

THEORY SELECTION IN MODAL EPISTEMOLOGY

Robert William Fischer

ABSTRACT

Accounts of modal knowledge are many and varied. How should we choose between them? I propose that we employ inference to the best explanation, and I suggest that there are three desiderata that we should use to rank hypotheses: conservatism, simplicity, and the ability to handle disagreement. After examining these desiderata, I contend that they can't be used to justify belief in the modal epistemology that fares best, but that they can justify our accepting it in an epistemically significant sense. Finally, I outline a modal epistemology that seems to do well by the standards I've recommended.

I. INTRODUCTION

In the seventies, modality made its comeback. Its epistemology wasn't far behind. And some forty years later, accounts of modal knowledge are many and varied, providing a bewildering array of options to any newcomer.¹ How should we choose between them? If we go by the literature, it may seem as though there is little consensus on this matter. Some seem concerned to preserve our intuitions about cases; others don't (e.g., Yablo 1993 vs. van Inwagen 1998). Some argue for a theory based on an antecedent metaphysical framework; others don't (e.g., Peacocke 1999 vs. Kung 2010). Some appear to be motivated by more general epistemological principles; others look for principles unique to the domain (e.g., Biggs 2011 vs. Rasmussen 2013). But this diversity is only apparent. The modal epistemology literature isn't characterized by brute disagreement about method; rather, it's characterized by attention to different desiderata. (If it were otherwise, we'd expect to see people arguing that alternative theories are unmotivated. But that charge rarely

appears.) What we need, therefore, is some clarity about what the desiderata are and how we ought to employ them.

My aim here is to provide that clarity. My strategy is to do so by way of two assumptions: first, that we should understand modal epistemologies as competing explanatory hypotheses; second, that we should use inference to the best explanation (IBE) to decide between them. Before going further, I should make two points. First, while I'm registering these claims as assumptions, I don't mean to suggest that there is nothing to be said in their defense. We should understand modal epistemologies as competing explanatory hypotheses because, on the face of it, we offer them to explain philosophical data: the distribution of justified and unjustified beliefs; the varying degrees of justification that our beliefs enjoy; why certain disagreements seem to be intractable; why certain apparent defeaters are only apparent; and so on. And we should employ IBE to choose between them because we have no real alternative. Perception is of no help here; there is no sample from which to make an inductive

argument; it isn't an *a priori* truth that some particular modal epistemology is the best one. Moreover, it would be irresponsible to endorse a modal epistemology *just* because it fits well with one's preferred semantic theory; or *just* because it implies that some theory of mind is correct; or *just* because it is intuitive or elegant or otherwise appealing. To take any such path would be to ignore the many other factors that are plainly relevant—that is, fit with all our *other* commitments. At least in this case, holistic evaluation is the only game in town. Given that the evaluation is of explanatory hypotheses, IBE is the obvious framework with which to begin.

The second point is this: as will be plain from what follows, we aren't settling much by deciding to use IBE. Its usefulness consists in the structure it provides for *raising* questions, not dodging them. Our two assumptions allow us to think of the desiderata in modal epistemology as being akin to the explanatory virtues—that is, the criteria for assessing the best explanation. Then, we can use debates about the merits of IBE to be more precise about the significance of various desiderata, as well as how they might be used in tandem to guide our selection.

In the next section, I discuss IBE in more detail, explaining how we can use it to illuminate issues in modal epistemology. Sections 3 and 4 are the heart of the paper. In the former, I map out the desiderata; in the latter, I suggest how they should be employed jointly. I conclude with some tentative remarks about the sort of theory that these desiderata seem to support.

2. IBE

Let's begin with a fairly standard formulation of IBE:

1. Hypothesis *h* would, if true, adequately explain facts f_1 – f_n .²
2. *h* would, if true, explain f_1 – f_n better than any of its competitors.
3. So, probably, *h* is true.

This formulation raises two issues. First, in what sense is *h* better than its competitors? Second, what's the appropriate attitude to have toward the argument's conclusion?

The classic answer to the first question is that *h* fares better than its competitors on the explanatory virtues: for example, simplicity, generality, conservatism, modesty, and predictive power.³ The classic problem with this answer is that the virtues aren't obviously truth-conducive.⁴ Thus Richard Fumerton:

A theory which is more consilient and simple than alternatives . . . is certainly more desirable than its competitors in the sense that it would be *nice* if it turned out to be true. But this not being the best of all possible worlds (some theologians aside) what would be nice is not always so. (1980, p. 596)

And if the virtues aren't individually truth-conducive—that is, if it's not the case that, in general, simpler hypotheses are more likely to be true, and likewise for each of the other virtues—then it isn't clear why we should think that IBE is truth-conducive.⁵

There are a few moves we might make here. First, we might offer an independent argument for IBE's truth-conduciveness, or at least for the conclusion that we are entitled to assume as much. (See, e.g., those like Harman 1986 and Lycan 1988, who take IBE to be the fundamental form of ampliative inference, and so epistemically legitimate if any ampliative inferences are.) This maneuver would allow us to concede that, for all we know, the explanatory virtues aren't individually truth-conducive; what matters is their collective use. Second, we might attempt to defend the truth-conduciveness of the individual virtues *as classically conceived*—that is, as general features of hypotheses. So, for example, we might specify a sense of simplicity—number of entities, number of kinds, and so on—and then present arguments for thinking that that feature is, relative to a field of hypotheses, an indicator of truth. (Swinburne takes this line, at least with respect to simplicity. He

maintains that “it is a fundamental *a priori* principle” that simplicity is evidence of truth (2001, p. 102).⁶ Third, we might take a page from Timothy Day and Harold Kincaid, who think of an appeal to a virtue as shorthand for an appeal to a certain domain-specific desideratum. On this view, we can’t assess the truth-conduciveness of simplicity *simpliciter*; we can only assess the truth-conduciveness of simplicity-in-population-biology or simplicity-in-climatology. By their lights:

IBE names an abstract pattern whose force and success depends on the specific background assumptions involved. Without substantive assumptions both about explanation in general and about specific empirical details, IBE is empty. In short, appeals to the best explanation are really implicit appeals to substantive empirical assumptions, not to some privileged form of inference. It is the substantive assumptions that do the real work. (1994, p. 282)

That is, insofar as we’re right to appeal to a particular virtue in a context, that’s because that virtue is shorthand for a set of background assumptions that guide inferences in that context. (They make the point in terms of specifically *empirical* assumptions, but what matters is that attending to those assumptions increases the likelihood of discovering the truth, not that they be empirical per se.)

