

2 | Modal Knowledge

Modal beliefs are woven into the fabric of our conception of the world. Still, despite our inveterate modalizing, it is not obvious how our opinions about the possible and the necessary are justified. Claims about how some part of the world actually is are made true by looking at and theorizing about it. But (1) how do we rationally discern the *possible* truth of some claim that is actually false, or the *necessary* truth of some claim that is actually true? And (2) what, ideally, is the overall structure of our modal beliefs, and how do they inferentially connect with other beliefs?

I seek to offer an answer to these questions that, while highly sketchy and incomplete, is more plausible than the going alternatives. While it has various elements in common with other accounts, traditional and contemporary, their combination is perhaps original. My point of departure is Stephen Yablo's "Is Conceivability a Guide to Possibility?" (1993). Yablo's account is designed to answer the first of our questions, and when we consider that question in isolation from the second it can easily seem appealing. But when we do consider our second question, the appeal of Yablo's picture diminishes. I then consider a more recent attempt by Christopher Peacocke to provide the needed comprehensive perspective without relying upon the broadly Aristotelian view of the metaphysics of modality that I hinted at in the previous chapter. After arguing that Peacocke's theory is fundamentally unsatisfactory, I develop some of the elements of a neo-Aristotelian metaphysical view and then suggest a more *au courant* view of the epistemology of our modal beliefs.

Conceivability As Our Guide?

Yablo begins by claiming the indispensability of Hume's maxim that whatever is conceivable is possible – or rather, of a more modest variant, that

conceivability is evidence of possibility, that it provides *prima facie* justification for accepting the possibility thesis in question. Later, we will ask whether this maxim is worthy of the confidence Yablo and many others place in it. We need first to ask just what its central notion of conceivability comes to. Yablo nicely distinguishes several ideas (1993: 26):

p is conceivable just in case

- it is believable that p
- it is believable that possibly, p
- one can imagine justifiably believing that p
- one can imagine believing p truly
- one can imagine believing something true with one's actual p -thought.

He then identifies his own notion ("philosophical conceivability"), not with any of these, but thus:

p is conceivable for me if I can imagine a world that I take to verify p .

Yablo elucidates the analysans by saying that the content of one's imagining will be highly detailed (though not of course fully determinate) with respect to a specific situation within the world that bears directly on the truth or falsity of p , "while leaving matters visibly irrelevant to p 's truth value unspecified" (29). That is, I try to imagine a coherent, connected totality, most of whose matters, though taken to *be* determinate, are not determinately represented in my imagining; and this totality contains a situation I represent with sufficient determinacy to verify that p is true of it. If I can do so, p is conceivable for me, and I am *prima facie* justified in believing that p is possible.

Philosophers often do use something like this procedure in judging possibility claims, yet their results are sometimes at odds. (Descartes was able to conceive himself existing without his body, but Hobbes could not. Hume conceived an event occurring without a cause; not so, Leibniz.) *Other* basic methods of belief formation, such as perception and memory, yield conflicting results, but we have some theoretical understanding of how they go wrong when they do. One reason some doubt the legitimacy of many or all *modal* beliefs is the apparent lack of a plausible account of how modal error arises when one uses the ostensibly proper method. Yablo contends, however, that a natural account (for most cases) is in fact ready to hand: I wrongly conceive that p – that is, I conceive that p when p is in fact impossible – when there is a proposition q such that

- (a) q ;
- (b) if q , then necessarily not- p ; and
- (c) that I find p conceivable is explained by my unawareness or outright denial that a, and/or my unawareness or outright denial that b. (34)

Consider, for example, a chemist of a bygone era who, accepting an erroneous account of water's molecular structure, finds it conceivable that water exists in the absence of oxygen molecules. His use of the method of conceivability leads him astray because of his factual error in disbelieving that water is partly constituted by oxygen molecules. Or consider a counterpart who accepts the correct account, but still conceives that there might be water in an oxygen-less world because he denies that water has its structure essentially. Here, his error is rooted in a failure to recognize a theoretical claim about essence (corresponding to the schema given in clause b of Yablo's model).

Yablo further claims that we may explain erroneous verdicts of *inconceivability* in precisely parallel fashion. There is a proposition q such that:

- (a) q ;
- (b) if q , then possibly p ; and
- (c) that I find p inconceivable is explained by my unawareness or outright denial that a, and/or my unawareness or outright denial that b. (39, n. 73)

Sticking with our previous sort of example, suppose this time a chemist who wrongly judges it inconceivable that water exist in the absence of carbon (p). He might do so because he disbelieves that water is H_2O (our q), or because, despite his acceptance of this identity, he thinks any world having hydrogen and oxygen must also have carbon (and so denies that if q , then possibly p).

Let us agree that no account of the justification of modal beliefs, including Yablo's, will or need satisfy the skeptical demand for external proof of reliability; nor need one suppose perfect reliability in the modal sphere. Still, there is reason to be unsatisfied with Yablo's account of both the justification of modal beliefs and of how error arises. Consider first his account of modal error. Despite the examples Yablo presses, it is not typically misinformation about nonmodal fact (his a condition), such as the chemical structure of water molecules, that leads to conflicting judgments. Instead, the typical disputes are of two other types:

1. Disagreement over the modal implications of nonmodal facts (the b conditionals in Yablo's error schemas): Could Socrates, who is actually an embodied human, have been an alligator, or disembodied? Are his origins essential to him? Could the property of unit negative charge have played the causal role actually played by unit positive charge?
2. More general, theoretical disagreements where the role of imagination is far from clear: Could an event occur without any causal antecedent whatsoever? Is time travel possible?

Yablo does not illuminate why disagreements of the first sort occur while failing to even take account of the second sort. Central to Yablo's picture of how we determine modal truth and falsity is an act of imagining. However, in neither of these typical sorts of cases where philosophers obtain conflicting results is the problem a failure of imagination. Even when imagination of a *p*-verifying totality is employed, what may be in dispute is whether a situation that is more or less determinately represented in one's imagination is *possible*. (So says the mind-body identity theorist in response to Yablo's imagining himself in a disembodied state.) Alternatively, thinkers may reasonably dispute the possibility of the *wider* totality that an imagining conceives as embedding a determinate situation – that is, one may dispute that one is truly imagining a possible world.¹ One can hardly *see* that it is, given that almost all of its details are not represented. It may be that unrepresented matters Yablo himself would judge to be “visibly irrelevant” to *p*'s possibility are not so in fact. (So say Leibniz and Spinoza, but also theorists who are less extreme in their embrace of a more organic conception of whole possibilities.) Such disputes do not turn on empirical matters or imaginative failures or confusions, but are instead theoretical in nature, concerning the truth or falsity of candidate necessity principles.

The persistence of such disagreements engender a reasonable worry about our access to, or the very existence of, modal facts in the absence of any theoretical perspective on how some of us, at least, go wrong in these ways. Yablo fails to shed any light here. Why do our judgments fall out as they do, even when we are correctly apprised of the nonmodal facts and are sensitized to differences between conceivability in Yablo's sense and, for example, believability? Theorists disagree on certain alleged necessities that constrain possibility – necessities of identity, of origin, and of constitution, for example. What prompts these differing judgments? How might I go about seeking to improve my own judgments? And note further Yablo's blithe acceptance (admittedly, he has plenty of company nowadays) that many wider matters about a totality's overall structure and content are “visibly irrelevant” to the possibility of some localized scenario. This clearly

presumes the correctness of some sort or other of Humean ‘recombination’ or ‘duplication’ principle, on which, roughly, anything can coexist with anything else.² Rationalist philosophers of various stripes deny such principles. Even supposing the Hume crowd were right on this score, we’d like to be able to see why one might think so. How does giving a wide scope to possibility hook up with our other theoretical beliefs? Yablo does nothing beyond inviting us to accept that that’s the way it is with possibility, and those who think otherwise are simply mistaken.

Finally, let us return to Yablo’s central thesis: conceiving that p is *prima facie justification* for believing that p is possible. He does not tell his readers what, in his view, epistemic justification consists in. But it is hard to resist supposing that he has something like a deontological conception in mind, on which believing that p is justified just in case it is within one’s epistemic rights, or one is epistemically without fault or blame. Now the believings of Yabloesque modalizers may well be blameless in this way, while still leaving us without answers to the questions worrying the modal skeptics: what are the truthmakers for modal facts, and what accounts for our (presumed) rough reliability in ascertaining them?³

Modality a Matter of Principle?

Before setting out my own view, I will consider one other account of the epistemology of modality currently on offer. Christopher Peacocke advances a novel approach in an intricate, lengthy chapter of his recent monograph, *Being Known* (1999: ch. 4).⁴ Peacocke believes it enables him to forge a transparent link between the metaphysics and epistemology of modal theses – something he takes to require eschewing primitive modality as a basic feature of the world – without falling prey to the temptation of thinker-dependent strategies. On what he terms a “principle-based conception,” the understanding of modal concepts involves an implicit grasp of certain principles of possibility, which principles specify the truth conditions for modal statements in terms of nonmodal properties of our concepts and of the properties and objects in the world. (These are the Modal Extension Principle, a family of Constitutive Principles, and a Constrained Recombination Principle.) To the extent that I correctly judge the content of these principles and grasp all relevant concepts, I am immediately in a position to judge the truth or falsity of specific modal statements. None of this involves positing or ‘receiving an impression of’ primitive necessity in the world.

