Naturalism and the A Priori

Penelope Maddy

The naturalism I aim to practise is a descendant of Quine's.[1] My goal here is to place this naturalism in what I hope to be an illuminating historical context, to trace the status of the a priori through its various twists and turns, and eventually to draw some tentative conclusions about the naturalistic status of the a priori. To do this, I first return to Kant. While it's surely no surprise that an examination of the a priori should start from Kant, perhaps his relevance to naturalism is less obvious. Let me begin, then, with an introductory word on that connection.

Though my naturalism differs from Quine's in a couple of significant ways, these disagreements won't matter until the very end, so we can begin with Quine's leading idea: the 'abandonment of the goal of a first philosophy' (Quine 1975: 72). The interconnections between Quine and Carnap will take centre stage later, but for now we need only note that the bare rejection of first philosophy can be seen as evolving out of Carnap's classification of many traditional metaphysical claims as 'pseudo-statement[s] without cognitive content' (Carnap 1950: 250). Carnap's idea is that legitimate scientific questions, 'theoretical questions', are asked within the linguistic framework of scientific language, with its associated principles of evidence; in contrast, metaphysical pseudo-questions are posed outside of all linguistic frameworks; perhaps as preamble to the adoption of a linguistic framework, as such, they are asked without the backing of associated evidential rules that would make them answerable, and indeed, that would give them sense.

Now Kant also had a keen nose for the pseudo-question:

To know what questions may reasonably be asked is already a great and necessary proof of sagacity and insight. For if a question is absurd in itself and calls for an answer where

[1] For details, see my (1997).

none is required, it not only brings shame on the propounder of the question, but may betray an incautious listener into absurd answers, thus presenting, as the ancients said, the ludicrous spectacle of one man milking a he-goat and the other holding a sieve underneath. (A58/B82–3)

Kant is particularly concerned to warn against applying concepts outside their proper range:

[I]t is very tempting to use these pure modes of knowledge of the understanding and these principles by themselves, and even beyond the limits of experience, which alone can yield the matter . . . to which those pure concepts of understanding can be applied. (A63/B87–8)

Or, more poetically

The light dove, cleaving the air in her free flight, and feeling its resistanc might imagine that its flight would be still easier in empty space. (A5/B8)

Kant's goal is to warn against this 'transcendental illusion', to 'expose the false, illusory character of these groundless pretensions' (A63–4/B88). Thus I think it is no stretch to suppose that an understanding of Kant's critical philosophy and its fortunes might shed light on the sources of naturalism.

Admittedly, Kant's critical philosophy is not the sort of thing at which it's easy to take a brief look. If I've learned anything from my foray into the secondary literature, it is that the game of isolating the correct interpretation of Kant is not one a novice like myself can profitably play. What I propose to do instead is to sketch two well-known interpretations of transcendental idealism that lie at opposite extremes. My hope is that this contrast will provide some useful tools for our advance into the twentieth century.

For all the customary focus on Hume in discussions of Kant, our story begins with Locke and Berkeley. 'The celebrated Locke', as Kant calls him (A86/B119), held a representative theory of perception: we are directly aware of our sensory experiences; some features of those sensory experiences match actual properties of external things—these are the primary qualities, like shape—while other features of sensory experiences are merely the result of the actions of those external things on our sensory apparatus—these are the secondary qualities, like colour. This sort of theory immediately gives rise to sceptical worries, as Locke himself realized:

It is evident the mind knows not things immediately, but only by the intervention of the ideas it has of them. Our knowledge, therefore, is real only so far as there is a *conformity* between our ideas and the reality of things. . . . How shall the mind, when it perceives nothing but its own ideas, know that they agree with things themselves? (Locke 1690: IV. iv. 3)

As we might put it: how are we to infer the properties of external things from the properties of our sensations?

Unfortunately, Locke's attempts to answer this question were less than satisfactory.[2] The good Bishop Berkeley then solved the sceptical problem by 'the simple, but extraordinarily dramatic expedient' (Pitcher 1971: 92) of eliminating the extra-mental world altogether and identifying objects with collections of sensations:

By sight I have the ideas of light and colours, with their several degrees and variations. By touch I perceive hard and soft, heat and cold, motion and resistance . . . Smelling furnishes me with odours; the palate with tastes; and hearing conveys sounds to the mind in all their variety of tone and composition. . . . And as several of these are observed to accompany each other, they come to be marked by one name, and so to be reputed as one *thing*. Thus, for example, a certain colour, taste, smell, figure and consistence having been observed to go together, are accounted one distinct thing, signified by the name apple. (Berkeley 1710: '1)

There is no difficulty as to how we know about apples, so characterized, as our sensations are precisely the things of which we are directly aware.

With this backdrop in place, let me sketch the first of the two promised interpretations of transcendental idealism; this might be called the 'harsh reading'.[3] The general outlines of Kant's Copernican revolution are undisputed: as space and time are the forms of our human intuition and the categories are the conditions of any discursive intellect, all possible human experience will conform to them; as a result, we can know a priori various mathematical and scientific facts about the world as it is experienced. The harsh reading continues from here by concentrating on the distinction between the appearance, phenomenon, or representation on the one hand, and the thing in itself, noumenon or the transcendental object on the other.[4] That these are distinct items is shown by Kant's insistence that the former are spatio-temporal while the later are not:

Space does not represent any property of things in themselves . . . Space is nothing but the form of all appearances of outer sense. (A26/B42) . . . we are here speaking only of an appearance in space and time, which are not determinations of things in themselves but only of our sensibility. (A493/B522)

This spatio-temporality, enjoyed by appearances, but not by things in themselves, is purely ideal:

[T]his space and this time, and with them all appearances . . . are nothing but representations, and cannot exist outside our mind. (A492/B520)

These appearances are the only objects of our knowledge, but the extra-mental transcendental objects are somehow responsible for them:

---

[2] Bennett (1971: 65–7) gives a summary.

[3] Prominent variations on the harsh reading can be found in Prichard (1909), Strawson (1966), and Guyer (1987).

[4] Considerable subtlety is lost by ignoring the differences in Kant's uses of the alternative terms on either side of this dichotomy (see e.g. Bird (1962): 76–80); Allison (1983: 242–6), but it would take us too far afield to attend to them here.