Let’s turn to the second question that I identified above: what’s the appropriate attitude to have toward the conclusion of an IBE argument?

Belief is the obvious answer, but it isn’t the only one. There are various grades of acceptance that involve weaker epistemic commitments. We can see this by considering three related differences between belief and acceptance. First, belief aims at truth; acceptance need not. Instead, it can have either a pragmatic or an epistemic aim. Second, belief commits the believer to the truth of what’s believed; acceptance does not. For a given purpose, we may well accept a proposition that we believe to be false. And third, belief

should be based on evidence; acceptance need not be. The norms governing the acceptance of propositions depend on the purpose of accepting those propositions.⁷

With these contrasts in view, one thing should be plain: should it turn out that we are not justified in *believing* the conclusions of IBE arguments, we may still be justified in accepting them. Whether we are, of course, depends on the details of IBE, the merits of particular IBE arguments, and the relevant version of acceptance.

Of course, given some ways of construing acceptance, it isn’t an epistemologically interesting propositional attitude. But not on all. Take, for example, Greg Dawes’s version of acceptance. On his view, to accept a proposition is “to employ it as a premise in one’s reasoning, whether theoretical or practical, in any domain to which it might apply, with the goal of attaining knowledge” (2013, p. 68). The goal still involves truth—since knowledge entails it—but there is no claim to have achieved it. So we might be convinced that the best explanation is not, therefore, more likely to be true. Still, we may maintain that reasoning *from* the best explanation is more likely to lead to truth down the road. And if so, we may accept that explanation—in Dawes’s sense of that term.

To sum up: Our answers to the two questions (concerning the comparative evaluation of hypotheses and the appropriate attitude toward the best one) are related. Suppose we bracket the project of giving a general defense of IBE. Then, if we can defend the truth-conduciveness of the individual virtues, we may maintain that it is permissible to believe the conclusions of IBE arguments. If we can’t, then we still may be in a position to accept those conclusions—though it remains an open question whether it will be a variety of acceptance worth wanting. We must examine the virtues to say more. So, to make progress in our understanding of theory selection in modal epistemology, we

need to articulate and examine the various desiderata. Some may be applications of more general epistemic desiderata; others may be domain-specific. And once we've articulated them, we'll be in a position to begin assessing their roles in theory choice, as well as the appropriate attitude toward our selection.

3. THE VIRTUES

In what follows, I sketch three virtues that a modal epistemology can have: conservatism, simplicity, and the ability to handle disagreement. Granted, there are others: for example, theories are better insofar as they are internally coherent; they are better insofar as they make successful predictions. But these are not the interesting cases: not the former, because I can't think of a modal epistemology that's been scuttled by internal contradiction; not the latter, because predictive power isn't a virtue that philosophical theories tend to enjoy. Moreover, it's worth seeing how far we can get with a short list of virtues—which is, of course, the counsel of simplicity.

3.1. *Conservatism*

First, a modal epistemology is preferable insofar as it gets the cases right; it's less preferable insofar as it doesn't. That is, we want our theory to agree with our considered judgments about those modal beliefs that are and aren't justified.⁸ Insofar as our epistemology does this, and all other things being equal, it's preferable to one that doesn't.⁹ This is an expression of conservatism.

There have been various attempts to defend conservatism's epistemic credentials.¹⁰ By my lights, the most promising strategy is due to Quine and Ullian:

[A]t each step [opting for the conservative hypothesis] sacrifices as little as possible of the evidential support, whatever that may have been, that our overall system of beliefs has hitherto been enjoying. The truth may indeed be radically remote from our present system of beliefs, so that we may need a long series

of conservative steps to attain what might have been attained in one rash leap. The longer the leap, however, the more serious an angular error in the direction. For a leap in the dark the likelihood of a happy landing is severely limited. Conservatism holds out the advantages of limited liability and a maximum of live options for each next move. (1978, pp. 67–68)

Quine and Ullian aren't claiming that more conservative hypotheses are therefore more probable than not (construing probability objectively). Rather, they are claiming that conservatism is our best long-term strategy for acquiring true beliefs—or at least for getting as close as possible. Their justification for this claim resorts to the “step vs. leap” metaphor, but it isn't hard to unpack. Suppose that we're trying to decide between two hypotheses, one of which conflicts with our current system of beliefs, the other of which dovetails nicely. If all else is equal, then it would be absurd to choose the former, since we have evidence *against* that hypothesis in the form of the evidence *for* the beliefs with which it conflicts. So, *if* we are going to take on a new belief, then, *ceteris paribus*, it should be the more conservative one.

It's important to note that conservatism—as understood here—doesn't provide reason to take on a new belief. (It matters that the conclusion of the last paragraph is conditional.) If we're justified in believing a hypothesis, then we have good evidence for it. An explanatory virtue is not evidence. Rather, after we've collected the evidence and identified the hypotheses that could explain it, conservatism helps us rank those hypotheses by their vulnerability to a defeater. In particular, it operates as a total evidence condition, discouraging us from ignoring the evidence we have for whatever else we believe. Hence, it discourages us from believing a hypothesis that's vulnerable to a *rebutting* defeater—that is, evidence that supports the falsity of the hypothesis in question, thereby casting doubt on the adequacy of our original reason for

believing the hypothesis. Granted, we will go wrong; we will reject true hypotheses in favor of preserving false beliefs. Nevertheless, conservatism remains a good strategy for getting closer to the truth—if only because the alternative (*not* attending to our total evidence) is certainly worse.

To see how this might play out in modal epistemology, consider Peter van Inwagen's assessment of Stephen Yablo's conceivability-based theory—namely, that it leads, with the help of a few assumptions, to modal skepticism (1998, pp. 77–81). In short, Yablo says we are justified in believing that *p* if we can objectually imagine (imagine *de re*) a world that verifies *p*.¹¹ Based on several considerations, van Inwagen argues that we can't objectually imagine worlds that are significantly different from our own. So if Yablo is correct, then, according to van Inwagen, we aren't justified in believing extraordinary possibility claims. If this is right, then the *fit with the cases* criterion counts against Yablo's epistemology.