Peacocke’s Principles seek to elucidate the concept of possibility in terms of a notion of *admissible assignment*. Ignoring various points of detail, an

assignment accords each atomic concept a semantic value of the appropriate category (e.g., an object to a singular concept). Assignments, then, will be total specifications conforming only to broad, categorematic norms. As a result, some such assignments will assign, for example, married men to the concept *bachelor*. The crucial question, then, is how to spell out the class of *admissible* assignments (excluding ones such as those just noted) in a way that does not simply presuppose an intuitive notion of what is possible.

Peacocke assumes that the actual extension of a given concept is determined by a rule. The rule for the concept *bachelor*, for example, is that it is the intersection of the actual extensions of the concepts *man* and *unmarried*. The Modal Extension Principle (1999: 134–6) has the same rules constrain admissible assignments in one of two ways, depending on whether the concept is, in Kripke’s terminology, *de jure* rigid (such as those associated with proper names and natural kind terms):

For concepts which are *not de jure* rigid – An assignment *s* is admissible for concept C only if the semantic value of C according to *s* is the result of applying the same rule as is applied in the determination of the actual semantic value of C.

For concepts which are *de jure* rigid – Where the actual semantic value of concept C is A, an assignment *s* is admissible only if the semantic value of C is A.

So, assuming Kripke is right on the *de jure* rigidity of natural kinds, the semantic value of the concept *water* is H₂O on *all* admissible assignments. For the non-*de jure*-rigid concept *bachelor*, the semantic value on all admissible assignments will vary, as a function of the values of its constituent concepts of *man* and *unmarried*, where these likewise conform to the same rule governing their applicability in the actual world.

A second set of constraints on admissibility are given by a family of Constitutive Principles. These require that an assignment respect “what is constitutive of the objects, properties, and relations it mentions” (Peacocke 1999: 144). For objects and properties, these will include kind-essential features; for individual objects, these will also concern actual origins. Peacocke does not advance a definitive set of such principles. He instead contends that our very notion of possibility is bound up with the idea that some such principles constrain admissibility, and deems it a matter of difficult *metaphysical* research – not conceptual analysis – that such principles are true.

Finally, the Constrained Recombination Principle asserts that any assignment respecting the above principles is admissible. (149) A proposition is

possible, then, if and only if it is true according to some admissible assignment, and a proposition is necessary if and only if it is true according to every admissible assignment. (150)

Will positing principles that are bound up with our grasp of the concepts of possibility and necessity suffice for Peacocke's aim? Note that what Peacocke is trying to do is walk a fine line between thinker-dependency and primitive modality. He wants to ground *objective* modal truths in facts about the *actual world*, on a fully nonmodal conception. (Modal truths could then be a 'conservative extension' of the nonmodal. I say that this is Peacocke's intent, but he wavers here, claiming that even if our concept of the actual is partly modal, his approach will still do its intended job. I will defer discussion of this claim until later.) These facts are the rules determining the actual extension of concepts and the fundamental principles that determine what an object, property, or relation is. The crucial issues are whether ascertaining these facts can be explanatorily prior to our understanding of modal concepts and independent of an acceptance of primitively modal features, such as are posited by a traditional realist conception of causality.⁵

Let us consider first, and very briefly, the rules that are said to determine concepts. We will assume that this is the correct picture of the nature of concepts and ask: What does this picture imply or presuppose, by way of modality? One might worry that talk of rules involves, *inter alia*, some notion of what it is to conform to a rule, and this appears to be a modal notion. But Peacocke contends that there are plausible available accounts of what it is to follow a semantic rule (and so to employ a concept), and of the identity of a rule or concept, that do not presuppose any modal claims (159). As he notes, one might suppose the actual extension of the concept *diamond-shaped* to be precisely those objects "which, as they actually are, would produce a certain kind of experience in a properly perceiving subject" (135). Accordingly, a subject might be said to be employing the concept *diamond-shaped* just in case they are, for example, disposed to employ certain terms under these circumstances.⁶ Here, I wish only to note the crucial role of counterfactual claims in justifying our attributing such a disposition. As we will now see, counterfactuals likewise play an important role in the other half of Peacocke's actualist determination of modal properties – the consideration of the identity of objects and properties.

Peacocke's Constitutive Principles determine the identities of objects and properties (conceived realistically) in terms of select actual properties, those that make them to be the entities they are. These principles are not a reflection of our concepts or categorizing proclivities, or in any other way something that is determined by human cognitive activity. Instead, they reflect

metaphysical truths about the entities themselves – though not primitively essential properties. About ascertaining such principles, he says only that they involve difficult metaphysical research, hinting that we settle on them for broadly theoretical reasons. But, *given a nonmodal conception of the actual*, what basis could there be for making an objective distinction between feature *F* as constitutive of the *x*'s and *F* as merely accidentally present in all the *x*'s? The constitutive truths, it seems, could be operative only in our counterfactual reasoning.⁷ An urgent question, therefore, with respect to the constitutive principles as well as the rules governing concepts, is the status and nature of counterfactual reasoning. Peacocke himself seems to recognize this, when he suggests that the central point of having a concept of necessity is just its connection to counterfactuals within theoretical and practical reasoning (172–3). If this suggestion is correct, one had better not go on to explain, even in part, the nature of our practical and theoretical interests in counterfactual matters in terms of an interest in the genuine absolute possibilities (or relative possibilities that require definition in terms of absolute possibilities). But, the needs of Peacocke's own approach aside, it seems natural to do just that. In theoretical contexts, for example, we sometimes explore an assumption whose truth is uncertain for the purpose of deciding whether it has impossible implications. We do so because we believe that anything that implies a contradiction cannot be true, and so is not true. In a more practical vein, I wonder about what might have happened had I chosen a different profession just because I believe that alternative course to have been a real possibility for me, and I care about the possible value, whether a net gain or loss, that I relinquished. Whether theoretical or practical, counterfactual reasoning is significant because the possibilities (and impossibilities) have direct significance for our understanding of, and engagement with, the actual.

A further problem arises for Peacocke's view when we consider the modal status of the principles of possibility themselves, which are said to map facts about the (allegedly nonmodal) actual into truths about the possible and the necessary. Peacocke draws our attention to the recursive nature of the Modal Extension Principle. The Principle (in conjunction with the others) specifies the rule for determining the *actual* extension of the concept of an admissible assignment. This concept, in turn, helps to specify the rule for determining the actual extension of the concept of necessity: a proposition is necessary if and only if it is true according to all admissible assignments. When we then apply the modal extension principle to this last rule fixing the concept of necessity, it returns that in *any* admissible assignment, the rule for fixing the concept of necessity will be the very one just noted. Peacocke concludes, "Hence, according to any admissible

assignment, the above Characterization of Necessity will be true. But this is precisely what it is for it to be necessary, under the principle-based conception” (152).

But surely this is an illegitimate kind of bootstrapping. What such a self-subsuming principle lacks is a truthmaker – an independent ground of its truth. Suppose we delineate several rival sets of candidate principles of possibility, each possessing a recursive element, and ask, in virtue of what is Peacocke’s the correct set and the others wrong? I take it that Peacocke’s answer will be simply that his set captures our modal concepts and the others (to varying degrees) do not: “The theorist of the principle-based conception will say, though, that if something is counted as true by an admissible assignment, that is all there is to its being possible. Nothing more is required” (181). But this, surely, undercuts the significance of true modal statements. Yes, they are, on the proffered account, determined *in part* by independent features of the world – features of our other concepts and of ‘constitutive’ facts about objects and their properties. These facts, however, do nothing to ground their own extension into truths of possibility and necessity; *that* work is done solely by the principles of possibility. These principles do not answer to any further facts, but simply specify the modal concepts. So we are led to wonder, as we did in the previous chapter when considering deflationary accounts of modality, what difference modal concepts could make to our understanding of the world if Peacocke’s account of them were true.⁸

That Peacocke indeed presents, in the end, a radically deflationary account of modality can be seen by his remarks on the implications of a ‘partially modal conception of the actual’ itself. Would the vindication of such a conception undermine his principle-based account of necessity, or his theory of understanding? Not at all, he contends – all that a partially modal conception would establish is a kind of ‘local [meaning] holism’, such that modal and other concepts are interlinked (175). It will remain that the principles of possibility are ungrounded determiners of modal statements.

But, contrary to Peacocke, a modally tinged conception of actuality is not fundamentally about our conceptualizations but about the nature of the world – the character of its constituents objects, properties, and (thereby) causality itself. Such a conception is guided by the conviction that only an enriched metaphysics of the actual (one encompassing primitively modal features) can adequately sustain the work our concepts of causation and of object and property identity are called upon to do, including within contexts of counterfactual reasoning. The interlinking of modal and nonmodal concepts is a reflection of the fact that the concepts are meant to link to the

world by doing explanatory work. Again, without underlying truthmakers, this work is negated and the project becomes deflationary.⁹

Peacocke's project was to account for the epistemology of modal belief by tying our very grasp of modal concepts to an acceptance of principles specifying modal truth conditions in nonmodal terms. Its central failing is the deflationary character of its metaphysics, on account of which it falls prey to criticisms much like those of the metaphysical views given in the previous chapter.

The Theoretical Roles of Modal Claims: Towards a Modal Epistemology

Reflection on the failed attempts by Yablo and Peacocke underscores the need to show how modal and nonmodal beliefs and methods of belief formation may be integrated while not retreating from the evident fact that modal beliefs play a unique role within our theoretical understanding of things, which fact apparently requires that at least some modal truths are metaphysically primitive. To determine how to go from there in outlining a plausible account of modal truth and modal knowledge (or at any rate a robust degree of warrant or rational justification), let us review the core 'data' that an integrated theory of modality ought to illuminate:

General level

We have a strong propensity to accept certain very general (*sometimes* formal) propositions and to regard them as obvious and yet transcending obvious empirical truths about the world. And our *explanatory* schemes appear to be built on such general propositions, so that theoretical rationality is attainable for us only if our modal notions are fundamentally legitimate.