How things may be in themselves, apart from the representations through which they affect us, is entirely outside our sphere of knowledge. (A190/B235) The non-sensible cause of these representations is completely unknown to us. (A494/B522)

Things in themselves are unknown, but they affect our senses to produce appearances, which we can and do know.

The result of this analysis is a mixed and unsavoury stew. Appearances are experiences, strictly mental—essentially Berkeleian congeries of ideas. Things in themselves are like Locke's external objects, except that even the properties Locke counted as primary—extension, figure—are now taken to be secondary, so that nothing in our experience corresponds to their actual features:

Long before Locke's time, and assuredly since him, it has been generally assumed and granted without detriment to the actual existence of external things that many of their predicates may be said to belong, not to the things in themselves, but to their appearances, and to have no proper existence outside our representation. Heat, colour, and taste, for instance, are of this kind. Now, . . . I go farther and, for weighty reasons, rank as mere appearances the remaining qualities of bodies, also, which are called primary—such as extension, place, and, in general, space, with all that which belongs to it (impenetrability or materiality, shape, etc.). (Kant 1783: 13, Remark II)

We are left with a combination of Berkeleian idealism with a Lockian representative theory, a combination which seems to preserve the bitter of Berkeley—his idealism—without the sweet—his reply to scepticism. A supporter of this interpretation remarks that 'Kant . . . is closer to Berkeley than he acknowledges' (Strawson 1966: 22), and a detractor concludes that 'Kant is seen as a . . . skeptic malgré lui' (Allison 1983: 5–6).

Now even the staunch opponents[5] of the harsh reading admit that there is much in Kant that might seem to support it; e.g. one writes, '[i]t would be foolish to deny that Kant can be interpreted in this way' (Matthews 1969: 205).[6] In response, they point to passages in which Kant rejects Berkeleian idealism (B69–71, B274–5) and the representative theory of perception (B332), and defends himself against the charge of scepticism (B333–4). The trick is to interpret these and similar passages in such a way as to suggest reinterpretations of the passages most supportive of the harsh reading. I will try to do this in the course of sketching out the second promised interpretation of transcendental idealism, which I'll call the 'benign reading'.

The first point to notice is that Kant's appearances are not mere sensations:

That in appearance which corresponds to sensation I term its *matter*; but that which so determines the manifold of appearance that it allows of being ordered in certain relations, I term the *form* of appearance. That in which alone the sensations can be posited and

---

[5] For a sampling, see Bird (1962) in reply to Prichard (1909); Matthews (1969), Bird (1982), and Allison (1983) in reply to Strawson (1966).

[6] Cf. Bird (1962:3): 'it is also true that some things which Kant says appear quite strongly to support Prichard's interpretation'.

ordered in a certain form, cannot itself be sensation; and therefore, while the matter of all appearance is given to us *a posteriori* only, its form must lie ready for the sensations *a priori* in the mind. (A20/B34)

These necessary forms of appearances include space and time, so appearances are spatio-temporal, external to us.[7] But,

When I say that the intuition of outer objects . . . represent[s] the objects . . . in space and in time, as they affect our sense, that is, as they appear, I do not mean to say that these objects are a mere illusion. . . . when I maintain that the quality of space and of time, in conformity with which, as a condition of their existence, I posit . . . bodies . . ., lies in my mode of intuition and not in those objects in themselves, I am not saying that the bodies merely *seem* to be outside me. (B69–70)

The key to understanding this perplexing position lies in the distinction between two perspectives: the empirical and the transcendental.

The central idea here is that questions about appearance and reality are ambiguous: they can be posed, considered, and answered either empirically or transcendentally. At the empirical level, we draw a distinction between the real and the illusory:

The rainbow in a sunny shower may be called a mere appearance, and the rain the thing in itself. This is correct, if the latter concept be taken in a merely physical [empirical] sense. (A45/B63)

The rainbow is an optical phenomenon while the rain itself is real. But, if we now inquire into the reality of the rain at the transcendental level,

We then realise that not only are the drops of rain mere appearances, but that even their round shape, nay even the space in which they fall, are nothing in themselves, but merely modifications or fundamental forms of our sensible intuition, and that the transcendental object remains unknown to us. (A46/B63)

In short, the rain drops are empirically real, but (to a large extent)[8] transcendentally ideal.

Once this distinction is drawn, we see that the appearance and the thing in itself are not two separate objects, but a single thing regarded in different ways:

[W]e can . . . have no knowledge of any object as thing in itself, but only in so far as it is an object of sensible intuition, that is, an appearance. (B xxvi)

The object as appearance is subject to our human forms of sensibility—space and time—and the categories of the discursive understanding—like substance and causation; it is spatio-temporal, external to our minds, and subject to causal laws. Thus Kant is an empirical realist, as opposed to Berkeley's empirical idealist

---

[7] I am speaking, here and in what follows, of objects of outer sense. Objects of inner sense are another story, one I will not pursue.

[8] The 'matter' of the appearance is not transcendentally ideal.

(who holds, e.g. that objects of experience only appear to be extra-mental). On the other hand, the object as it is in itself is not subject to our forms and categories; spatio-temporality and the categories are impositions of our minds, necessary for the object to be experienced, not features of the thing in itself. This is transcendental idealism, as opposed to Locke's transcendental realism (which holds, e.g. that things in themselves are spatio-temporal). Because we are discursive intellects, that is, because all our knowledge arises from the application of concepts to the products of intuition—appearances—we can know nothing of the things in themselves.

This reading goes a long way towards interpreting passages favoured by the harsh theorists, passages in which Kant refers to 'mere appearances' or objects 'in the mind': he is to be read as speaking transcendentally, not empirically. And there is considerable textual support for this distinction:

Our exposition therefore establishes the *reality*, that is, the objective validity, of space in respect of whatever can be presented to us outwardly as object, but also at the same time, the *ideality* of space in respect of things when they are considered in themselves through reason, that is, without regard to the constitution of our sensibility. We assert, then, the *empirical reality* of space, as regards all possible outer experience; and yet at the same time we assert its *transcendental ideality*—in other words, that it is nothing at all, immediately we withdraw the above condition, namely, its limitation to possible experience, and so look upon it as something that underlies things in themselves. (A28/B44)

This reading clearly undermines the harsh interpretation of Kant's appearances as similar to Berkeley's congeries of ideas.