(It doesn't count against it decisively. Conservatism only offers clear guidance when *ceteris is paribus*. And that may not be the case when we consider Yablo's view. At this stage, the point is just to see how the criterion works, not to advocate for or against a particular modal epistemology.)

Conservatism already functions this way in modal epistemology, though it's rare to see it cited explicitly. Reflect for a moment on the many reactions to the problem of *a posteriori* necessities. (In short: if there are *a posteriori* necessities, then surely there are some of which we're unaware; and if there are some of which we're unaware, then we have an undercutting defeater for taking *p*'s apparent possibility to support *p*'s genuine possibility.) Most modal epistemologists don't throw up their hands at this problem. Rather, they offer (alleged) solutions. But why bother? It makes sense to offer solutions if we have some reason to avoid skepticism—for example,

that we took the challenged beliefs to be justified. Otherwise, the move seems unmotivated. Correspondingly, when philosophers *embrace* skepticism, they argue as though this is a cost to be justified.¹² Why? Because all parties are sensitive to the importance of conservatism.

3.2. *Simplicity*

W. D. Hart (2003) suggests that we should posit a primitive epistemic relationship between persons and modal facts—"grasp"—roughly akin to Russellian acquaintance. Hart is forced to make this suggestion because he is staunchly committed to (a) our having ample modal knowledge, (b) a neo-Quinean form of empiricism, and (c) the view that the truthmakers for modal propositions are abstracta. Whatever the merits of Hart's position, it's not simple. (The simpler conclusion is the disjunctive one: either (a), (b), or (c) is false.) And his position seems less plausible for that reason.

Simplicity is the foil of conservatism. Conservatism tells us to preserve our system of beliefs; simplicity advises us to modify it, and this to rid ourselves of the complexities we've harbored.¹³ Should we trust this advice? Some balk because they think of appeals to simplicity as being predicated on a baseless conjecture about the world. Consider, for example, Paul Benacerraf and Hilary Putnam's criticism of an argument for mathematical platonism:

[W]hy should the simplest and most conservative system (or rather, the system that best balances simplicity and conservatism, by our lights) have any tendency to be *true*? . . . It is hard enough to believe that the natural world is so nicely arranged that what is simplest, etc., by *our* lights is always the same as what is *true* (or, at least, *generally* the same as what is true); why should one believe that the universe of sets (or the totality of modal truths) is so nicely arranged that there is a pre-established harmony between *our* feelings of simplicity, etc., and *truth*? (1983, p. 34)

Others understand simplicity as a “merely” pragmatic consideration. Thus Bas van Fraassen:

Judgments of simplicity and explanatory power are the intuitive and natural vehicle for expressing our epistemic appraisal. [But these] are specifically human concerns, a function of our interests and pleasures, which make some theories more valuable or appealing to us than others. Values of this sort . . . provide reasons for using a theory, or contemplating it, whether or not we think it true, and cannot rationally guide our epistemic attitudes and decisions. For example, if it matters more to us to have one sort of question answered rather than another, that is no reason to think that a theory which answers more of the first sort of questions is more likely to be true (not even with the proviso ‘everything else being equal’). It is merely a reason to prefer that theory in another respect. (1980, p. 87)

However, simplicity need not be cashed out in a way that makes it vulnerable to these challenges. As with conservatism, the advantage of simplicity is not that it marks a particular hypothesis as probably true. Rather, it’s the best way to pursue truth in the long run.

There are a number of reasons for this. First, a simpler hypothesis has a greater prior probability than its more complex cousin, at least *ceteris paribus*. The probability of p is invariably higher than is the probability of p and q ; so, if we can get by without q , our hypothesis will be more probable as a result (i.e., p ’s posterior probability will be greater than the posterior probability of p and q —again, *ceteris paribus*). Second, simple hypotheses reduce the odds of making computational errors that might derail inquiry. This worry is not (just) about multiplying incorrectly. The more complex the theory, the greater the risk of our getting confused about the relationships between its parts, of smuggling in unstated assumptions, of misjudging how the theory should be applied in unusual contexts, and so on. We minimize these risks by preferring simplicity, making it more likely that if our

theory has to be amended or rejected, it’s for the right sorts of reasons. Third, simple hypotheses leave us in a better position to handle new data. Quine makes this point nicely:

If we encompass a set of data with a hypothesis involving the fewest possible parameters, and then are constrained by further experiment to add another parameter, we are likely to view the emendation not as a refutation of the first result but as a confirmation plus a refinement; but if we have an extra parameter in the first hypothesis and are constrained by further experiment to alter it, we view the emendation as a refutation and revision. Here again the simpler hypothesis, the one with fewer parameters, is initially the more probable simply because a wider range of possible subsequent findings is classified as favorable to it. (1963, p. 105)

In this last respect, simplicity amounts to a rejection of Popper’s call to make bold conjectures. But that’s as it should be. Bold conjectures aren’t well-supported by evidence; they’re designed to be falsified. When we’re trying to decide which hypothesis to believe, we want to avoid risk, not embrace it.

Finally—and perhaps in light of the above—we can see “prefer the simple” as a fast and frugal heuristic that guards against ad hoc hypotheses. We don’t need justification to *abstain* from saying that the world is a certain way; rather, we need justification to make such claims. So if we can explain the data with a weaker theory, and we take it on a stronger one anyway, we should worry that we’ve done so for the wrong reasons—namely, preserving a pet theory. Thus, “prefer the simple” checks the temptation to resist theoretical upheaval. This isn’t evidence that in a field of competitors, the simplest hypothesis is more probable than not. But it does suggest that, if we want to reach the truth, it would be foolish not to take the road paved with simple hypotheses.¹⁴

In the case of Hart’s “grasp,” we might not immediately see its problems in the terms outlined above. Instead, we might want to

express our reservations with a series of questions: Do we ever grasp the wrong thing, or only partially grasp the right thing? If so, what explains this? And can we detect such errors? If we can't, then what's the epistemic significance of this fact? Furthermore, how did grasp evolve? What advantages did it confer on our ancestors? Is there reason to suppose that it's a uniquely human trait, or should we expect other species to have it as well? I think these questions just illustrate the costs of increasing qualitative complexity—in this case, postulating a new kind of relation. Hart's view is bound to involve features that we find it difficult to deploy; it is a bold conjecture (and yet one not obviously falsifiable); it appears to be ad hoc. So, simplicity counsels against it.