Object level

The notion of necessity is bound up with our grasp of things, properties, and causation. Thus, there is no *perception of the actual* that is not imbued with an implicit grasp and acceptance of a modal character to reality.¹⁰

We should do justice to these data while also:

1. Allowing that in ordinary contexts, people reliably track the instantiation of modal features in the world – in a very rough and ready way –

without attempting to develop or draw upon a theoretical perspective on modality that would enable them to correct their judgments or set them within a coherent explanatory framework. Sound theoretical perspective is surely capable of improving the reliability of modal judgment, but theoreticians and ordinary folk alike work from much the same base of basic modal judgments.

2. Explaining how there can be rough yet reliable tracking of modal features at a ‘commonsense’ level that nonetheless is matched by systematic error at a more ‘theoretical’ level. (Faulty theoretical perspectives are capable of *worsening* the reliability of modal judgment.)
3. Responding to the contention that nonaccidental belief requires perceptual or causal contact with the facts believed, in application to both object-based and general modal judgments.

Theoretical modal judgments

Let us first consider the epistemology of general theoretical modal judgments. As I argued in considering Quine’s skeptical position, some judgments concerning formal necessities are necessarily at work at the very outset of theoretical inquiry. So, assuming that rational theoretical belief is attainable at all, there must be rational or warranted a priori beliefs concerning these necessities.¹¹ Such *prima facie* warrant extends to virtually every competent thinker, as even the theoretically unsophisticated are dimly aware of the role of logical and mathematical necessities within explanation – that of structuring the space of possibilities to be considered in formulating empirical theories in response to experience.

For such unsophisticated thinkers, however, this awareness is commonly accompanied by an ignorance of rigorous procedures for precisely determining the facts so necessitated. The *honing and extending and occasional creation of formal methods*, as frequently occurs in the history of mathematics, allows errors to be corrected and judgments to be refined. Judgments are also properly modified by a wide variety of types of *indirect reflection*. One variety involves taking into account work that sets within a more general framework what hitherto had been seen as a largely complete and independent formal domain. Consider, for example, the impact of Riemann’s celebrated 1854 lecture that generalized the very conception of geometry to the study of n -dimensional manifolds in any kind of space.¹² This allows one to probe the source of previously foundational convictions, and on occasion come to see them as having been confused or conflated with true propositions that previous theorists lacked the resources to correctly identify. Here again, developments in geometry provide an unusually

dramatic instance of this with Hilbert's insight, itself building on the creative work of Riemann and others, that geometry is properly understood to be a formal or axiomatic discipline dealing with abstract relational structures rather than an a priori theory of the nature of physical space. In another kind of case, much the same improvement of perspective is provided by simple analogy with better-developed disciplines. Philosophy is rife with instances of creative thought being spurred in this way through reflections on scientific models, leading to a better grasp on possibilities of philosophical concern. A modest example of this is the range of 'Chaos' models in Chapter 4 of this book.

Finally, attention to the history of mathematics shows that revision of a priori beliefs is not always a simple matter of shifting from one set of beliefs to another. The development of key concepts has often been a long and arduous process, one littered with proposals that proved unfruitful or incoherent and were quickly abandoned. One salient example of this was the long process of attempting to clarify the notion of continuity, which eventually moved from geometric to algebraic characterizations, culminating in Dedekind's work in the nineteenth century. The long constructive phase was carried forward through a mixture of strongly and weakly justified a priori beliefs, confused notions used as guides, hunches, and the like. But the later mathematician can study its completed product, with distinctions properly drawn, and have as a result a rich, stable set of strongly justified a priori beliefs.¹³

The procedure of refinement and development of theoretical modal judgments is thus one of reflective equilibrium, much emphasized in other contexts, but here at work within the a priori realm. This immediately implies fallibilism, of course, which is contrary to what most philosophers of tradition have held.¹⁴ That older view seemed inevitable for those taking at face value the metaphor of 'directly seeing the necessity' of at least some elementary truths. But once we suppose that the sober truth of the matter is that there is but a strong disposition to believe of certain propositions that they are not just true but necessary, and that this disposition is fed by an often hidden, underlying commitment to the explanatory role of such necessities (or of related claims), the notion that we sometimes err is no more disturbing in the a priori realm than in the a posteriori – unless one is still plumping for the heroic, skepticism-slaying project of the classical foundationalist.¹⁵

It is striking how few traditional proponents of the a priori as a rich source of important truths seem to appreciate, much less squarely acknowledge, that our a priori judgments come with varying degrees of confidence, and rightly so.¹⁶ We all believe that *if Derrida is wise, then Derrida is wise*

is necessarily true with near maximal confidence (despite the considerable doubts of many of us about the truth of its antecedent). Likewise with $2 + 2 = 4$, for any positive integers x , y , and z , if x is greater than y and y is greater than z , then x is greater than z , and if an object is red, then it is colored. Many of us also believe in the necessary truth of *it is possible for there to be a disembodied thinker*, *there are universals*, and *W. V. Quine exists in more than one possible world*, but appropriately not with the same degree of confidence. Much of our philosophical theorizing consists in weighing the relative confidence we have in apparently conflicting a priori judgments in order to retain only those that we deem to be more likely. It would be unreasonably dogmatic, then, not to allow that I might – it is *epistemically* possible that I – be warranted in the future in revising some of my present a priori beliefs. (Note that in some cases, such as when we set certain beliefs into a newly developed and more comprehensive mathematical framework, the result could be a proper *increase* in our confidence in them.) This fallibilism is wholly consistent with my believing, at any given stage, that sundry propositions that I believe with less than maximal confidence nevertheless are unconditionally *necessary*.

Some who proclaim the downfall of the a priori upon clearer recognition of fallibility in the target domains are simply failing to distinguish the idea that the warrant for these beliefs is in some sense not essentially dependent on empirical information or experience – apart from whatever experience may be necessary for thinkers like us to form the relevant concepts – from the traditional, but here rejected theory that the underlying belief-forming mechanism is an unmediated, direct accessing of truths. Other critics of the a priori point to the fact that revisions in formal domains often occur in the course of *empirical* theory formation – or even, apparently, in direct response to empirical data.¹⁷ It is one thing, such critics will say, to allow that purely a priori reflection might occasion adjustments in some of our a priori beliefs. But once we recognize that reflection on empirical evidence has done so, we have effectively gutted the idea that the truth of such claims is independent of contingent, empirical facts, and so the *necessary* status of even those logical and mathematical propositions that withstand empirical theorizing is rendered implausible.

In responding to this challenge, it is important to separate cases. I have repeatedly emphasized the importance of what we might term ‘core’ a priori beliefs. While the borders of this notion are not sharp, the core will certainly include at least the validity of some central rules of inference (and necessary truth of their associated propositions) in first-order logic with identity, elementary arithmetic, and some methodological truisms guiding any form of empirical inquiry. *Pace* Quine, it is not conceivable that empirical evidence

could lead us to abandon any of these beliefs, and for two reasons. First, each of these beliefs are rationally warranted to a greater degree than could be the plausibility of any complicated argument from a conjunction of experience and well-supported theory to their falsity. When faced with an incongruence (contradiction being the limiting case), the rational course is to give up the less certain. Secondly, and most decisively, our ‘core’ a priori beliefs are essential to our understanding of empirical theoretical confirmation. Demonstrating a tension, incongruence, or contradiction between a core belief and a theoretical framework that has empirical support requires tacit acceptance of that very core belief.

But what of slightly less entrenched formal beliefs? Haven’t there already been instances where experience have led us either to discount or at least to question beliefs that were held with nearly the same degree of confidence as, for example, the Law of Noncontradiction? I suggest that a careful look at the cases typically cited in fact reveals that empirical evidence per se does not (properly) lead to the revision, or contemplation of revision, of putatively a priori beliefs. Instead, the empirical enterprise occasions continued reflection on the formal tools it uses, leading to developments sparked by purely a priori advancements. Rather than trying to argue the point at length here, I will just make a couple of particular observations that the reader, hopefully, will recognize as sensible, despite being overlooked by the more cavalier critics of the a priori.

The first observation concerns everyone’s favorite example of geometry. We are reminded that Kant was just the latest in a long line of the best mathematical thinkers who championed the a priority and necessity of Euclid’s flat, three-dimensional geometry (though for Kant this necessity attached to human thought rather than a mind-independent external space). But Kant and the rest were wrong, and we now have a well-confirmed empirical theory that shows that it likely isn’t even a contingent truth. But naysayers who trumpet this example curiously fail to highlight the obvious fact that the pre-nineteenth-century tendency to overproject necessity within geometry was fundamentally a result of failure of a priori imagination – a failure to grasp possibilities afforded in part by the independence of the parallel postulate from the other axioms of Euclidean geometry. (And note that even before the development of other systems of geometry based on many-parallels and no-parallels postulates – developments unprompted by physical evidence or theorizing – theorists regarded Euclid’s one-parallel postulate as less evident than his other axioms and sought to derive it from them.¹⁸)

My second point is that much of the impetus to reconsider certain parts of classical logic comes from a priori reflections on, for example, the nature

of concepts and their relationship to natural properties. (For more on this relationship, see the following section.) So some are led to reconsider the Law of Excluded Middle on the grounds that most of our concepts are vague – even if this is not so, or could not be so, for fundamental natural properties in the world.¹⁹