What about the second component of the harsh reading: the representative theory of perception with things in themselves somehow producing appearances by their action on our senses? Though the two-object interpretation has been abandoned, we could still imagine that the way things are in themselves affects us so as to produce the way things appear. First, notice that from the empirical point of view, something very like a Lockian analysis is possible: objects, in the empirical realm of appearances, cause certain responses in the sensory organs of human beings, again regarded empirically, some of these responses produce veridical beliefs, others illusions; these facts can be described and explained scientifically (see, e.g. A28; (Allison 1983: 249)). It's important to realize that Kant's transcendental inquiry does not compromise or even affect the ordinary practice of empirical science (e.g. A39/B56, A393, A30/B45).[9]

But back to the transcendental perspective: do things as they are in themselves affect us to produce appearances? The various Kantian passages that suggest an affirmative answer to this question raise two of the oldest problems for his interpreters: how can we know that things in themselves do this when we can know nothing about them?, and how can things in themselves 'affect' us or 'cause'

---

[9] I come back to this point in a moment.

appearances when causation is a category of the understanding, and as such, only applicable to the world of phenomena? Harsh theorists tend to cite these passages as evidence of deep inconsistencies in Kant, while their opponents give other readings.

To see how one of these other readings might go, consider the following continuation of a passage quoted a moment ago:

> The non-sensible cause of these representations is completely unknown to us, and cannot therefore be intuited by us as object. For such an object would have to be represented as neither in space nor in time . . ., and apart from such conditions, we cannot think any intuition. We may, however, entitle the purely intelligible cause of appearances in general the transcendental object, but merely in order to have something corresponding to sensibility viewed as receptivity. (A494/B522)

The picture is this: sensibility is a capacity to be affected in a certain manner; the affecting agent provides the matter for the resulting appearance; this matter becomes an appearance only after being submitted to the forms of intuition. Under these conditions, what can the affecting agent be? Not an appearance, which is already imbued with spatio-temporal form, but something as yet unaffected by that intuitive processing, in other words, a thing as it is in itself. But to say this is not to gain any contentful knowledge of things in themselves, but only to follow out the concepts involved.[10] To speak of the transcendental object is a harmless way of alluding to the receptive character of sensibility.[11]

The third and final feature of the harsh interpretation is the accusation of scepticism: Kant admits that we can only know appearances, not things as they are in themselves. To this objection, Kant gives a direct reply:

> If by the complaints—*that we have no insight whatsoever into the inner [nature] of things*—it be meant that we cannot conceive by pure understanding what the things which appear to us may be in themselves, they are entirely illegitimate and unreasonable. For what is demanded is that we should be able to know things, and therefore to intuit them, without senses, and therefore that we should have a faculty of knowledge altogether different from the human, and this not only in degree but as regards intuition likewise in kind— in other words, that we should be not men but beings of whom we are unable to say whether they are even possible, much less how they are constituted. (A277–8/B333–4)

The important point is that Kant takes our knowledge of external, spatio-temporal objects to be direct; there is no scepticism here. By contrast, to ask to know

---

[10] In other words, the claim is merely analytic. See Allison (1983: 247–54). Allison (1976) gives a similar reply to another puzzle which I won't discuss, namely: how we can know that the thing in itself is not spatio-temporal (as opposed to the agnostic position that we can't know whether it is or isn't)? See also Allison (1983: 27, 104–14). For a spirited reply, see Guyer (1987: 336–42).

[11] Another benign reading would be to understand the idea (or Idea) of the transcendental object as operating in a regulative, rather than factual sense: to say that the transcendental object causes appearances is to commit ourselves to an unending pursuit of deeper and deeper causal factors. See also Bird (1962: 68–9, 78–9).

what things are like in themselves is to ask to know what things are like when they are not known. To reject this demand is good judgement, not scepticism.

In sum, then, the benign interpretation goes like this: two levels of discourse must be distinguished—empirical and transcendental; from the empirical perspective, appearances are perfectly real, contrasted with sensory illusions; from the transcendental perspective, they are (partly) ideal, contrasted with things as they are in themselves; neither appearances nor things in themselves are Berkeleian congeries of ideas and neither is a Lockian object inscrutable behind its veil of perception. Kant's notions of thing in itself, noumenon, and transcendental object have various uses, but never to give us contentful knowledge of matters beyond the world of experience. The empirical world is the world, and we have direct knowledge of it.

Leaving aside the question of how Kant is best interpreted, I think most of us would find the benign view a more attractive piece of philosophy than the harsh, so let's press on a bit further. At the empirical level, the appearance–reality distinction remains robust—we give the usual account of sensory illusions—but at the transcendental level, we find the transcendental companion to appearances an increasingly empty notion;[12] the thing in itself is not a separate entity, but a sort of empty placeholder[13] that calls attention to the receptivity of our senses, to the role of the forms of intuition and the categories in our knowledge, perhaps to certain methodological principles.[14] In fact, Kant sometimes says that our inquiries at the transcendental level, our critique of pure reason, gives us no knowledge at all, properly speaking:

Its utility, in speculation, ought properly to be only negative, not to extend, but only to clarify our reason, and keep it free from errors—which is already a very great gain. (A11/B25)[15]

Now, given the wide range of transcendental claims that pepper the *Critique*, a rank outsider like myself is tempted to wonder at their status. They certainly sound like knowledge claims: e.g. there are two forms of human sensibility; human knowledge is discursive; or one of my personal favourites

The transcendental unity of apperception is that unity through which all the manifold given in an intuition is united in a concept of the object. (B139)

If these are not knowledge claims, we need some general instruction on how they should be read; if they are knowledge claims, we need to know how they can be

---

[12] See A255/B310: 'the domain that lies out beyond the sphere of appearances is for us empty.'

[13] See A104: 'we have to deal only with our representations . . . that x (the object) which corresponds to them is nothing to us.'

[14] See n. 11.

[15] See also A12/B26: 'It is upon this enquiry, which should be entitled not a doctrine, but only a transcendental critique, that we are now engaged. Its purpose is not to extend knowledge, but only to correct it.'

fitted into Kant's own account of human knowledge: are concepts being applied to intuitions, as is required of all human (i.e. discursive) knowledge? are the claims analytic or synthetic, a priori or a posteriori?