For the same reason, simplicity may also counsel against accepting views that involve a fortuitous correlation between our modal beliefs and modal reality. For example, Scott Sturgeon argues that modal rationalism—like that of (Chalmers 2002)—is vulnerable to “the Objection from Magic.” That is, such accounts posit a perfect and inexplicable correlation between primary, positive, ideal conceivability, and metaphysical possibility (see Sturgeon 2010; cf. Kallestrup 2006).¹⁵ If Sturgeon is right, then modal rationalism might not be much of an advance over “grasp.” Similarly, simplicity might counsel against the theistic reliabilism that O'Connor (2008) develops based on Plantinga (1993). According to O'Connor, God ensures that our modalizing produces true beliefs (when performed properly in the right environment). This view can't be faulted for positing an *inexplicable* correlation between our modal beliefs and modal reality—since God stands in the gap—but insofar as theism is itself an unduly complex hypothesis, simplicity isn't on O'Connor's side.

On one level, all these examples highlight the difficulty of handling “the integration challenge”—the term coined in Peacocke

(1999) for the task of reconciling our epistemological and metaphysical convictions. This challenge is especially difficult if we deny that we are causally related to the truthmakers for modal claims (because we take them to be abstracta, or Lewis's concrete worlds, etc.) yet remain optimistic about the scope of our modal knowledge.¹⁶ It wouldn't be surprising, then, if simplicity tended to push us toward various ways of reducing modality, thereby dissolving the problems that animate the modal epistemology literature. But insofar as we resist the reductive move, I suspect that simplicity recommends either (a) greater epistemic humility, or (b) beginning with views that extend the epistemic resources we use in other domains—for example, induction, abduction, analogical arguments, and so forth—rather than postulating new ones.¹⁷ (I'll return to this thought later on.)

3.3. *Handling Disagreement*

Thirdly, a theory is preferable inasmuch as it provides tools with which to explain—and ideally to resolve—modal disagreements. Handling disagreement is not a standard explanatory virtue. In the present context, however, it deserves special mention.¹⁸ First, we might argue for its importance by maintaining, with Yablo (1993), that this helps to protect against *brute* disagreement, which can be evidence that the disagreement isn't about an objective fact (in some sense of that vexed phrase; Yablo's line on this draws on Wright 1986). After all, if our modal beliefs aren't about modal facts, then our work in modal epistemology is almost done; after we develop our error theory, we can go home.¹⁹ Thus handling disagreement is a way to avoid a potential defeater: namely that we have reason to think that there are no modal facts, and thus reason to doubt that our modal beliefs can be either true or false.

But even if we part ways with Yablo, maintaining that brute disagreement does *not* count against the objectivity of the fact in question,

it's surely the case that a brute disagreement about p can undermine our justification for believing p . (We can doubt our methods for acquiring beliefs about x without doubting that x is real.) Again, this is not to present handling disagreement as a mark of truth, but only as a way to avoid a criticism. Insofar as we take ourselves to have justified beliefs about certain modal matters, and insofar as those matters are contentious, we have cause to seek a modal epistemology that will help us to diagnose the controversy. Otherwise, it may be difficult to avoid the conclusion that the disagreement is brute, and thus that our justification was illusory. This isn't just an expression of conservatism—though it's that as well. Additionally, it's a way of replying to the charge that *if* we've struck on the truth, it's only by chance. If we can explain where our peers go wrong, then we're in a better position to explain what it takes to go right, and thus why we might be reasonable in trusting our judgments about what is and isn't possible.

There are plenty of proposals about the nature of modal error, which can usually be reframed as accounts of how to negotiate modal disagreements: for example, Yablo (1990, 1993), Chalmers (2002), Bealer (2004), Gregory (2010), and Sturgeon (2010). But my own view is that insofar as any account is promising, it's because it draws our attention to the larger doxastic context in which modal judgment occurs.

To see what I mean, let's begin with the observation that there is something appealing about the perceptual metaphor we sometimes employ when discussing certain modal phenomena. When I contemplate red and blue, I just "see" that they are incompatible, that it is *impossible* for any wholly red surface to be wholly blue. But we should admit that the story has to be far more complicated than the metaphor suggests. Consider a *thermochromic* surface—that is, a surface composed of a substance such that its color varies depending on its temperature. (Liquid crystals, the

substances used to make many thermometers and mood rings, are prime examples.) There certainly are thermochromic surfaces that can be wholly red at one time and wholly blue at another. Is this a counter-example to the impossibility claim above? Perhaps not. Perhaps this example only shows that the original formulation of the claim was sloppy, and the correct thing to say is that no surface can be wholly red and wholly blue *simultaneously*. But now we're faced with a hard question: on what basis do we deny that the surface is wholly red and wholly blue simultaneously? Why, for example, do we think that the relationship between *being wholly red* and *being wholly blue* is not akin to that between *being fragile* and *being resonant*? Those properties may not be manifested simultaneously either, but there is no obvious reason to deny that, say, a wineglass has them simultaneously. The point here is not that we should adopt some quirky version of color dispositionalism. Rather, the point is that, when we say we "see" that being wholly red is incompatible with simultaneously being wholly blue, we're assuming—*inter alia*, and probably without realizing it—that being a particular color is *not* like a thing's having a dispositional property.²⁰

Such assumptions often sit in the background when we modalize; it's rare that we make them explicit and put them on the table for consideration. They also tend to be the differences that make disagreements so hard to resolve. Consider, for example, the various Kripke-style arguments for dualism.²¹ Such arguments rely on certain modal premises, all of which have faced sustained criticism, though this fact seems not to demoralize their advocates. As far as I can tell, this is because one philosopher's view of x might be so different from another's that the first sees no reason to think that Fx is possible, while it seems obvious to the second that it is. For example, the first thinks that pain is characterized by its feel; the second that

it's characterized by its functional role in an organism.²² These sorts of standoffs have made us wary of such modal arguments, since they appear to circumvent, rather than face squarely, the substantive issues: namely what's the right theory of pain? Progress with respect to modal disagreement, then, may involve shifting our attention to background theories, implicit assumptions, and their ilk. A modal epistemology can aid us in making this shift by drawing attention to the ways in which our justification is mediated, at least in part, by their doxastic context.