Finally, we come to the proposal of Putnam (1965: 75–101) and others (following the mathematician Neumann) that adequately characterizing surprising quantum phenomena requires us to adopt a ‘quantum logic’ that rejects the classical laws of distributivity for conjunction and disjunction. This claim is bound up not only with highly technical issues within the physical theory but also with the question of scientific realism with respect to quantum mechanics under one’s favored interpretation. It would take us too far afield to consider these matters – and needlessly, since the “science has uncovered the need for a revision of our logic” suggestion has not received widespread support. Here a fundamental procedural point suffices: insofar as this goes beyond the mere suggestion that one might usefully employ nonclassical algebras for modeling certain phenomena to the contention that these ‘odd’ data (within the context of a well-supported theoretical framework) ought to lead us to change our pattern of *reasoning*²⁰ in accordance with a nonclassical logic, it seems indefensible on *nonempirical* grounds. The proponent must make his case for such a revision either in classical logical terms – reasoning classically both in showing the supposed ‘incongruence’ of classical logic and certain quantum phenomena and in defining the quantum logical connectives and rules of inference – or in terms of the new logic. If the former, he hasn’t really repudiated the normative status of classical reasoning. If the latter, he fails to persuade us, as pointing out that there are difficulties with our system of reasoning if one accepts another system of reasoning is necessarily underwhelming.²¹

These points notwithstanding, I grant that broadly empirical considerations could lead one to give up an a priori belief somewhere outside the core, but in a way that will not do much to advance the case of the Quinean vision that all beliefs ultimately stand before the tribunal of experience. Plantinga (1993: 112) plausibly suggests the following simple case: I believe a sophisticated mathematical claim upon a priori reflection, but I am a little out of my depth. (His example is a person’s coming to the belief that a function cannot be everywhere continuous and nowhere differentiable.) I am told by a trustworthy mathematician that there are such functions, though he does not provide the proof. In such a case, I will (or ought) abandon my belief, and for a clearly empirical reason. However, my warrant for abandoning the claim here derives from the superior a priori warrant the mathematician has for its denial. What if the mathematician is sincerely

mistaken or even pulling my leg, and there is no such proof? Of course, in that case, I may not be warranted in affirming the denial of my original belief, but I am surely warranted in abandoning the original belief. Here, we must grant there is no a priori judgment of anyone that is the basis of my rational revision of my a priori belief. My sole basis is the (reasonable but mistaken) a posteriori belief that there *is* an a priori basis in the community of inquirers for thinking my belief is mistaken. But it also seems that this is a special kind of case that cannot do the work of showing that significant a priori beliefs are subject to the possibility of direct disconfirmation by future scientific developments. It admittedly highlights the large and complicating role of shared knowledge, though, regrettably, I cannot delve into its implications for a priori justification here.²²

I want to consider briefly a final, original suggestion concerning the nature of a priori justification and its role in empirical inquiry. Michael Friedman has argued in sophisticated detail for a “relativized and dynamical” conception of the a priori insofar as it bears on scientific theorizing (2001). This view is based on a careful reading (from a perspective steeped in the logical empiricist tradition) of the revolutionary changes in physics from the Aristotelian paradigm to Newton’s mathematical physics and the latter to general relativity and quantum mechanics. Friedman argues that we must replace a two-place picture of formal apparatus plus empirical theory with a tripartite structure (35–8):

1. Formal mathematical structures (Euclidean geometry, calculus, and tensor calculus, or the general theory of manifolds)
2. Coordinating principles (Newton’s laws of motion, Einstein’s constancy and source independence of the velocity of light)
3. Empirical laws (Newton’s law of gravitation, Einstein’s field equations)

The mathematical structures define a space of purely logical possibility. But in similar fashion, Friedman argues, the coordinating principles function to define a space of ‘empirical’ or ‘real’ possibility. There simply can be no empirical testing of the empirical laws unless the coordinating principles are assumed, or treated as justified a priori. This is not for the Quine–Duhem thesis that one cannot test one branch of a theory in total independence of the other parts (as reflected in the fact that testing the influence of gravity on the trajectory of a beam of light requires one to assume the theory of optics underlying your use of telescopes). Rather, and more strongly, this is because the coordinating principles, by serving to connect the formal apparatus needed to make precise predictions to the empirical concepts and

laws, are necessary conditions on the meaningfulness of the latter. Without these principles, the empirical laws “are empirically meaningless,” having no empirically defined consequences. Contra Quine, there is an asymmetrical epistemological dependency of empirical laws on the coordinating principles. (Friedman 2001: ch. 2)

At the same time, the status of these coordinating principles differs from what tradition assigns to propositions justified a priori. While they cannot “become empirically false” by disconfirmation in a crucial test or experiment with respect to their theoretical alternatives, at a later stage of empirical progress involving a changed overall structure they can be so tested and indeed in some cases disconfirmed. Hence, they require a relativized and dynamical conception of the a priori (86–7).

Now, this is all as may be as an account of the structure of mathematical theories and how their *empirical* parts (coordinating principles and dynamical laws) are subject to (dis)confirmation. But for the purpose of evaluating the place of a priori beliefs *in the ordinary sense* in the practice of forming and assessing empirical theories, the real action is at the top level of mathematical structures. Does Friedman believe that formal mathematical and logical beliefs, too, have only a relativized and (indirectly) empirical justification? It is hard to say with confidence. He makes the usual nod to the demise of Euclidean geometry as an account of the necessary structure of space. He emphasizes that the calculus and tensor calculus were “controversial,” not “part of mainstream mathematics” when Newton and Einstein, respectively, introduced them – *just as* the associated coordinating principles were revolutionary in the then-current physics (39–40). And he looks favorably on Neumann’s quantum logic proposal, whatever its ultimate merits, as “[indicating] a way in which the idea of a relativized and dynamical a priori can even extend to fundamental principles of logic” (123). All of these contentions suggest that Friedman *is* thinking of at least much of mathematics as but relatively a priori. On the other hand, he speaks at one point of the theorems and principles of the abstract theory of Riemann manifolds as being “justified purely mathematically” (80). And further complicating matters is the obscure talk in one place of the field equations of relativity theory being “logically possible *as soon as* we have Riemann manifolds within pure mathematics” (84). (Were they not always possible, with the mathematical innovation simply making it possible for us to *appreciate* their possibility?)

So, it is not clear how radical a thesis lies behind Friedman’s picture of a “relativized and dynamic” a priori. In the light of our earlier discussion, however, what is clear is that the remarks just noted that might indicate a thoroughgoing commitment to the relativity of the a priori make for an

uncompelling case. Nothing about Friedman's description of theory development and revision in basic physics supports the claim that even core principles of logic and mathematics are (properly) taken as given only provisionally (within the context of Kuhn's 'normal' as opposed to 'revolutionary' science). Nor does he give reason to suppose that reflection on data in the light of our best going theories of the empirical world might directly induce, or rationally require, change even in our defeasible a priori beliefs (outside the core). Friedman's picture of a relativized and dynamical a priori is at best a correct (albeit misleadingly termed) account of the nature of theoretical confirmation and revision in physical science proper.

There is a good deal more to be said about the particular examples noted above and, more generally, about the relationship of a priori and empirical justification. My aim here is more modest than working out a full-blown picture. In the last chapter, I argued that there is no way to make sense of empirical knowledge of the world that does not rest in part on some metaphysical necessities. My aim here is merely to point out that there is an attractive, broad picture of warranted a priori acceptance of such necessities that differs from the traditional 'rational insight' account. This alternative picture is fallibilist and does not invoke quasi-perceptual rational capacities.

Objectual modal judgments

When we turn from general to specific modal judgments (concerning objects and their kinds), it is natural to begin with Kripkean considerations. Our prescientific ignorance of the underlying nature of things easily leads us to overplay the scope of possibilities – for example, supposing that there might have been water in the absence of hydrogen. (These are instances of Yablo's first condition for faulty conceivings.) Error in this regard is compounded by the subtlety of the distinction between epistemic and metaphysical possibility.²³

These loci of error in commonsense judgment are well recognized and easily resolved. More difficult to sort out are the consequences of the constantly changing, multilayered approach of modern science for ontology in general, and so modality as one of its aspects. Obviously, changes in the *categories* of our best empirical theories requires corresponding adjustment in modal claims that quantify over those categories. But crucial problems remain even were our theories at various levels to be taken as more or less finished products. First, do we deem only certain scientific kinds to be fully objective, in the sense that any description of the world that fails to quantify over such kinds is necessarily incomplete, failing to recognize certain of Nature's 'basic joints'? If so, by which criteria? Second, how does one

distinguish essential features of such natural kinds from those inessential features, if any, which happen to be had by all actual instances of the kind?

As I will now show, our answers to these questions will flow, in part, from our responses to the traditional *metaphysical* problem of universals and the modern problem of reductionism. These problems, in turn, are connected to the analysis of causation. I believe it is most efficient to begin with this last matter, and then work our way back. After our foray through these basic metaphysical topics, we will draw some lessons for modal epistemology.

Causation

Earlier, I contended that induction would lack a rational basis apart from a tacit acceptance of causal and formal necessities as primitive structural features of the world. Those traditional philosophers who agreed with this assessment have cashed out causal necessity in terms of an object's causal powers, which are closely linked to the object's intrinsic properties. Since Hume, however, it has been common to repudiate this account on the grounds that the posited powers are unobservable and mysterious. More recently, David Armstrong has added to the chorus of criticism by charging that in cases where, *ex hypothesi*, causal powers go unrealized, the dispositions appear to be akin to the absurdity of a one-sided relation: the potential cause tends towards a circumstance that does not exist (1997: 79).²⁴

I believe these criticisms of the causal powers account are unfounded. However, Armstrong (1983) and Tooley (1987) suggest that there is a way to have primitive causality without objectionable dispositions. As their approach has gained some currency, it is worth showing that it cannot deliver the goods before turning to defend the traditional account.