Questions of this sort have led even a sympathetic commentator like Lewis White Beck to a less than satisfying 'meta-critique of pure reason'; he writes:

Not only is it [the claim that 'the only intuition available to us is sensible'] not proved, it is not even a well-formed judgment under the rubrics allowed in the *Critique*, for it is neither analytic nor a posterior, and if it is synthetic yet known a priori, none of the arguments so painfully mounted in the *Critique* to show that such knowledge is possible has anything to do with how we know this (if indeed we do know it). (Beck 1976: 24)

H   Matthews, a leading benign theorist,[16] discusses another such claim:[17]

The statement . . . can certainly not be given any factual content, since the conditions for the empirical application of the concept . . . cannot be met. But the statement is not self-contradictory, and may well have a function, that of expressing the limitations of our experience, which gives it some kind of meaning. The difficulty which is met here is one which arises whenever one tries to talk about the limits of human knowledge . . . The only way in which one can really present the limits of human thought is by showing the confusions and contradictions that arise when one tries to overstep the limits (as Kant does in the Antinomies). But if one does try to *state* the limits (rather than just *showing* them), then the statement, despite its factual appearance, should be interpreted as having a different function. (Matthews 1969: 218–19)

The implicit appeal to Wittgenstein's position in the Tractatus is made explicit in Matthew's discussion.

Given the inherent difficulties associated with these discussions at the transcendental level, the rank outsider might go on to wonder whether the gains are worth the costs, to wonder, that is, what would be lost by resting content with empirical realism neat? Of course, the Kantian opus provides various fascinating and much debated answers to this question—e.g. the role of transcendental idealism in his solutions to the cosmological antinomies and in the transition to his moral philosophy—but the most fundamental role is the one that takes centre stage here:

This transcendental consideration is a duty from which nobody who wishes to make any *a priori* judgments about things can claim exemption. (A263/B319)

Only from the transcendental perspective can we identify conditions on human experience that are not merely empirical, psychological, and only transcendentally necessary conditions of this sort can deliver a priori knowledge of the world.[18] So the transcendental perspective is essential if we are 'to supply a

---

[16] See Ameriks (1982) for a slightly dated, but still useful survey of the literature.

[17] This one concerns the purely intuitive intellect, with which Kant contrasts our human discursive intellect.

[18] Psychological conditions, if they could be isolated, would be part of our empirical theory

touchstone of the value, or lack of value, of all *a priori* knowledge' (A12/B26).

Armed with this thumbnail account of Kant's project and theory, let's now move forward in time and consider the impact of scientific developments since his day. Notice that on the harsh reading, it is hard to see how our empirical knowledge could ever extend beyond what we actually perceive (as in Berkeley) or, perhaps, beyond what we could possibly perceive (as in phenomenalism), and predictably, there are passages that suggest this interpretation:

[F]rom the perception of the attracted iron filings we know of the existence of a magnetic matter pervading all bodies, although the constitution of our organs cuts us off from all immediate perception of this medium. For in accordance with the laws of sensibility and the context of our perceptions, we should, were our senses more refined, come also in an experience upon the immediate empirical intution of it. (A226/B373)

This may sound like phenomenalism, but the benign theorist[19] takes the counter-factual suggestion to be inessential, because Kant also writes:

[T]he knowledge of things as *actual* does not, indeed, demand immediate perception . . . of the object whose existence is to be known. What we do, however, require is the connection of the object with some actual perception, in accordance with the analogies of experience, which define all real connection in an experience in general. (A225/B272)

The key phrase is 'in accordance with the analogies of experience', which include, e.g. the law of causality. Thus, even Newton's findings result from 'the empirical employment of the understanding' (A257/B313), and an early twenti-eth-century Kant could grant empirical reality to, say, atoms, despite our inability to perceive them, because they are suitably connected to what we can perceive by the postulates of empirical thought.[20]

So far so good. Unfortunately, as is well known, further developments in twentieth-century physics cannot be so easily incorporated into the Kantian account. The theory of relativity, for example, denies the objectivity of Kant's time sequence (that is, its observer-independence), replaces his a priori forms of intuition—space and time—with a new conception of space-time, and most famously, forgoes his a priori Euclidean geometry in favour of an a posteriori non-Euclidean geometry. These problems are confined to the forms of intuition, but with the development of quantum mechanics, difficulties spread to the categories of pure understanding; there the law of universal causation is undermined and even the category of substance may be called into question.[21]

---

of ourselves as knowers. What we would learn would be how we are bound to regard objects, not how they really are; part of the payoff of the theory would be to help us better understand how they really are. Transcendental conditions, on the other hand, tell us how the things we experience must be, a priori.

[19] See Allison (1983: 32–3).

[20] This should help explain how considerations at the transcendental level have no effect on the practice of empirical science.

[21] See e.g. Reichenbach (1949), or Körner (1955: 87–9).

Though we have seen that extrapolation using the analogies can carry us to knowledge of matters beyond what is perceived or even perceptible, it is quite another matter to come to know things that conflict with the necessary conditions of experience. There is no particularly Kantian way of dealing with this problem. One reaction would be to suppose that modern physics has taken us beyond the world of appearances—saddled as they are with Kant's forms and categories—and closer to the things in themselves.[22] Another would be to admit that Kant's preferred principles are not necessary after all, that they can, and sometimes should be, revised. Either way, propositions once taken to be true will come to be modified, and the grounds for these modifications will be empirical.[23] At this point, the very distinction Kant's transcendental idealism was designed to found—the distinction between a priori and a posteriori truths about the world—emerges as a distinction without a difference.[24] Once again, pure empirical realism, without the transcendental perspective, presents itself as an alternative worth considering.[25]

My use of the phrase 'distinction without a difference' is meant to recall Quine's critique of Carnap. What I want to suggest now is that this echo is not misplaced, that it is illuminating to view Carnap's position on the a priori and Quine's reaction to it as a variation on the rise and fall of Kantianism as just rehearsed. Perhaps this is not implausible, given that Carnap began his career as a neo-Kantian. In any case, by the early 1930s, Carnap had produced his theory of logical syntax, which the historian Alberto Coffa describes as 'the first genuine alternative . . .[26] to Kant's conception of the a priori' (Coffa 1991: 259).

---

[22] See e.g. Gödel (1946/9: 240, 244–6, 257–9).

[23] Körner (1970: 19) attempts what he calls a 'relativized' version of Kant's synthetic a priori principles. But in the end, he admits that philosophical argumentation has had and can have no evidential influence on changes in these principles (see 69–74). Presumably, then, these changes are made on empirical grounds.