Again, an epistemology's ability to handle disagreement is not a mark of its truth. Nevertheless, like conservatism and simplicity, it's a valuable constraint in our search for the truth. For without it, we might well adopt a theory that lacks the resources to answer two allegations that should worry anyone who takes modal knowledge to be a worthy subject of inquiry: first, that there are no modal facts; second, that we aren't justified in believing any controversial modal claims.

4. BELIEF OR ACCEPTANCE?

Before going further, let's recall the formulation of IBE with which we began:

1. Hypothesis h would, if true, adequately explain facts f_1 – f_n .
2. h would, if true, explain f_1 – f_n better than any of its competitors.
3. So, probably, h is true.

This formulation raises two questions. First, in what sense is h the best explanation? Second, what's the appropriate attitude toward the conclusion of an IBE argument? In the last section, I offered three virtues that a modal epistemology can have: conservatism, simplicity, and the ability to handle disagreement. On my view, then, if h is a modal epistemology, it qualifies as the best explanation in virtue of earning the best overall score on these three virtues. I can't precisely unpack the "score" metaphor I just

deployed. Nevertheless, I take that we are somewhat familiar with making such comparative evaluations in practice.

This still leaves the problem of balancing three desiderata. We can make some progress on this issue by reflecting on the virtues we've discussed. Conservatism, recall, is akin to a total evidence condition; it guards against adopting hypotheses that face rebutting defeaters. So, we can think of it as having a regulatory function, to be deployed *after* the other virtues, not before. (Deploying it first would be to risk dogmatism; deploying it simultaneously creates the problem rather than solving it.) This two-stage approach makes our task somewhat easier. Now, our aim can be to look for simple hypotheses that can handle disagreement. Once we have a field of competitors meeting this description, we can assess them in terms of conservatism. This is hardly a decision procedure, but it's nevertheless an advance over a single-stage process, which provides no real advice about how to adjudicate the demands of competing desiderata.

This leaves us with our question about the appropriate attitude. In the preceding section, I was careful in the way that I articulated the epistemic significance of the virtues. By my lights, being the simplest hypothesis in a field of competitors is not evidence that the hypothesis is in fact true. So, I grant that we shouldn't believe a hypothesis on the basis of its simplicity. Still, when choosing between competing hypotheses, we have every reason to prefer the simplest—at least if *ceteris paribus*—because our aim of securing the truth in the long run is better served thereby. I say the same of the other virtues. Therefore, we needn't settle for an epistemologically uninteresting attitude toward the conclusions of IBE arguments in modal epistemology; the discussion warrants more than using-for-the-sake-of-convenience. Indeed, we can recommend in modal epistemology what Dawes recommends generally: if a hypothesis fares

the best on the explanatory virtues, then we should accept it—*not* believe it. Nevertheless, we haven't crossed over into the merely pragmatic. Truth, after all, is the aim: to accept a proposition is “to employ it as a premise in one's reasoning, whether theoretical or practical, in any domain to which it might apply, with the goal of attaining knowledge” (Dawes 2013, p. 68). Dawesian acceptance is a properly epistemic propositional attitude.²³

5. THE WAY FORWARD

In the preceding sections, I've identified the various explanatory virtues in modal epistemology, I've indicated how they should be used, and I've taken a position on the attitude they can warrant. In this final section, I want to make some very tentative remarks about the sort of view that I take the virtues to support.

Consider the roughest sketch of an *argument-based modal epistemology* (AME)—inspired by Roca-Royes (2007) and Leon (2009).²⁴ On this view, if we are justified in believing that *p* is possible, then either:

1. we justifiably believe that *p* is actual, and we inferred its possibility on this basis, or
2. we could (in a relatively modest sense of “could”) give an argument for *p*'s possibility based on premises we justifiably believe (where that argument is such that we would be justified in believing that the premises support the conclusion).²⁵

Likewise, if we are justified in believing that *p* is necessary, then:

3. we could (in a relatively modest sense of “could”) give an argument for *p*'s necessity based on premises we justifiably believe (again, where that argument is such that we would be justified in believing that the premises support the conclusion).

AME doesn't require us to have an argument for, say, *p*'s necessity to be justified in believing that it's necessary. It only requires that we be able to offer one should we be pressed. Moreover, AME places no

restrictions on the sorts of arguments that we can offer for modal claims. We might deduce that an object has a modal property because it falls under a certain kind, and we justifiably believe (on some basis or other) that all instances of the kind have that property; we might make an inductive argument based on what we've observed about other instances of the kind; we might posit a modal property as the best explanation of some data; we might make an analogical argument based our observations of instances of related kinds. What matters isn't the form of the argument, but our being able to offer a reason for our belief.

In principle, then, AME is perfectly compatible with conceivability, imaginability, or intuition-based arguments. However, it insists that we justifiably believe the premises of the argument we could offer. So, if we justifiably believe the following premises:

1. If we can conceive that *p*, then *p* is possible.
2. We can conceive that *p*.

. . . then we are justified in believing the conclusion:

3. Therefore, *p* is possible.

No one should object to that conditional. Rather, people should deny the antecedent: if they aren't convinced of (3), they should insist that we aren't justified in believing either (1) or (2) (or both)—and then they should offer an argument for that conclusion (as they doubtlessly will).²⁶

One appealing feature of AME is that it provides us with a battery of questions to ask when faced with disagreement over a possibility claim—namely all those questions we ask when we encounter *any* argument of the relevant type. Suppose the argument is deductive. Is it valid? Are the premises true? Suppose the argument is inductive. Is the sample large enough? Do we have good reason to think that it's representative? Suppose the argument is analogical. How good is the analogy? Would we expect the mechanism

that explains x 's being F to be operative in y 's case, so that it's probably F too? And so forth. AME treats modal disagreement as the familiar problem of argument evaluation.

AME is also a simple theory. It invokes epistemic tools with which we are already familiar; it requires no new primitives; it involves no commitment to a pre-established harmony between our beliefs and modal reality; it is compatible with a variety of metaphysical frameworks. For these very reasons, AME also has some claim to be a conservative theory. Additionally, it can claim that title on the basis of getting the cases right. If Jones takes himself to be justified in believing that p is possible, it would be awfully surprising to discover that Jones denies he has an argument for p 's possibility—or thinks he has an argument, but isn't justified in believing the premises; or he doubts that the premises support the conclusion. And if we think that Jones is wrong about p 's possibility, it would be awfully surprising were we to be impressed by the merits of Jones's argument for that conclusion. So if we imagine ourselves in Jones's shoes, the odds are good that we'll agree with how AME sorts

cases; if we imagine ourselves as the critics, the odds are no worse.