Armstrong and Tooley posit a primitive relation of causality, or necessitation, that *contingently* links certain universals. They develop this basic idea in different ways, in part reflecting their differences on the immanence or transcendence of universals. I will confine my attention to Armstrong's account, as I believe his understanding of universals as immanent is the more defensible. (I return to this matter below.) Start by thinking of the world as largely constituted by an enormous array of spatiotemporally related, particular matters of fact, which we may term the 'first-order states of affairs'. We will have particle *a*'s having *spin* \emptyset and *charge* -1 , *b*'s having $1/2 \emptyset$ and *charge* $+1$, and so on. Suppose we observe a number of occasions in which the instantiation of a structured universal *F* (a big conjunctive fact about some entity or system of entities) is accompanied by the instantiation of *G*, and through the usual scientific methods this pattern is taken to reflect

a direct causal connection. Here, says Armstrong, we should posit an irreducible *second-order* state of affairs, $N(F, G)$, where N is a nomic relation. Doing so, he contends, will enable us to *explain* why the F - G pattern is observed, why ever so many particular instances of F are accompanied by instances of G . In all these instances, there are literally common elements – the universals F , G , and the linking relation, N . Because of this single, second-order relation, no instance of F *can* occur without G (and so none *does* so). The second-order relation provides structure to the world, permitting some first-order state of affairs sequences and precluding others.

Armstrong originally conceived the second-order state of affairs, $N(F, G)$, as ontologically distinct from particular instances of it: a's being F at time t_1 *brings about* b's being G at time t_2 , which fact is explained by the logically prior second-order state of affairs. As several have pointed out, however, it is not transparent how a second-order relation among universals such as $N(F, G)$ constrains the character of particular first-order sequences (van Fraassen 1988: ch. 5; Lewis 1999b: 40). Recognizing the difficulty, Armstrong (1997) now suggests that the second-order necessitation relation among universals is in no way distinct from its instances: causation is a relation among *types* of states of affairs. (So when I experience the causal force, say, that is exerted on my toe by a heavy object, what I am experiencing is nothing particular, but rather causation in general, or nomicity.) This move succeeds in making transparent the relevance of the second-order relation to first-order sequences, but at a steep cost, as the explanatory advantage over the brute conjecture of the Humean regularity theorist has vanished. For each occurrence of G is ontologically and so explanatorily prior to the co-occurring $N(F, G)$ fact. (The existence of a second-order N -relation between F and G as part of the fabric of the world, multiply instanced, implies but is not implied by the constant conjunction of first-order F and G instances.) So the posit of the N relation is gratuitous, as it can only be put into the world consequent upon the regularity – unless one makes a stronger claim to the effect that F -ness by its very nature is disposed to bring about G -ness (bringing the second-order constraint into first-order facts), in which case we are back to the primitive dispositional thesis Armstrong has sought to avoid.²⁵

What our discussion shows, I think, is that the Armstrong–Tooley picture is simply second-order Humeanism, adding a bit more structure to the inert Humean picture of the world at no explanatory gain whatsoever. Therefore, let us return to the causal powers account. It ascribes a fundamental duality to properties: they are qualities that are disposed to produce other qualities in particular circumstances.²⁶ Since Hume, it has been common to excise the dispositional from the world, leaving only pure, qualitative 'suchness',

on the grounds that the posited *powers* are unobservable and mysterious, something we cannot receive a direct ‘impression’ of. But the claim that the dispositional aspect of things makes no impression on us is question-begging, and the assumption that we can make do without them is dubious. I observe events without, such as rushing water, and am aware of states within, such as my believing that Hume was badly mistaken. Of what am I aware – sequences of purely qualitative and categorical states only, or states that also dispose towards activity? We may take this question two ways: (1) What is the nature of my perceptual and introspective processes and of their objects? (2) How do I conceptualize these processes? As for the first question, if the world is as the anti-Humean claims, then to look at rushing water (or flying birds, or . . .) *is* to perceive dispositionality: I perceive intrinsic characteristics, and these are, on this view, partly dispositional in character. As for the second question of how I *conceive* what is going on in perception, Hume’s anti-dispositionalist account of the matter is unconvincing. As William James long ago observed, the Humean picture of a patterned sequence of mental ‘snapshots’ of independent external states is not how anyone without a tendentious philosophical axe to grind would characterize our perceptual life: “Every examiner of the sensible life in concreto must see that relations of every sort, of time, space, difference, likeness, change, rate, cause, or what not, are just as integral members of the sensational flux as terms are . . .” (1987: 757).

Might we nonetheless, on philosophical grounds, shake ourselves of our actual, dispositionalizing conception of what we experience? Not without emptying our scientific theories of any content concerning the intrinsic character of fundamental kinds, all of which are individuated in functional terms.²⁷ (The concept of unit negative charge is the concept of a feature that interacts with other fundamental qualities in specified regular ways – for instance, and put intuitively, to repel instances of positive charge.)

Basic properties

I’ve been exploring the metaphysics of causation for the broader purpose of showing how fundamental natural kinds posited by scientific theories might be objectively basic joints in the physical world and, still further, for developing a plausible recipe for distinguishing the essential from the inessential features of such kinds. We’ve seen that a satisfactory account of how we are justified in making causal inferences requires a robust realism about causal dispositions, on which properties necessarily confer tendencies to act in specific ways. Hence, our template for how we may come reasonably to form detailed beliefs about the modal properties of things has as a core

building block the thesis that modal properties are woven into the basic fabric of the world. But surely these do not include *all* ‘properties’, on an indiscriminate use of that term. My ‘property’ of being such that Ron Artest is a tenacious basketball player does no causal work. Nor even does my ‘disjunctive property’ of being 6 feet tall or being a Martian. And it is unsatisfactory to simply declare that some properties do causal work and others do not.

It is often assumed, quite implausibly, that a single category of entities (properties) serves a variety of disparate functions, from underwriting statements about the structure and dispositions of objects to being the meanings of terms. Causal realists, however, will wish to sharply separate the truth-makers needed for metaphysical and semantic purposes into the categories of property and concept, respectively.²⁸ Our concepts for things are abundant and can be as useless as you please for explanatory purposes, but properties (as I will use the term) are sparse and earn their keep: they are those aspects of things which ground objective, intrinsic similarity and difference among things and that confer basic capacities to act. Many concepts might truly apply to an object in virtue of a single one of its properties.

Armed with a metaphysics of properties, I will now outline a congruent approach to the essential properties of objects. We should divide our topic into two categories, corresponding to basic and composite individuals. I set aside the difficult question of individual essence and focus exclusively on the essential feature of object kinds.

Basic object kinds

We do not yet know what our world’s fundamental kinds of particulars are. Among the posits of present-day particle physics are electrons, quarks, messenger particles, and fields of interaction among them. It identifies these kinds in terms of limited clusters of features that are themselves characterized dispositionally. Electrons, for example, have unit mass, unit negative charge, and spin 1/2. Importantly, it does not just fall out of physical theory that these features happen to co-occur ubiquitously at the subatomic level. Rather, these features are treated as *constitutive* of the (possibly fundamental) kind, *electron*. It is quite plausible to suppose more generally that all universally shared features of the fundamental object kinds (whatever these might turn out to be) are essential to that kind. Granted, we could perhaps imagine epistemic circumstances in which this supposition would not be well founded. (Imagine that the roles of distinct fundamental particle kinds overlapped a great deal, and that for each kind, there were some properties had by nearly, but not quite, all of its instances.) But given our actual circumstances, in which physical theory take certain property-clusters of

fundamental particle kinds as basic units of explanation, and does so with great success across a wide spectrum of experimental circumstance, the hypothesis that universally shared features of the basic object kinds are alike essential is a reasonable, if defeasible, one.

Composite object kinds

I said earlier that our theory of properties should be sparse, countenancing only those candidate kinds that ‘earn their keep’ by doing nonredundant causal work. Just how many properties are there? And – what is a related question – how many kinds of object possessing wholly objective persistence conditions are there? Answer to questions concerning essence, alternative possible histories, and the like depend on answers to these more basic questions – and these, in turn, depend on how the properties of composite ‘objects’ relate to those of their fundamental constituents. So this last matter is where we must begin.