[24] The historian Alberto Coffa (1991: ch. 10) describes the reluctant progress of twentieth-century neo-Kantians to this conclusion. See also Friedman (1983: ch. 1).

[25] This seems a fair description of one strain of Reichenbach's thinking. 'The development of science . . . has led away from Kantian metaphysics. . . . the synthetic principles of knowledge which Kant has regarded as a priori were recognized as a posteriori, as verifiable through experience only and as valid in the restricted sense of empirical hypotheses' (Reichenbach 1949: 307). According to Reichenbach, though Kant claimed to be giving an analysis of pure reason, his methods actually led to an account of (the conceptual underpinnings of) the scientific knowledge of his time (see M. Reichenbach 1965: p. xix). Reichenbach aimed to do the same job for Einsteinian physics (Reichenbach 1949: 310). Finally, on the methods available, he writes, 'modern science . . . has refused the authority of the philosopher who claims to know the truth from intuition, from insight into a world of ideas or into the nature of reason or the principles of being, or from whatever super-empirical source. There is no separate entrance to truth for philosophers. The path of the philosopher is indicated by that of the scientist.' (ibid. 310). This is clearly a version of proto-naturalism.

[26] The full quotation says that Wittgenstein's and Carnap's 'theories of philosophical grammar and of logical syntax may well be regarded as the first genuine alternatives to Kant's conception of the a priori.' The two 'trains of thought were largely independent of each other' (1991: 260).

As indicated earlier, the leading idea of Carnap's position is the linguistic framework. Consider, for example, a linguistic framework for talking about observable things and events in space and time, what Carnap calls 'the thing language' (Carnap 1950: 242–4). Such a language would include the usual logical apparatus–variables, predicates letters, function symbols, names—some logical axioms, and evidential rules specifying what counts as evidence for what. It would be part of this linguistic framework that, for example, certain sensory experiences would count as evidence for certain physical object claims. For speakers of this language, then, it is fairly easy to resolve the truth or falsity of a statement like 'there is a tree outside my window', and from this statement it will follow by pure logic—part of the language—that there are physical objects.

But this is not the resolution philosophers have in mind when they ask whether or not there are physical objects. What the philosopher wants answered is not a question asked inside the thing language—what Carnap calls an 'internal question'—but a question asked outside the thing language, perhaps prior to a decision on whether or not we are justified in speaking the thing language at all. The prior philosopher's question Carnap calls an 'external question'; he writes:

Realists give an affirmative answer, subjective idealists a negative one, and the controversy goes on for centuries without being resolved. And it cannot be solved because it is framed in a wrong way. To be real in the scientific sense means to be an element of the system [i.e. the linguistic framework]; hence this concept cannot be meaningfully applied to the system itself. (Carnap 1950: 243)

In other words, the evidential rules for resolving reality claims about things are to be found inside the thing language; to ask a reality question outside that framework is to pose a pseudo-question.

There is, however, a perfectly reasonable question that can be asked outside the thing language, and that is: should we adopt the thing language?[27] The philosopher imagines that there is a right or wrong answer to this question, based on the prior external question, but Carnap considers it a matter for purely conventional decision in which there is no issue of truth or falsity and tolerance should reign:

In logic, there are no morals. Everyone is at liberty to build up his own logic, i.e. his own form of language, as he wishes. All that is required of him is that, if he wishes to discuss it, he must state his methods clearly, and give syntactic rules instead of philosophical arguments. (Carnap 1934: 52)

Naturally, our reasons for choosing one linguistic framework over another will involve various factual considerations—its effectiveness, fruitfulness, simplicity—

---

[27] Carnap recognizes that the question is odd in the case of this particular linguistic framework because 'there is usually no deliberate choice because we all have accepted the thing language early in our lives as a matter of course' (1950: 243). But this isn't so for other frameworks e.g. the atom framework (see below).

but these are pragmatic reasons in favour of adopting the new framework, not theoretical evidence of its correctness. In Carnap's words:

[I]t would be wrong to describe this situation by saying: 'The fact of the efficiency of the thing language is confirming evidence for the reality of the thing world'; we should rather say instead: 'This fact makes it advisable to accept the thing language.' (Carnap 1950: 244)

To see how this bears on the question of a priority, suppose we've adopted the thing language, and we are considering whether or not to add a new linguistic framework that combines the thing language with something new, the number language. This would mean adding relation symbols for things and numbers, number-theoretic names, predicates symbols, function symbols, axioms (say, first-order Peano Arithmetic), plus evidential rules that would allow us to make assertions like 'there are three apples on the table.' We can well imagine that pragmatic considerations would prompt us to adopt these new conventions.

Now compare the grounds on which our new linguistic framework allows us to make various knowledge claims: if we have experiences so-and-so, we can combine those with the relevant evidential rule and conclude that 'there is a tree outside my window'; if we form a numerical term 5 (SSSS0), we can combine it with '0 is a number' and 'for all x, if x is a number, then Sx is a number' to conclude that '5 is a number.' To justifiably assert that there is a tree outside my window, I need to have certain experiences, but to justifiably assert that 5 is a number, I need nothing beyond the linguistic framework itself. In other words, '5 is a number' is true by virtue of our linguistic conventions; it is, in Carnap's sense, analytic. Thus, '5 is a number' is a priori, and 'there is a tree outside my window' is not. (Notice also that the evidential rule that allows me to infer from my experiences to the assertion about the tree is part of the linguistic framework, and hence a priori.) This is Carnap's new account of a priority.

To the extent that it makes sense to ask questions about Carnap's view in terms borrowed from Kant's, I think it is fair to say that this new account of the a priori does not deliver synthetic truths, but purely linguistic ones. And, obviously, Carnap's a priori truths are not absolute or necessary; they are dependent on conventional choices. Still, there is an important structural similarity between the two positions: in both cases, we have to distinguish two perspectives from which a question can be asked. For Kant, these are the empirical and transcendental perspectives, which deliver different answers to a question like 'is the rain real?': empirically real, transcendentally (at least partly) ideal. For Carnap, these are the internal and external perspectives, which deliver different answers to a question like 'do physical objects exist?': internal to the thing language, the existence of physical objects follows logically from the fact that there are, e.g. trees; external to the thing language, the question is a pragmatic one about the advisability of adopting certain conventions. And for both philosophers, there is a danger of confusion at the higher level: for Kant, we must not be led into thinking that we are gaining knowledge of transcendental objects, or things in themselves; for

Carnap, we must not confuse reasons for adopting a linguistic framework with reasons for believing that its objects really exist, in some external sense. So there is a strong analogy between the old and the new.