AME is a deflationary theory in the sense that it discourages us from focusing on a single route to modal knowledge; on this view, it's a mistake to assume that, if our beliefs are justified, then they're backed by primary positive ideal conceivability, or principles of plenitude, or any other one-size-fits-all approach to our knowledge of possibility. Instead, it invites us to consider many kinds of arguments, sorting through them using the methods of argument evaluation that are our stock-in-trade. In practice, of course, we do this anyway. Perhaps we should bring theory into alignment.

Space doesn't permit comparing this approach with its competitors, so the above doesn't constitute much of an argument for AME. Moreover, while AME appears to fare quite well on the explanatory virtues, perhaps they don't ultimately support it. My guess, though, is that the best hypothesis is in this neighborhood.

Texas State University

NOTES

A few years ago, in the wake of a talk in which I claimed to be defending a modal epistemology, Peter van Inwagen told me that I was actually defending a meta-epistemology of modality. He was right; I thank him for making the observation. I'm also grateful to Eric Gilbertson for reading various drafts and patiently discussing these issues with me. Finally, I owe a great deal to two anonymous reviewers.

1. See McLeod (2005); Vaidya (2011); and Evinne (2008) for overviews of the literature.
2. The adequacy requirement in P1 ensures that h isn't merely the best of a bad lot. Granted, it matters a great deal how we articulate the relevant standard of adequacy, though I won't explore that issue here. For more on the "best of a bad lot" objection, see van Fraassen (1989); and Lipton (2004).
3. For this list, see Quine and Ullian (1978, pp. 64–82). See also Lycan (1988, p. 130); and Lipton (2004, p. 122).
4. For criticisms along these lines, see Gunner (1967); Benacerraf and Putnam (1983); van Fraassen (1989); Day and Kincaid (1994); Bueno and Shalkowski (2004); and Shalkowski (2010).
5. For more on this line of argument, see Fischer (2014).
6. For other routes to simplicity's truth-conduciveness, see Quine (1963); Sober (1981); and Kelly (2007).

7. See Tuomela (2000) for further discussion of the differences between belief and acceptance.

8. An example on the justified side: my coffee could be warmer than it is; an example on the unjustified side: the liquid I used to make my coffee (i.e., water) could be XYZ.

9. Likewise, I submit that a modal epistemology is preferable insofar as it maintains that the methods by which we acquire justified modal beliefs are the methods by which we *tend* to acquire modal beliefs. So whatever we believe about the methods by which we acquire justification about modal matters, a theory is better insofar as it spares us from revising these beliefs.

10. See Harman (1986); Lycan (1988); Kvanvig (1989); Huemer (2001); and McCain (2008).

11. As he puts it:

Imagining can be either propositional—imagining that there is a tiger behind the curtain—or objectual—imagining the tiger itself. To be sure, in imagining the tiger, I imagine it as endowed with certain properties, such as sitting behind the curtain or preparing to leap; and I may also imagine that it has those properties. So objectual imagining has in some cases a propositional accompaniment. Still the two kinds of imagining are distinct, for only the second has alethic content—the kind that can be evaluated as true or false—and only the first has referential content—the kind that purports to depict an object. (Yablo 1993, p. 27)

12. On the non-skeptical side, see, for example, Peacocke (1999); Chalmers (2002); and Williamson (2007); on the skeptical side, see van Inwagen (1998); Nozick (2003); and Hanrahan (2007).

13. There are, of course, various ways in which one hypothesis can be simpler than another: it postulates fewer objects; it postulates fewer kinds; it is syntactically simpler; it employs fewer argument patterns to achieve the same explanatory ends (as in Kitcher's brand of unification); and so forth. Some ways are more important than others, but this isn't the place to explore which and why. For further discussion, see Nolan (1997).

14. For a formal argument for the same conclusion, see Kelly (2007).

15. Of course, we might also doubt that the correlation obtains; on this point, see Vaidya (2008).

16. If the modal facts are as mind-independent as are the facts about the half-lives of isotopes, then there is no apparent reason to think that modal reality is organized in a way that would be epistemically convenient for us. Hence it's hard to see how we could be in a good epistemic position with respect to the modal facts. I would think that, at best, our epistemic situation vis-à-vis the modal facts would be akin to our epistemic situation vis-à-vis the facts outside our light cone. (On the assumption that no cause can travel faster than the speed of light, we can say, very roughly, that our light cone is the region of spacetime containing both (a) the events that could causally affect us, and (b) the events on which we might have some causal impact. See Huggett (2010). Thus, we'll be able to make a few claims with a high degree of confidence, but most we'll need to hold rather loosely, and there will be a great deal about which we'll need to confess ignorance. We can motivate a different form of skepticism with an evolutionary argument. There are tolerably good reasons to think that we are built to assess whether a given state of affairs is physically possible: our ancestors needed to be able to assess risks and opportunities, and our modal notions allow us to do so. If *p* is impossible, then we can ignore it in our deliberations; if it isn't, then we need to factor it in. And, depending on *p*'s valence for an agent at *t*, *p* qualifies as either a risk or an opportunity for that agent at *t*. But it's not entirely clear why our ancestors would need to be able to assess yet stronger modalities; there is no obvious adaptive advantage to getting things right in all worlds, since, *ex hypothesi*, there were no physically impossible selection pressures. For more on this line of reasoning, see Nozick (2003) and Nichols (2006).

17. See, for example, Roca-Royes (2007); Leon (2009); and Biggs (2010).

18. There may well be a way to construe this virtue in terms of conservatism or simplicity, but in the present context, little would be gained from trying.