Aristotle supposed, quite commonsensically, that *horse* was a basic natural kind. Horses appear to share numerous salient *and distinctive* characteristics that explain all manner of phenomena. But modern biology clouds this picture considerably. We know that many macroscopic features of living organisms, including horses, are wholly constituted by the enormously complex activity of their parts. It is an open possibility – hardly established, but far more plausible since the mid-twentieth century than it ever was previously – that *every* observable feature of a given horse at a given time has such a microstructural constitution. If this were so, and we were to embrace a parsimonious ontology of universals, we should conclude that, appearances notwithstanding, *horse* is not a true natural kind. There is a concept of *horsehood*, and a useful one at that, but an ontological inventory of the natural world that fails to employ it is not thereby incomplete. It will be ignoring a certain derivative macroscopic pattern running through parts of nature, but not, again, any of its basic joints.²⁹

The reductionist program in biology, then – not to mention chemistry and neuroscience – complicates the metaphysician’s task of identifying true natural kinds. In my judgment, a plausible criterion will involve the notion of ontological *emergence*. An ontologically emergent property is *structurally simple*, (therefore) *primitively efficacious*, and a *causal consequence* of a composite’s having the requisite type of intrinsic and functional complexity. A property is ‘nonstructural’ if and only if its instantiation does not even partly consist in the instantiation of a plurality of more basic properties by the object or its parts. It is ‘primitively efficacious’ in that it has a causal tendency that is neither constituted by nor supervenient upon the fundamental causal tendencies of the object’s microphysical features, includ-

ing the relations among its parts. Instead, the emergent state appears as a causal consequence of the object's having a certain complex configuration – one that could not be anticipated through analysis of microphysical dispositions in contexts lacking the required complexity.³⁰

If there are any properties that are emergent in this sense, they clearly will make a fundamental difference to how the world works, even from the vantage point of microphysics. Any attempt to give a general account of the world's evolution in purely microphysical terms – an account holding at all regions of space-time, no matter its occupants – would falter in regions where emergent features are instanced. Even a parsimonious ontology would have to quantify over them, and so over the *objects* that have them. Such objects would thus fall under genuine natural kinds, with the holistic features conferring the 'natural unity' that is the hallmark of genuine composites, as against mere aggregates. Furthermore, there is a plausible general criterion for identifying the essential properties of emergent objects: whatever structural features are causally necessary to sustain the persistence of the emergent feature. At any rate, this is a start. (We can imagine lots of complications on this simple emergentist model.)³¹

I will not address the question of whether it is plausible to suppose there are any such emergent properties, although I have done so elsewhere.³² I note only that it is plausible that questions about the essential properties of composites hang on this prior question. It is important to see that some sort of conventionalism about many or all mid-sized objects is compatible with a robust commitment to nonconventional essentialism about the fundamental entities, whatever they may be. Indeed, conventionalism is extremely plausible for many objects in commonsense ontology, beginning with artifacts, as it is very likely that none of the features of such 'objects' are ontologically emergent.

This completes my very sketchy and incomplete treatment of objectual modal judgments. I hope to have given the reader some sense of the direction I would take a fuller treatment, enough to make claims by modal deflationists about the hopelessness of the realist project appear overblown.

A modal skeptic briefly considered

Peter van Inwagen is a realist about modality who nonetheless is moderately skeptical about claims to modal knowledge (1998: 67–84). He allows that we can know modal truths that are relevant to everyday life ("this table could have been over there") or to well-confirmed theories in science, and even that we can justifiably believe some that are strictly philosophical. But, he thinks, modal claims that are far removed from everyday life ("there

could be a three-inch-thick sheet of iron that is transparent to light”) or which are controversial and have substantial philosophical implications (“it is possible that I exist and nothing immaterial exists”) are simply beyond our ken (69–70). He draws an analogy to judgments of distance made by the naked eye, which are quite reliable within certain limits, but which are wildly unreliable when, say, judging the distance of the sun.

Furthermore, while van Inwagen thinks we do have some modal knowledge, he thinks it is mysterious how we come by much of it. The nonmysterious variety concerns those modal propositions whose truth or falsity can be established by logic, mathematical reasoning, and reflection on the meaning of words. The mysterious variety concerns (or depends upon) knowledge of possibility for claims which are not known to be true or knowledge of necessity for claims that are known to be true.

We know that a table in front of us could have been two feet to our left and that there could not be liquid wine bottles – somehow, but in a way that we do not, at bottom, really understand. That we do not understand how we know these things is seen through reflection on a picture of how basic knowledge of this sort is acquired, a picture which he claims is at least on the right track: Yablo’s method of justifying the claim that p is possible by imagining a scenario which one takes to verify p . Even if this is roughly right, says van Inwagen, it doesn’t really enable us to see *how* we know the possibility claim. For it requires us to know that the scenario we imagine is itself a possible one (as I noted earlier in criticizing Yablo), and we are not given a way to see how we know that such a scenario is possible, which seems required for us to know *how* we know what we know in this sphere (75–6). Nonetheless, van Inwagen supposes that we do (somehow) know some such possibility claims, and Yablo’s method, or something like it, is part of the story. *Skepticism* about knowledge of possibility comes in when we consider the limits on using Yablo’s method. We cannot imagine in sufficient detail potentially verifying scenarios involving ‘remote’ or overly complicated matters (e.g., the transparent iron case), or controversial philosophical possibility claims, which implicate contentious and not fully imaginable factors about the large-scale character of any world in which these alleged possibilities would obtain. As for controversial (nondemonstrable) *necessities*, van Inwagen claims that it is “less clear whether we know of any proposition that it is a necessary truth if it cannot be shown to be true either by reflection on logic and the meaning of words or by mathematical reasoning” (74).

In my judgment, van Inwagen should have pressed further in querying our ‘basic knowledge’ of possibilities, as it would have revealed its partly

theoretical basis. Consider our confident judgment that the table before us might have been two feet to our left. Underlying such a judgment, surely, is the knowledge that this table and various features of its environment are intrinsically like other objects in relevant respects and so (we judge) like them with respect to certain modal *constraints* they induce. Judging there to be no relevant constraints on the possibility of *this* table's occupying a certain location when similar objects have occupied similar locations, we endorse the possibility. In other words, judgments of the kinds and scope of necessity enter into our judgments of possibility, which fact appears to run counter to van Inwagen's suspicion that there is no "very close connection between" knowledge of possibility and knowledge of necessity (74).

In some ways, van Inwagen's most careful statement of the skeptical side of his position on modal knowledge (appearing in a final note) is quite sensible and congenial to my own view. For it merely questions whether we can know truths involving remote and complicated unactualized *possibilities*, on the assumption that Yablo-esque imaginability is the only method available for deciding the matter in question. Even on the present view, which does not share that assumption, it is highly plausible that we are simply not in a position to confidently determine possibilities for kinds of entities, structures, or situations whose properties are not clearly grasped. The (to my mind unwarranted) skepticism in van Inwagen's position really comes in when this statement is conjoined with the strong inclination to deny knowledge of *necessities* beyond demonstrable ones, which inclination is expressed elsewhere in the text (74). As I have argued above, we can have warranted claims about necessity where these play clear explanatory roles in our best overall conception of the world. (I leave aside what we might say about which of our true modal beliefs might be *known* as opposed to possessing a more limited degree of warrant. My interest, again, is not to provide a thorough epistemology of modal beliefs. It is, instead, to sketch a plausible picture of how we reasonably come by modal beliefs in best-case scenarios and to explore important implications of this picture in the chapters that follow.) To be sure, confidence in claims of necessity must be tempered by recognition of the fallibility of our theoretical judgments and the likely need to revise aspects of our current picture of the natural kinds that structure the world's dynamics and are crucial to determining its ontology. And we should always allow for the possibility of our coming to discern motivation for hitherto unrecognized constraints on possibility – something to which I will return in the final section of this chapter.

A causal role for theoretical necessities?

Earlier, I suggested that the justification of *theoretical* modal beliefs increases through a process of complex inference and reflective equilibrium, assuming that our foundational modal convictions are reasonably on track in the first place. But a familiar reflection suggests this can't be the whole story. Surely it matters what the causal origin of these fundamental modal convictions is. Our acceptance of all such claims had better not be 'accidental' from the point of view of evolutionary history, being explained in a manner that is causally independent of the truthmakers of such claims. (Contrast our disposition to accept the deliverances of memory, say, which disposition is presumably explained in part by the fact that memory is usually reliable in the main.) Here we return to the problem raised back in Chapter 1, in the section "An Epistemological Worry about Modality: Causal Contact with Modal Facts," which problem has animated many deflationary and anti-realist theorists of modality.

This problem of accidental truth appears to afflict the position David Lewis staked out in *On the Plurality of Worlds* (1986), for example. Lewis attempted to sidestep it by exploiting the difference between contingent and necessary facts. A causal condition on knowledge, he suggested, is warranted only where it is possible that our beliefs not track the facts – only, to be precise, where we can sensibly ask what we would have believed were the facts believed to have been otherwise. But necessary facts could not be otherwise, and so the idea of 'tracking' the facts cannot apply (111–12). In the modal sphere, he suggests, we are fundamentally guided by a recombination principle (as we saw with Yablo), which we apply to possibilities we already accept to extend our beliefs. We properly modalize when we engage in piecemeal revision, balancing the virtue of theoretical conservatism with the pursuit of theoretical unity. "If we are prepared to expand our existential beliefs for the sake of theoretical unity, and if thereby we come to believe the truth, then we attain knowledge" (109).

Charles Chihara argues that this reply is insufficient in a way analogous to the inadequacy of the justified true belief (JTB) analysis of knowledge (1998: 90–3). So long as our acceptance of modal claims on the basis of theoretical unity considerations is causally independent of the truthmakers of those claims, we will still have something similar to accidental truth that figures in Gettier problems for the JTB analysis. Chihara presses this worry by suggesting we have no reason to accept the truth-conduciveness of accepting modal claims for reasons of theoretical unity. In that, I think, his objection misfires, as all that need be the case for Lewis's position on modal knowledge to be plausible is for it to be *true* that the structure of modal

facts is such that using the principle tends towards truth. A better way to cast the worry is the one I give above, in terms of the causal origin of our fundamental modal convictions, including the disposition to believe that theoretical unity conduces towards truth in the modal sphere. Given Lewis's own wider views on modality (the truthmakers for modal truths are worlds causally disconnected from our own), it seems not only possible but unavoidable that our acceptance of all such claims is 'accidental' from the point of view of evolutionary history. There will be a (complex) contingent historical explanation for why we modalize as we do that is causally independent of these truthmakers.

What might one say here? Bear in mind that one need not *demonstrate* the connection of modal beliefs to their truthmakers. It is enough that one tell a just-so story on which they are in fact so connected, a story that is not wildly at odds with settled fact. I will not offer such a story, but will state a story *beginning* that I think is worth considering. It is this: An evolutionary advantage accrued to cognizer-types that readily assent to the *actual* truth of core logical and mathematical principles and that systematize the world in terms of natural kinds; some such cognizers in our ancestral history were selected in part owing to this fact; and *the truthmakers for these actual truths are none other than their modalized counterparts*.