In fact, the analogy carries over into some of the difficulties associated with the higher level of inquiry. For Kant, the nature and status of transcendental knowledge proved a mystery. For Carnap, a similar question arises about the status of our meditations on linguistic frameworks: where do these take place? Inside another linguistic framework? If not, we need an account of the non-conventional meta-language; if so, we need to understand why we should adopt the views Carnap espouses over various alternative meta-languages. For example, why should we adopt a linguistic framework that includes the principle of tolerance? What is it to be tolerant about tolerance? As in the case of Kant, these meta-questions are disputed by interpreters of Carnap.[28] Once again, we might well wonder about the tenability of the higher perspective.

Let me leave that problem for now and turn to Quine's criticism of Carnap; I hope to isolate one particular strand from their debate. To see how, consider: when it comes to judging linguistic frameworks, Carnap tells us that 'The purposes for which the language is intended to be used . . . will determine which factors are relevant for the decision' (Carnap 1950: 244).

So the factors involved in choosing linguistic frameworks will be as varied as the goals those frameworks are intended to serve. Our focus, however, is on linguistic frameworks intended for scientific uses; presumably the thing language is the first such language. Carnap tells us that its purpose is 'communicating factual knowledge', that it is judged by such factors as its 'efficiency, fruitfulness and simplicity of use' (ibid.).

Now suppose that we have adopted a scientific framework—complete with logical and evidential rules—and we are considering, internally, a new scientific hypothesis. What sorts of considerations will be raised? Quine suggests that the virtues of such hypotheses are roughly these: simplicity, familiarity of principle, scope, fecundity, and consistency with experiment (see Quine 1955). There may be a subtle trading off of one virtue against another, but an hypothesis that is successful by these criteria is considered well confirmed.

Now suppose, by way of contrast, that we have adopted a scientific framework, and we are considering, not the confirmation of an internal hypothesis, but a move to another, more inclusive linguistic framework; this time, we are considering what Carnap calls an 'ontological question', that is, the addition of syntactic apparatus and evidential rules for a new sort of entities, like atoms, or sets. What sorts of considerations will be raised? Quine's suggestion is that the very same theoretical virtues will be relevant this time around:

---

[28] See Coffa (1991: ch. 17) for the view that 'behind the first level semantic conventionalism there is a second level semantic factualism' (ibid. 322). For an opposing perspective, see Goldfarb (OAC).

Our acceptance of an ontology is, I think, similar in principle to our acceptance of a scientific theory, say a system of physics: we adopt, at least insofar as we are reasonable, the simplest conceptual scheme into which the disordered fragments of raw experience can be fitted and arranged. Our ontology is determined once we have fixed upon the over-all conceptual scheme which is to accommodate science in the broadest sense; and the considerations which determine a reasonable construction of any part of that conceptual scheme, for example, the biological or the physical part, are not different in kind from the considerations which determine a reasonable construction of the whole. To whatever extent the adoption of any system of scientific theory may be said to be a matter of language, the same—but no more—may be said of the adoption of an ontology. (Quine 1948: 16–17)

In other words, the methods used to evaluate ordinary scientific hypotheses are the same as those used to decide Carnap's ontological questions.[29] Thus, Quine argues, there is no methological basis for the distinction Carnap draws.[30]

The parallel with the Kantian saga so recently rehearsed is obvious. The problem for the twentieth-century Kantian is that the criteria for modifying our intuitive and categorical principles are indistinguishable from those for modifying our scientific beliefs generally. This Quinean objection to Carnap is that the criteria for adopting linguistic frameworks are indistinguishable from the criteria for adopting scientific hypotheses generally. In both cases, the cherished distinction seems groundless.

Quine's response to this situation is analogous to the response that tempted us in the Kantian case: forgo the transcendental level; fall back on empirical realism neat. In the more recent context, this is Quinean naturalism: 'the recognition that it is within science itself, and not in some prior philosophy, that reality is to be identified and described' (Quine 1981a: 21) Carnap's external pseudo-questions—'are there physical objects?' or 'are there numbers?'—are actually ordinary scientific questions: 'Ontological questions, under this view, are on a par with questions of natural science' (Quine 1951a: 45).

Epistemological questions are also to be addressed from within science:

Naturalism does not repudiate epistemology, but assimilates it to empirical psychology. Science itself tells us that our information about the world is limited to irritations of our surfaces, and then the epistemological question is in turn a question within science: the question how we human animals can have managed to arrive at science from such limited information. (Quine 1975: 72)

---

[29] Carnap sometimes argues that the methods are different because the question of whether or not to adopt a linguistic framework is based on considerations that are a matter of degree, while the question of the truth of a scientific claim is all or nothing: 'these questions [as to the pragmatic virtues of the thing language] cannot be identified with the question of realism. They are not yes-no questions, but questions of degree' (Carnap 1950: 244). But the comparison is faulty: the considerations that lead us to adopt a scientific claim are also matters of degree.

[30] At least, insofar as it applies to scientific frameworks. If we imagine a framework with a very different goal, say to generate sentences with particular aesthetic qualities, then the internal criteria for deciding which sentence to 'assert' might be very different from the external criteria for deciding which such framework to adopt.

To describe naturalistic philosophy in general, Quine appeals to a favourite image:

> Neurath has likened science to a boat which, if we are to rebuild it, we must rebuild plank by plank while staying afloat in it. (Quine 1960: 3) The naturalistic philosopher begins his reasoning within the inherited world theory as a going concern. He tentatively believes all of it, but believes also that some unidentified portions are wrong. He tries to improve, clarify, and understand the system from within. He is the busy sailor adrift on Neurath's boat. (Quine 1975: 72)

For the naturalist, there is no higher perspective, where transcendental or other extra-scientific considerations hold sway. The naturalist operates 'from the point of view of our own science, which is the only point of view I can offer' (Quine 1981b: 181).

