a naturally purple cow, or brain-state transfers, or teletransportation devices.

Why this skepticism? Van Inwagen defends it in several ways. But the worry relevant here is that, when we try to justify these claims

based on their imaginability or conceivability, we have no reason to think that we're imagining what we claim to be imagining. Granted, we're familiar enough with tables and their parts—and have seen enough of them disassembled—that we can be reasonably confident in imagining a particular table never having been assembled. But when we turn our minds to stranger cases, we lack the background knowledge that could underwrite confident claims to imagine. We can, of course, conjure a scenario that *looks* like a bar of iron floating on water, but since we can't imagine the case in much detail, we have no reason to think that we're imagining iron and water *per se*, rather than an opaque, gray object bobbing in a clear liquid. After all, actually *seeing* such objects would give you reason to think either that you aren't seeing iron or that you aren't seeing water—or both—given what we know about them. Why should the presumption be any different in the imagination?

Now recall how I criticized Hale's claim that conceptual competence involves knowing the conditions that are necessary for its application. I argued that, at least in many cases, it's sufficient to know conditions that are necessary-in-a-context; we needn't know anything stronger. And crucially, I said that when we're trying to conceive of perfectly ordinary objects in ordinary circumstances—e.g., my mug being several inches to the right of where it is on my desk—we may be justified in believing that we've conceived, since there's no reason to doubt that we can tell whether our concepts apply. However, I said that we may *not* be justified in weirder circumstances. This is because there's no reason to think that our concepts are much more precise than they need to be for the ends they serve, which means that there's no reason to think that we'll be able to tell whether we're really imagining iron and water, for example, once we begin considering sufficiently unusual situations. In short, there won't be enough information packed into those concepts to settle what would happen in sufficiently unusual situations, and the information that *is* there doesn't support the claim that we are indeed imagining iron floating in water. Hence, we won't be in a position to determine whether we're imagining what we're claiming to imagine. Presumably, we could make the same points about imagining my existing without anything material existing, or brain-state transfers. (In some cases, expertise might make a difference. There may, for example, be animal scientists who know enough about livestock pigmentation to make informed judgments about whether there could be naturally purple cows. It's less plausible that there are experts about perfect beings.) If this is right, then while Hale's conceptual-competence argument doesn't show that possibility-based approaches aren't viable, it does show that they'll require an extra measure of epistemic humility from their defenders.

Van Inwagen, of course, won't mind this conclusion, and nothing I've said suggests that anyone should. Still, it seems that possibility-

based approaches don't escape Hale's arguments unscathed. Even if Hart's move works, it looks like possibility approaches can't secure much knowledge of philosophically interesting necessity. And absent a reply to Craig's argument, there's a worry that these approaches lead to conventionalism about necessity. Finally, absent a reply to Hale's conceptual competence argument, possibility-based views will need to take modest positions on the scope of our modal knowledge even concerning possibility. Examining the architecture of modal knowledge may not support Hale's view, but doing so does offer a clearer picture of his rivals.

## References

- Chalmers, David. *The Character of Consciousness*. New York: Oxford University Press, 2010.
- Cimpiana, Andrei, and Erika Salomon. "The Inherence Heuristic: An Intuitive Means of Making Sense of the World, and a Potential Precursor to Psychological Essentialism," *Behavioral and Brain Sciences* 37 (2014): 461–480.
- Craig, Edward. "Arithmetic and Fact." *Exercises in Analysis: Essays by Students of Casimir Lewy*. Ed. Ian Hacking. New York: Cambridge University Press, 1985. 89–112.
- Giaquinto, Marcus. "The Linguistic View of *a Priori* Knowledge." *Philosophy* 83 (2008): 89–111.
- Hale, Bob. "Absolute Necessities." *Philosophical Perspectives* 10 (1996): 93–117.
- Hale, Bob. "Knowledge of Possibility and of Necessity." *Proceedings of the Aristotelian Society* 103 (2002): 1–20.
- Hale, Bob. "What is Absolute Necessity?." *Philosophia Scientiae* 16 (2012): 117–148.
- Hale, Bob. *Necessary Beings: An Essay on Ontology, Modality, and the Relations between them*. New York: Oxford University Press, 2013.
- Hanrahan, Rebecca. "Imagination and Possibility." *Philosophical Forum* 38 (2007): 125–146.
- Hart, WD. *The Engines of the Soul*. New York: Cambridge University Press, 1988.
- Horvath, Joachim. "Lowe on Modal Knowledge." *Thought* 3 (2014): 208–217.
- Kung, Peter. "Imagining as a Guide to Possibility." *Philosophy and Phenomenological Research* 81 (2010): 620–663.
- Leon, Felipe. *Armchair Knowledge and Modal Skepticism: A Rapprochement*. Philosophy: University of California, Riverside, Riverside, 2009.
- Lowe, EJ. *The Possibility of Metaphysics: Substance, Identity, and Time*. New York: Oxford University Press, 1998.
- Lowe, EJ. "What is the Source of Our Knowledge of Modal Truths?" *Mind* 121 (2012): 919–950.
- Quine, WV. "Two Dogmas of Empiricism" *The Philosophical Review* 60 (1951): 20–43.
- Roca-Royes, Sonia. "Mind Independence and Modal Empiricism." *CEUR-WS Proceedings* (2007): 117–135.
- Roca-Royes, Sonia. "Conceivability and *De Re* Modal Knowledge." *Noûs* 45 (2011): 22–49.
- Shalkowski, Scott. "Essentialism and Absolute Necessity." *Acta Analytica* 12 (1997): 41–56.
- Vaidya, Anand Jayprakash. "Understanding and Essence." *Philosophia* 38 (2010): 811–833.
- van Inwagen, Peter. "Modal Epistemology." *Philosophical Studies* 92 (1998): 67–84.
- van Woudenberg, René. "Conceivability and Modal Knowledge." *Metaphilosophy* 37 (2006): 210–221.
- Weisberg, Michael. "Water is Not H<sub>2</sub>O." *Philosophy of Chemistry*. Eds. D. Baird, E. Scerri, and L. McIntyre. Netherlands: Springer, 2006. 337–345.
- Williamson, Timothy. "Conceptual Truth." *Aristotelian Society Supplementary* 80 (2006): 1–41.
- Williamson, Timothy. *The Philosophy of Philosophy*. New York: Blackwell, 2007.
- Yablo, Stephen. "Is Conceivability a Guide to Possibility?" *Philosophy and Phenomenological Research* 53 (1993): 1–42.