19. There are alternate accounts of modality on which this isn't true; see Thomasson (2007).

20. Someone might object that this isn't an assumption—it's a conclusion, and one that's warranted based on the seeing. I find this implausible. First, if it's a conclusion, then it appears to be one that's reached without considering alternatives; and, if that's right, then one wonders why we should take it to be warranted. Second, note that we have a fairly clear account of how we can "perceive" sophisticated facts, thanks to relevant prior knowledge: the neurologist "sees" frontal cortical atrophy on a CT scan, whereas I see a Rorschach test, because of what the neurologist knows about brains, CT scans, and the like. And absent such knowledge, the perception would be completely mysterious. So, barring considerations to the contrary, we should offer a model of our perception of modal facts that's in line with our model of the perception of actuality. (Of course, in the modal case, the assumption is no doubt implicit, but that makes no difference: what matters is that background information constrains the interpretation of the data.)

21. See, for example, Kripke (1980); Taliaferro (1986); Hart (1988); Yablo (1990); and Chalmers (1996).

22. This is, essentially, the criticism leveled in Jackson, Pargetter, and Prior (1982).

23. That said, I don't mean to imply that the case for acceptance is watertight. Our judgment to accept, rather than believe, is based on our threshold for epistemic risk. Our threshold of epistemic risk is the point at which we are unwilling to venture belief based on the evidence; it's the point at which we judge that the risk of believing falsely outweighs the possibility of believing truth. If we put thresholds on a simple spectrum, then it would be one that runs from skeptic to sucker. Skeptics have a very low threshold for epistemic risk, and hence they deny that we know much of anything. Suckers, by contrast, have a high threshold; they are eternal optimists when it comes to the scope of their knowledge. Let's grant that we shouldn't set our threshold at either extreme. Might there be a range of acceptable thresholds in the middle? Perhaps so. And in the present case, I see no obvious reason to deny that, with the right argument behind it, tentative belief could be an appropriate attitude. For more on ranges of epistemic risk, see Riggs (2008).

24. None of these philosophers endorse the simple version of AME presented here.

25. I'm here interested in primary sources of justification, so I'm ignoring testimony, memory, and so on.

26. We might worry that the debate here is about the truth of a particular modal epistemology; hence, AME just pushes back the question. I don't think that this is correct. We usually offer a modal epistemology as *the whole story* about modal knowledge (or justification, etc.). But the debate here is not about a theory that might tell the whole story about modal knowledge. The debate is about the truth of an epistemic principle and a psychological claim. Both the principle and the claim are subject to various interpretations and could be motivated in any number of ways. The controversy doesn't concern, say, the viability of Yablo's modal epistemology; rather, it's about whether we can find interpretations of the principle and the claim that will deliver the conclusion. Moreover, it's no part of the above argument that conceiving is the only route to knowledge of unrealized possibilities. And if the point of offering the argument is to defend *p*'s possibility, we wouldn't adopt this bold premise unless forced to do so. Granted, we *might* be forced to do so. And so it could turn out that, at the end of inquiry, AME delivers exactly the results we'd get from a (certain sort of) conceivability-based modal epistemology. But if so, that's because it turns out that a conceivability-based argument is the only good argument for believing that *p* is possible—not (just) because of some features intrinsic to conceiving.

REFERENCES

- Bealer, George. 2002. "Modal Epistemology and the Rationalist Renaissance," in *Conceivability and Possibility*, ed. Tamar Gendler and John Hawthorne (New York: Oxford University Press), pp. 71–125.
- . 2004. "The Origins of Modal Error," *Dialectica*, vol. 58, no. 1, pp. 11–42.

- Benacerraf, Paul, and Hilary Putnam. 1983. Introduction to *Philosophy of Mathematics: Selected Readings* (2nd edition), ed. Paul Benacerraf and Hilary Putnam (New York: Cambridge University Press), pp. 1–40.
- Biggs, Stephen. 2011. “Abduction and Modality,” *Philosophy and Phenomenological Research*, vol. 83, no. 2, pp. 283–326.
- Bueno, Otavio, and Scott Shalkowski. 2004. “Modal Realism and Modal Epistemology: A Huge Gap,” in *Modal Epistemology*, ed. Erik Weber and Tim De Mey (Brussels: Royal Flemish Academy of Belgium), pp. 93–106.
- Chalmers, David. 1996. *The Conscious Mind: In Search of a Fundamental Theory* (New York: Oxford University Press).
- . 2002. “Does Conceivability Entail Possibility?” in *Conceivability and Possibility*, ed. Tamar Gendler and John Hawthorne (New York: Oxford University Press), pp. 145–200.
- Dawes, Gregory W. 2013. “Belief Is Not the Issue: A Defence of Inference to the Best Explanation,” *Ratio*, vol. 26, no. 1, pp. 62–78.
- Day, Timothy, and Harold Kincaid. 1994. “Putting Inference to the Best Explanation in Its Place,” *Synthese*, vol. 98, no. 2, pp. 271–295.
- Evine, Simon. 2008. “Modal Epistemology: Our Knowledge of Necessity and Possibility,” *Philosophy Compass*, vol. 3, no. 4, pp. 664–684.
- Fischer, Robert William. 2014. “Why It Doesn’t Matter Whether the Virtues Are Truth-Conducive,” *Synthese*, vol. 191, no. 6, pp. 1059–1073.
- Fumerton, Richard. 1980. “Induction and Reasoning to the Best Explanation,” *Philosophy of Science*, vol. 47, no. 4, pp. 589–600.
- Geirsson, H. 2005. “Conceivability and Defeasible Modal Justification,” *Philosophical Studies*, vol. 122, no. 3, pp. 279–304.
- Gregory, Dominic. 2010. “Conceivability and Apparent Possibility,” in *Modality: Metaphysics, Logic, and Epistemology*, ed. Bob Hale and Aviv Hoffmann (New York: Oxford University Press), pp. 319–342.
- Gunner, D. L. 1967. “Professor Smart’s ‘Sensations and Brain Processes,’” in *The Identity Theory of Mind*, ed. Charles Frederick Presley (St. Lucia: University of Queensland Press), pp. 1–20.
- Hanrahan, Rebecca. 2007. “Imagination and Possibility,” *Philosophical Forum*, vol. 38, no. 2, pp. 125–146.
- Harman, Gilbert. 1986. *Change in View: Principles of Reasoning* (Cambridge, MA: MIT Press).
- Hart, W. D. 1988. *The Engines of the Soul* (New York: Cambridge University Press).
- . 2003. “The Music of Modality,” *Topoi*, vol. 22, no. 2, pp. 135–142.
- Hawke, Peter. 2011. “Van Inwagen’s Modal Skepticism,” *Philosophical Studies*, vol. 153, no. 3, pp. 351–364.
- Huemer, Michael. 2001. *Skepticism and the Veil of Perception* (Lanham, MD: Rowman & Littlefield).
- Huggett, Nick. 2010. *Everywhere and Everywhen: Adventures in Physics and Philosophy* (New York: Oxford University Press).
- Jackson, Frank, Robert Pargetter, and Elizabeth Prior. 1982. “Functionalism and Type-Type Identity Theories,” *Philosophical Studies*, vol. 42, no. 2, pp. 209–225.
- Kallestrup, Jesper. 2006. “Physicalism, Conceivability, and Strong Necessities,” *Synthese*, vol. 151, no. 2, pp. 273–295.
- Kelly, Kevin. 2007. “A New Solution to the Puzzle of Simplicity,” *Philosophy of Science*, vol. 74, no. 5, pp. 561–573.
- Kripke, Saul. 1980. *Naming and Necessity* (Cambridge, MA: Harvard University Press).
- Kung, Peter. 2010. “Imagining as a Guide to Possibility,” *Philosophy and Phenomenological Research*, vol. 81, no. 3, pp. 620–663.
- Kvanvig, Jonathan. 1989. “Conservatism and Its Virtues,” *Synthese*, vol. 79, no. 1, pp. 143–163.
- Leon, Felipe. 2009. “Armchair Knowledge and Modal Skepticism: A Rapprochement,” PhD diss., University of California, Riverside.