A similar rejection of the transcendental level is found in Arthur Fine's 'natural ontological attitude', or NOA.[31] The context here is the realism-anti-realism debates of the late 1970s and early 1980s, exemplified, for example, by Putnam's attack on 'metaphysical realism'[32] and van Fraassen's agnosticism about unobservables.[33] As Fine understands it, the impulse towards realism is actually based in 'homely' beliefs, which, he says

> I will put it in the first person. I certainly trust the evidence of my senses, on the whole, with regard to the existence and features of everyday objects. And I have similar confidence in the system of 'check, double-check, check, triple-check' of scientific investigation, as well as the other safeguards built into the institutions of science. So, if the scientists tell me that there really are molecules, and atoms, and y/J particles, and, who knows maybe even quarks, then so be it.[34] (Fine 1986: 126-7)

From this point of view, we can ask after the relations between humans, as described in psychology, physiology, linguistics, etc., and the world, as described in physics, chemistry, geology, etc., and draw conclusions about the relations between sentences and the world, an investigation that may result in a correspondence theory of truth or a deflationary theory of truth or some other theory of truth or no theory of truth at all, depending how things go.[35] But however they go, this theory will be just one part of our overall scientific theory of the world.

---

[31] See Fine (1986 and 1996a). In the latter (1996a: 176-7), Fine allies NOA with Quinean naturalism while rejecting some of its descendants. Much as I applaud the general outlines of NOA, I fear I must also distance myself from some of Fine's further views (e.g. see Fine 1996b).

[32] See e.g. Putnam (1981).

[33] See van Fraassen (1980).

[34] This quotation may suggest that NOA eschews the normative, but this isn't so (see Fine 1996a: 177-8). Whatever Fine's precise views on the subject, Quine explicitly locates the norms of science within science (see Quine 1981b: 181), and I follow him in this see my (1977: III.3).

[35] I take this to be the debate that began with Field (1972) and Leeds (1978). Horwich (1990) presents a recent summary. By taking this to be a scientific question on which the jury is still out, I find myself in disagreement with Quine (see Quine 1970: 10-13), and perhaps also with Fine (cf. Fine 1986: 130 and 134; 1996a: 184; and 1996b).

On these matters, Putnam and van Fraassen agree with the NOAer, but they don't stop here; each, in his own way, goes beyond science, to a higher level. There Putnam distinguishes metaphysical realism, which adds to NOA's core an extra-scientific correspondence theory of truth, and internal realism, which adds to the same core a Peircean analysis of truth as warranted assertability in the ideal limit.[36] Focused on the problem of ontology rather than truth, van Fraassen adds an extra level of epistemological analysis where we must abstain from belief in molecules and atoms and electrons, despite our acceptance of these same entities for scientific purposes. Here the holder of our homely beliefs will be tempted to object that atoms *really* do exist, thus embodying Kant's 'incautious listener', faced with 'a question . . . absurd in itself', who then gives 'an answer where none is required' (A58/B82–3): he wants to insist on the reality of atoms, but all the genuine scientific evidence, though accepted at the lower level, has been ruled out of bounds at the higher level; the frustrated Scientific Realist ends by stomping his foot. Fine's proposal is that we rest with the natural ontological attitude and resist the temptation to engage in extra-scientific debate.

To subject our naturalism to the same challenge put to both Kant and Carnap, we should ask: is naturalism itself a scientific thesis? I think the right answer to this question is that naturalism is not a thesis at all, but an approach. The naturalistic philosopher is the Neurathian sailor, working within science to understand, clarify, and improve science; she will treat philosophical questions on a par with other scientific questions, insofar as this is possible; faced with first philosophical demands—that is, questions and solutions that require extra-scientific methods—she will respond with befuddlement, for she knows no such methods; from her scientific perspective, she is sceptical that there are such methods, but she has no a priori argument that there are such methods, but she has no a priori argument that there are none; until such methods are explained and justified, she will simply set aside the challenges of first philosophy and get on with her naturalistic business. Naturalism contrasts with both Kantianism and Carnapianism in forgoing any 'higher-level' considerations.

Now while I am in hearty agreement with these naturalistic sentiments, I find myself less comfortable with the Quinean argument that got us here: the time has come for me to own up to my departures from Quinean orthodoxy. These are several,[37] but the one relevant here concerns the thoroughgoing pragmatism that

---

[36] There is some irony in Putnam's move to the extra scientific perspective, given that he was once a devout naturalist: 'it is silly to agree that a reason for believing that p warrants accepting p in all scientific circumstances, and then to add "but even so it is *not good enough*". Such a judgment could only be made if one accepted a tran-scientific method as superior to the scientific method; but this philosopher, at least, has no interest in doing *that*' (1971: 356).

[37] My (1997) describes additional disagreements over the force of indispensability arguments and the treatment of pure mathematics.

underlies the objection to Carnap we've been considering.[38] I've told this story in some detail elsewhere (1997: II.6), so I will be as brief as possible here.

Consider, for illustration, the case of atomic theory. Attention to the historical details of the debate over atoms, as it stood at the turn of the century, reveals some surprising features. First, by any account of confirmation that depends on an enumeration of general theoretical virtues (like Quine's), atomic theory was extremely well confirmed; it was central to chemistry in all theories of composition and combination and to physics in the wide-ranging kinetic theory. Second, despite this, there was still scepticism among respectable scientists for respectable reasons about the existence of atoms. In 1905, Einstein set out to provide the theoretical grounds for a proof of atomic reality, and soon thereafter, Perrin produced experimental results. Only then did the scientific community as a whole embrace atoms. The moral of the story is that no amount of general theoretical support was enough; the new unobservables had to be 'detected' by a sufficiently direct procedure.

Now consider this case from the perspective of the debate between Carnap and Quine. For Carnap, the question 'are there atoms?' is prime example of an external ontological question, just the sort of thing that van Fraassen and the foot-stomping Scientific Realist would vigorously debate. Carnap would replace this external pseudo-question with a legitimate one—should we adopt the atoms language?—and this question he would answer on pragmatic grounds. Quine, regarding the original ontological question as an ordinary scientific question, would decide it on the same pragmatic grounds. But historical analysis reveals that the question of atoms, however interpreted, was not, in fact, decided on such grounds. Atoms were still considered problematic after they had passed every pragmatic test imaginable.

Now consider, for contrast, another example, one not so neatly comprised in a single historical episode, but I think no less compelling for that. In the beginning, the motion of a medium-sized physical object was represented by a table: at time t, the object is at position x, at t', position x', and so on. Eventually, this crude approach gave way to our current notion of a function from real numbers to real numbers—with the independent variable representing time and the dependent variable representing position—which allowed the powerful methods of the calculus to be applied. Nowadays, space and time have been replaced by space-time, and ordinary calculus by more elaborate analysis, but the principle of using continuum mathematics for such purposes has remained the same. The success of this way of doing things can hardly be exaggerated.