Such are the beginnings of my story. But now what are these truthmakers for highly general modal facts? For now, I leave this as an unanswered challenge for the position I am sketching. I call attention to (without endorsing) a radical answer Robert Koons has recently proposed: in the course of developing an information-flow account of causality in the context of situation theory, he argues that we need logical and mathematical information to play a causal role (2000: ch. 15). This leads him to hold that these formal facts are literally embedded in the natural world, as part of the structure of situations that cause our judgments. If this view could be worked out within an attractive metaphysic, it would both provide a linkage between baseline modal facts and dispositions to modalize and allow for general modal facts literally to give structure to the actual world (as opposed to the one merely corresponding to the other). But can we plausibly integrate this formal structure with the contingent causal structure imposed by actual properties, such that the formal structure is still capable of being realized within possibilities with "alien properties"? I do not know. Even if we adopt this proposal, we still need an account of the epistemology of belief *revision*, such as the one I've sketched above. In any case, Koons's bold proposal may be gratuitous given the case for the unification of modal facts that I will make in the next chapter, a case that leads to a venerable but unfashionable view of their ontology.

The Spheres of Possibility

It is common for contemporary philosophers to think that there *several* distinct kinds of alethic modality, so that we have narrowly logical possibility (consistency with theorems of first-order logic with identity, perhaps); conceptual possibility (logical possibilities that are also consistent with conceptual or ‘nominal’ definitions); metaphysical possibility (also consistent with a posteriori identities of kind and individual essence); and causal possibility (also consistent with causal and structural laws of the actual world). If the broad conception sketched above is on the right track, then this common view is mistaken.³³ As an alternative, I suggest that we think in terms of concentric spheres (see Figure 1), denoting increasing theoretical constraints on possibility *tout court*, with merely formally consistent statements in the outermost sphere, and genuine possibilities in the innermost sphere.

Looking at things this way leads to the conclusion I have already urged that quick ‘recombination’ principles are misguided, but it also explains why they are tempting. What a philosopher is able to grasp in suitably constrained ‘modal thought experiments’ that the present view might deem to involve impossibilities are incomplete possibility-candidates, involving partial considerations which abstract away one or more grounds of their impossibility. There can be reliable reasoning about and with and such incomplete possibility-candidates, provided one is clear about the sphere by which such reasoning is to be constrained.³⁴

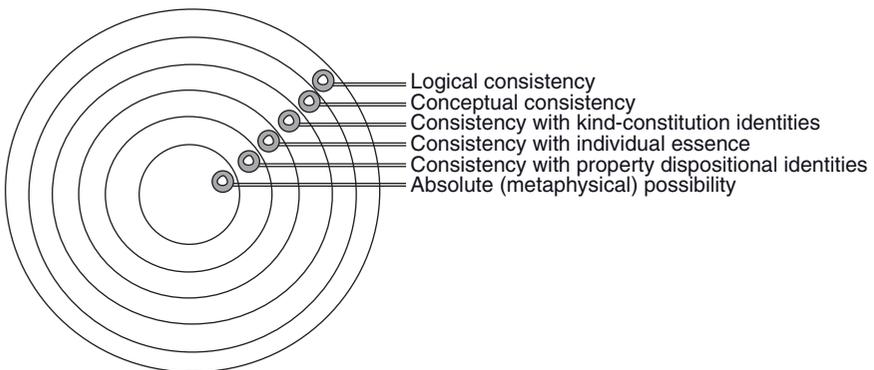


Figure 1

I've said that the innermost sphere is distinctively constrained by necessities of nature, pertaining to both kinds and individuals. I now raise the following important question: Are there extrascientific explanatory considerations that should lead us to posit yet further modal structure to the space of genuine possibility?

- 16 David Robb offered this suggestion to me in correspondence.
- 17 MacBride makes this point (1999: 471–501).
- 18 See the discussion in Lewis (1986: 184–9).
- 19 The claim that all we need to do to effect the reduction is to believe in ‘more of what we already believe in’ is nonsense. On any reasonable comparisons of commitments to theoretical entities, Lewis’s debt is unfathomably large: elves, fairies, dragons, and Greek gods are just the tip of the iceberg.
- 20 Plantinga (1987) makes a similar criticism of Lewis’s candidate modal truthmakers.
- 21 I say “it” for convenience, although there will be a plurality of sentences satisfying his account. Sider notes (2002: 295 n. 27) that this will be harmless provided they are logically equivalent, as intended.
- 22 For a careful comparison, see section V of Sider (2002).
- 23 Cf. the objection posed to Rosen: Why suppose that there is some *one* fiction whose content is the object of our disagreement?
- 24 A similarly conventionalist line run on behalf of a different view of the modality of identity statements is found in Gibbard’s “Contingent Identity” (1975: 187–222).
- 25 Perhaps there is room for the barest supposition of a real causal connection linking quantities of the world-stuff, as a referee suggests. However, this will still be inadequate to underwrite the laws we use that quantify over conventional kinds.
- 26 And mere bits of stuff cannot do the job, either. Convention-imposing thinking and speaking could not themselves be conventional. Such activities, therefore, imply the existence of thinkers with objective persistence condition.
- 27 It appears to be committed to rejecting S5, which plausibly accords with the structure of ordinary modal thinking.
- 28 Sympathy for the broader view that modality is to be understood in noncognitivist fashion is expressed by Wright (1980: ch. 23; 1989). See also Craig (1985).
- 29 Field defends a nonfactualist, ‘evaluativist’ view of the reasonableness of our adopting the a priori rules of inference that we do (2001: ch. 13). In his review of Field’s book, Koons (2003: 119–26) effectively makes the point that Field’s attempt to rebut the charge that his position leads to rank relativism concerning epistemic standards is plausible only to the extent that he presupposes a notion of objective propensity. Koons’s observation illustrates the bulge-in-the-carpet problem that is inherent in deflationary strategies concerning modality and a priori knowledge.

2. Modal Knowledge

- 1 Van Inwagen makes this point in assessing Yablo’s account (1998: 67–84).
- 2 For a more careful statement of this idea, see Lewis (1986: 87–9).
- 3 Yablo recognizes these worries without attempting to address them. He suggests that such questions are apt to lead us to reconsider “either the nature of

the target facts or the nature of our access to them,” citing in a note recent work with a strongly conceptualist flavor, including Sidelle (1989). I argued in Chapter 1 that such strategies are quite unpromising, and it is hard to see what other sort of strategy could be followed to bolster our confidence in the imagination-based account Yablo sets before us.

- 4 See also the symposium on his book (2002b: 636–79), all of which centers on Peacocke’s account of modality.
- 5 I discuss this conception below.
- 6 Peacocke at one point suggests that the Principle-based conception is not undermined even if it is true that the very notion of a concept involves the notion of metaphysical necessity. (Cf. my discussion in Chapter 1, in the section “Modal Reductionism and Deflationism.”) He points to the object- vs. meta-level distinction, saying that a theorist’s own modal understanding need not infect his characterization of another thinker’s mastery of modal concepts (159–60). This is true enough, but if the claim about the identity conditions of concepts stands, it shows that reductionist ambitions for Peacocke’s project cannot be achieved.
- 7 Fine urges us to replace a modal characterization of essence with one that answer’s to Aristotle’s notion of ‘real definition’ (1994: 1–16). But this, if accepted, does not lend aid to a thoroughgoing reductionist program, partly for just the reason indicated in the text: our grasping of essences, so construed, is rooted in counterfactual reasoning that presupposes objective modal facts.
- 8 Compare the criticism of Sider’s and Rosen’s metaphysical views in Chapter 1.
- 9 Another unsatisfying aspect of Peacocke’s project, typical of deflationary strategies, is its inability to provide a place for possibilities involving fundamentally ‘alien’ properties (properties not instantiated in the actual world). All nonactual possibilities must, on his picture, “be constructed from the materials of the actual world” (153). In his symposium response (2002: 676), Peacocke acknowledges the need to generalize the account to nonactual kinds, while not offering any suggestions concerning how this might be done. It seems that the resources of his theory can at best underwrite kinds consisting of novel combinations of actual world properties.
- 10 This is revealed *inter alia* in our persistent inclination to raise modal questions about the surrogate items featured in deflationary and thinker-dependent theories of modality. And note that the thesis that items in the world have a modal character must be distinguished from a mere functional-role view of our *concepts* that denies or is neutral on the metaphysical thesis. The latter view holds that it is essential to our *classifying* something as an electron that it exhibit a characteristic pattern of behavior associated with the terms “negative charge,” “mass,” and “spin,” but allows that an actual electron might have behaved entirely differently, in which case it would not be an electron. Ellis and Lierse (1994: 27–45) reject this view for much the same reason as I do. Mumford (1998: 232–8) criticizes their argument by defending the plausibility of combining an essentialism about dispositions with the view that objects belong to the

dispositionally determined kinds that they do only contingently. I briefly discuss the question of kind essentialism in the section “Objectual modal judgments” later in this chapter.

- 11 Here I agree with Bealer, who argues that any view that denies that a priori ‘intuitions’ count as ‘basic evidence’ is incoherent (1999: 29–55). We will consider later the question of the connection between such modal beliefs and the facts known.