- Lewis, David. 2001. *On the Plurality of Worlds* (Malden, MA: Blackwell).
- Lipton, Peter. 2004. *Inference to the Best Explanation* (2nd edition) (New York: Routledge).
- Lycan, William. 1988. *Judgment and Justification* (New York: Cambridge University Press).
- McCain, Kevin. 2008. "The Virtues of Epistemic Conservatism," *Synthese*, vol. 164, no. 2, pp. 185–200.
- McLeod, Stephen. 2005. "Modal Epistemology," *Philosophical Books*, vol. 46, no. 3, pp. 235–245.
- Nichols, Shaun. 2006. "Imaginative Blocks and Impossibility: An Essay in Modal Psychology," in *The Architecture of the Imagination*, ed. Shaun Nichols (New York: Oxford University Press), pp. 237–256.
- Nolan, Daniel. 1997. "Quantitative Parsimony," *British Journal for the Philosophy of Science*, vol. 48, no. 3, pp. 329–343.
- Nozick, Robert. 2003. *Invariances* (Cambridge, MA: Harvard University Press).
- O'Connor, Timothy. 2008. *Theism and Ultimate Explanation: The Necessary Shape of Contingency* (Malden, MA: Blackwell).
- Peacocke, Christopher. 1999. *Being Known* (New York: Oxford University Press).
- Plantinga, Alvin. 1993. *Warrant and Proper Function* (New York: Oxford University Press).
- Quine, W. V. 1963. "On Simple Theories of a Complex World," *Synthese*, vol. 15, no. 1, pp. 103–106.
- Quine, W. V., and J. S. Ullian. 1978. *The Web of Belief* (2nd edition) (New York: Random House).
- Rasmussen, Joshua. 2013. "Continuity as a Guide to Possibility," *Australasian Journal of Philosophy*, vol. 92, no. 3, pp. 525–538.
- Riggs, Wayne. 2008. "Epistemic Risk and Relativism," *Acta Analytica*, vol. 23, no. 1, pp. 1–8.
- Roca-Royes, Sonia. 2007. "Mind Independence and Modal Empiricism," in *CEUR-WS Proceedings, Latin Meeting in Analytic Philosophy Genova*, September 20–22, pp. 117–135.
- Shalkowski, Scott. 2010. "IBE, GMR, and Metaphysical Projects," in *Modality: Metaphysics, Logic, and Epistemology*, ed. Bob Hale and Aviv Hoffmann (New York: Oxford University Press), pp. 167–187.
- Sober, Elliott. 1981. "The Principle of Parsimony," *British Journal for the Philosophy of Science*, vol. 32, no. 2, pp. 145–156.
- Sturgeon, Scott. 2010. "A Priorism and Modality," in *Modality: Metaphysics, Logic, and Epistemology*, ed. Bob Hale and Aviv Hoffmann (New York: Oxford University Press), pp. 275–310.
- Swinburne, Richard. 2001. *Epistemic Justification* (New York: Oxford University Press).
- Taliaferro, Charles. 1986. "A Modal Argument for Dualism," *Southern Journal of Philosophy*, vol. 24, no. 1, pp. 95–108.
- Thomasson, Amie. 2007. "Modal Normativism and the Methods of Metaphysics," *Philosophical Topics*, vol. 35, nos. 1–2, pp. 135–160.
- Tuomela, Raimo. 2000. "Belief versus Acceptance," *Philosophical Explorations*, vol. 3, no. 2, pp. 122–137.
- Vaidya, Anand. 2008. "Modal Rationalism and Modal Monism," *Erkenntnis*, vol. 68, no. 2, pp. 191–212.
- . 2011. "The Epistemology of Modality," in *The Stanford Encyclopedia of Philosophy* (Winter edition), ed. Edward N. Zalta. <http://plato.stanford.edu/archives/win2011/entries/modality-epistemology/>.
- van Fraassen, Bas. 1980. *The Scientific Image* (New York: Oxford University Press).
- . 1989. *Laws and Symmetry* (New York: Oxford University Press).
- van Inwagen, Peter. 1998. "Modal Epistemology," *Philosophical Studies*, vol. 92, pp. 67–84.
- Williamson, Timothy. 2007. *The Philosophy of Philosophy* (Malden, MA: Blackwell).
- Wright, Crispin. 1986. "Inventing Logical Necessity," in *Language, Mind and Logic*, ed. Jeremy Butterfield (Cambridge, UK: Cambridge University Press), pp. 187–210.
- Yablo, Stephen. 1990. "The Real Distinction between Mind and Body," *Canadian Journal of Philosophy*, supp. vol. 16, pp. 149–201.
- . 1993. "Is Conceivability a Guide to Possibility?," *Philosophy and Phenomenological Research*, vol. 53, no. 1, pp. 1–42.1.