I will not undertake an analysis of Bealer’s own intricate theory of a priori justification set out in a series of papers that includes the one just cited. Instead, I simply note a basic disagreement underlying our respective pictures. He proposes that the reason that our basic intuitions reliably modally track the modal facts is that *determinate concept possession* (an idealized state to which we can sometimes approximate) entails that there is such a tie, that the subject would have a given intuition if and only if it is true. I see little to recommend such a strong claim, unless, perhaps, one is already committed to the transparency of modal truths to ideal inquirers. I am animated instead by the belief that theoretical explanations properly often invoke or tacitly assume necessities that are not transparent even given an idealized grasp of the relevant concepts. Nonetheless, we agree that (as I will go on to suggest in the text) reflective equilibrium governs the process of revising a priori beliefs even while it remains autonomous with respect to empirical inquiry.

- 12 See the brief discussion in Boyer (1991: 545–6).
- 13 See the summary of some of this process in Kitcher (1984: ch. 10). However, Kitcher himself draws the wrong lesson from this developmental process, rejecting the a priority of mathematical knowledge. His cardinal error lies in wedding the notion of a priority to infallibility. But for his response to this and other criticisms of his thesis, see Kitcher (2000: 65–91).
- 14 There is now a clear trend away from the infallibilist thesis. See, e.g., Plantinga (1993: 110–13), Goldman (1999: 1–28 [5]), Bonjour (1998: 111–13), and Bealer (1999: 29–55).
- 15 Whether there is a distinctive phenomenology to our conscious believings of necessities is a difficult question. Bealer (1999: 30) and Bonjour (1998: 102) speak of ‘intuitions’ (or ‘intellectual seemings’) and ‘rational insight’, respectively, and maintain that these are instances of a distinctive propositional attitude that ordinarily occasion belief, much as sensory experiences are distinct from but generally occasion empirical beliefs. Plantinga (1993: 106) agrees with the view in the text that they are a certain kind of belief, while arguing that they have a distinctive phenomenal character. We may suppose in either case that such a distinctive quality of intellectual seeming is a necessary condition on the beliefs having a priori warrant. (I note that Goldman [1999: 9] claims that varieties of rational thinking generally often have phenomenal aspects but doubts that there is a single distinctive form.)
- 16 Plantinga is a contemporary philosopher who is admirably clear on this point (1993: 110ff.).

- 17 Hilary Putnam's wholesale repudiation of a priori knowledge in response to the development and successful application of non-Euclidean geometry probably reflects both of these reactions (1965).
- 18 See Sklar (1976: 14–15) and the appendix to Bonjour (1998).
- 19 The supervaluationist alternative of course shows that rejection of this law is not a *required* response to vagueness.
- 20 And so also our affirmations and denials that certain propositions are necessary truths.
- 21 Here I am indebted to a similar argument made by Railton, though he does not draw as firm a conclusion from it (2000: 174). A referee suggests that we take the reasoning of the critic as a kind of *reductio* of classical logic, undermining it from within. But this fails to appreciate the point that our ability to *grasp* the rules of the nonclassical logic that is supposedly shown to be preferable requires us to reason classically and is thus dependent on the normative status of such reasoning.
- 22 Another kind of case where certain beliefs may receive *both* a priori and a posteriori justification is this: There are empirical results that are deeply puzzling, conceptually, in the context of well-confirmed empirical theories. On reflection, we are able to accommodate these results within an improved physical theory without abandoning any formal or philosophical commitments that had vaguely seemed 'threatened' by those results. Here, the successful accommodation seems to properly augment our confidence that we haven't made some sort of mistake within the a priori chain of thought leading to those commitments. Depending on how strongly we characterize a priori justification, a simpler instance of the same phenomenon may occur when my confidence in the belief in the necessity of a theorem increases upon my getting the same result after checking a complicated proof, or in having my proof confirmed by an expert referee.
- 23 Epistemic possibility might be thought of as consistency with what one confidently believes (or perhaps believed consistency with what one believes, for a still more agent-relative notion).
- 24 Armstrong is equating an unmanifested disposition with a one-sided *relation* attached to the dispositional circumstance on one end and nothing on the other. The proper reply is that this assimilation of dispositions to relations reflects a failure to keep firmly in mind the distinction between the dispositional and qualitative aspect of things. Dispositions are intrinsic and 'complete' entities quite apart from their manifestations.
- 25 To make matters worse, Armstrong has to give back the (spurious) advantage of the type-based account when dealing with probabilistic causality (1997: 237ff.). In such cases, he suggests, "the first state-of-affairs type will give a certain objective probability . . . that an instance of the second state-of-affairs type will be caused to exist in suitable relation to the first state of affairs" (237).

So the type-level fact is only a propensity – a propensity for a first-level instance of causation, which in some cases will not occur (such a 'misfiring' has a nonzero probability of occurring).

Tooley's response to the same problem of showing the relevance of second-order causal laws to first-order sequences of events is different from Armstrong's. He posits unusual features in the mereology of transcendent universals. If it is a law of nature that all things having property *P* have property *Q*, then we might suppose that *P* "exists only as part of the conjunctive universal, *P* and *Q*" (1987: 125). It would then follow that any time *P* is instantiated, *Q* is as well. But this response puts pressure on the official view that the second-order relations by means of which we specify the causal relation are only contingently associated with it. What more than the obtaining of *N* (*P*, *Q*) can be meant by Tooley's claim, if not that it is part of *P*'s nature to be bound up in the second-order fact? (And, like Armstrong, Tooley cannot extend this suggestion to the case of probabilistic causation. See Tooley 1987: 148ff.)

Both Tooley and Armstrong are driven into the original difficulty precisely because of their insistence, in opposition to the causal powers account, that the nomic structure among universals is merely contingent.

- 26 It would take us too far afield to consider an alternative 'pure dispositionalist' view, as defended, for example, in Shoemaker (1980: 109–35). Either view would serve my argumentative purposes in this chapter.
- 27 Ellis emphasizes this point in arguing for the dispositional conception of properties (2000: 329–51; 2001).
- 28 This point has been made recently by several authors, including Armstrong (1978, 1983) and Swoyer (1999: 100–31).
- 29 There are any number of other macroscopic patterns that human beings are not wired to be sensitive to. But it would be absurd to claim that one has fully characterized the world only when one has taken notice of every such supervening pattern. I discuss this matter more fully in "What Kim Should Have Said," currently unpublished.
- 30 This at least will characterize the *onset* of emergent states within a system. Since the initial emergent states will themselves help to determine similar subsequent states – possibly resulting in a complex, stratified range of such states in some systems – the microphysics alone will not determine these later states. Furthermore, emergent states will work in tandem with the underling micro-states to determine later micro-states, manifesting a sort of 'downward causation'. I discuss the concept of ontological emergence at length in O'Connor (2000a: 105–17) and O'Connor and Wong (2005: 659–79).
- 31 We shall still have to contemplate cases in which there are layers of emergent features, along a spectrum from those in which there is a single baseline feature accompanied by ephemeral accoutrements which that baseline feature partly sustains to those in which secondary and tertiary emergents are more robust and radically reconfigure the causal capacities of the system. See O'Connor and Jacobs (2003: 540–55).
- 32 I regard as quite plausible the claim that mental phenomena are emergent in this sense. See O'Connor (2000b: ch. 6).
- 33 Happily, a few recent thinkers have also begun to question this orthodoxy. (See, e.g., Shoemaker 1998: 59–77.) The reader who takes the standard view

of distinct kinds of possibility to be just obvious is urged to consider Peter van Inwagen's (1998) amusing and instructive analogy to someone who posits a distinctively 'cartographic' variety of uninhabitability: uninhabitability that is indicated by an imagined infallible Standard Atlas, which marks certain islands as uninhabitable, none as inhabitable, and does not claim to be complete in this matter (71).

- 34 Chalmers criticizes the acceptance of nontransparent necessities (which he dubs "strong metaphysical necessities") on the grounds that it leads to *modal dualism*, with the two wholly distinct modal primitives of logical and metaphysical possibility (1999: 473–96; 2002: 145–200.) The present view posits no such dualism, as the notion of 'narrow' logical possibility is here conceived as an abstraction resulting from restricting one's focus to the essential properties of the logical constants. Furthermore, the arguments of Chapters 1–2 imply the untenability of his own monist position, 'modal rationalism', which posits a notion of logical possibility defined in terms of a notion of ideal conceivability. (See also Jackson 1998.) Rather than seek to ground metaphysical possibility in an allegedly independent notion of logical possibility, we have seen reason to reverse this procedure.

3. Ultimate Explanation and Necessary Being: The Existence Stage of the Cosmological Argument

- 1 J. J. C. Smart, who evinces a great deal of sympathy for the Existence Question while regarding it as inherently unanswerable, is quite explicit that his rejection of the cosmological argument from contingency is based precisely on his failure to grasp any absolute modality other than that exhausted by formal logic. See his and John Haldane's contribution to the debate in Smart and Haldane (1996: esp. 180–3).
- 2 And whatever the ultimate merits of such objections in relation to Aquinas and Scotus, Leibniz's metaphysics is of course clearly beset with this problem, though for a very different reason, connected to his understanding of the conditions for adequate explanation.
- 3 Of course, Kant is correct that anyone who accepts the existence of a necessary being that has other traditional perfections must also accept that there is a sound version of the ontological argument – valid in form and with all true premises. But that is not the same as accepting that the ontological argument has independent force as an argument for the existence of a necessary being. (Cf. the atheist's position with respect to the following argument: [1] Everything is red or God does not exist. [2] Not everything is red. Therefore: God does not exist.) Cf. Forgie (2003: 364–70).
- 4 Walker alleges that the simple appeal to necessary existence is not even formally satisfactory in providing an ultimate stopping point of explanation (1997: 109–23). He writes, "Sugar is necessarily soluble, in the sense that to be soluble is part of its essence, but it is still reasonable to ask why it is; in that case we