Notice that this account of motion depends on a strong assumption about the

---

[38] This is not Quine's only objection; e.g. in his (1951*b*), Quine argues that Carnap's distinction between external and internal questions rests on a prior distinction between 'category questions' and 'subclass questions', which he dismisses as 'a distinction which is not invariant under logically irrelevant changes of typography' (ibid. 210).

structure of physical reality, namely that space-time is continuous, that it is, or is isomorphic to, a continuous manifold. Despite this, physicists still consider the question of the continuity of time and space to be open. On the face of it, this case parallels the case of atomic theory: the pragmatic success of a theory is not enough to fully confirm it. But, in fact, or so I would claim, there are striking disanalogies. On the one hand, the continuity of time and space have not been a subject of debate in the same sense that atoms were before Einstein and Perrin: no one worries over this use of continuum mathematics, no one demands experimental confirmation, no one demands reform.[39] On the other hand, the success of the structural assumption that space-time is continuous has not counted in favour of its truth: despite the overall confirmation of the theory, physicists consider the question to be wide open. In cases like these, it's hard to avoid the conclusion that these assumptions are not being treated on the same epistemic terms as other scientific hypotheses, that the conditions for adoption are not as stringent and the confirmatory benefits of their success are negligible. Indeed, it seems that scientists feel free to use whatever mathematical apparatus they find convenient and effective, without regard for the abstract ontology required, and more to the point, without regard for the physical structural assumptions presupposed.

Of course, I can't be considered to have argued for this claim here,[40] but supposing it is correct, then there is a distinction to be drawn, within science, between hypotheses that are subjected to the full range of empirical testing and available for full confirmation, and those that are not; indeed, hypotheses in the second category might well be described as having been adopted on pragmatic, conventional grounds. Essentially, what I am objecting to is Quine's blanket pragmatic account of the method of science; it seems to me that the practice of actual science is much more varied that any such general account allows. And, in particular, there does seem to be room for the sort of methodological distinction on which Carnap rests his two-level picture. The line between the conventional and the ordinary scientific hypotheses doesn't fall quite where Carnap would have it—the embrace of atoms was not a pragmatic linguistic change but an internal, scientific one—but there remains room to classify such hypotheses as the continuity of space-time as conventional/pragmatic, and for so distinguishing it from theoretical/empirical hypotheses like the existence of atoms. Given the centrality of the two-level model to Carnap's account of the a priori and given that Quine's views lead him to reject the distinction between a priori and a posteriori entirely,

---

[39] I am ignoring the fact that anomalies in quantum field theory, anomalies with roots as far back as classical electrodynamics, have led some physicists to question the small-scale continuity of space-time. (See my 1997: II.6, for discussion.) For our purposes, it is enough to notice that another sort of objection was advanced against atoms before Einstein and Perrin: that their existence was too theoretical, that it hadn't been sufficiently verified. No such objection has been raised to continuous space-time.

[40] Again, my (1997: II.6) contains some more detail.

we should take a moment to reconsider: could a viable two-level account of the a priori be based on this methodological distinction?[41]

As I've said, I myself favour a one-level view, a version of naturalism, largely because I don't believe in a perspective outside of science from which science, or even scientific language, is evaluated. The nagging difficulties that arise when we try to characterize the methods of that higher level of inquiry are symptoms of the disorientation that sets in when we imagine ourselves to be operating outside of science. In the case of Carnap, or rather, the neo-Carnap we're imagining, I think the interplay between conventional/pragmatic and the theoretical/empirical is too involved and too important to permit the two to be segregated as the two-level picture requires. By the time I'm done here, I hope to have made this claim plausible.

That said, let me suspend (for the moment) this question about the two-level neo-Carnapian to consider one last approach to the a priori. Quine, as we've noted, rejects the distinction when he moves to his naturalistic web of belief, explaining away our impression that some truths are a priori as an illusion of a difference in kind created by a difference in degree. But given that we've departed from Quine's uniform vision of scientific methods, perhaps this conclusion comes too quickly for us. We wonder: does the rejection of a two-level view, by itself, doom prospects for the a priori? To put it another way: can our naturalist give an account of the a priori without appeal to a 'higher level' of inquiry?

For the record, Quine figures that the external–internal distinction is not essential to Carnap's account:

No more than the distinction between *analytic* and *synthetic* is needed in support of Carnap's doctrine that statements commonly thought of as ontological . . .[42] are analytic or contradictory given the language. No more than the distinction between analytic and synthetic is needed in support of his doctrine that the statements commonly thought of as ontological are proper matters of contention only in the form of linguistic proposals. (Quine 1951b: 210)

Without the higher level, such an account will lack some of the virtues of Carnap's two-level version, simply because the two-level view allows us to ask after the status of a given statement in two distinct ways: at the higher level, we consider various options and plump for one on conventional/pragmatic grounds,

---

[41] I do not assume that a neo-Carnapian could, in fact, draw a workable distinction between the empirical and the conventional based on the methodological observations rehearsed here, but I will extend the benefit of the doubt for the sake of argument.

[42] In this hiatus, Quine gives examples: 'statements such as "There are physical objects", "There are classes", "There are numbers". The first of these seems to me problematic. On my understanding of the thing language, it is the evidential rules that are analytic—such-and-such counts as evidence that there is a tree outside my window—and it takes experiential input to draw the conclusion that the tree exists. 'There are physical objects' then follows from the existence of the tree, but this conclusion does not follow from features of the linguistic framework alone; it is not analytic.

but once we'd adopted the winning linguistic framework and begun to work within it, as long as we continue to speak that language, the framework principles are analytic and unrevisable. In other words, a statement that's optional and revisable when viewed at the higher level is necessary and unrevisable when viewed from the lower level, and thus, robustly a priori at that lower level. When the naturalist drops the higher level, all these considerations are placed side-by-side; the statement in question is recognized as the revisable result of a conventional/pragmatic decision.

What might be salvaged—and this is what Quine has in mind—is a distinction between revisions that are changes in theory and revisions that are changes in language, and it is this distinction that Carnap never gave up:

I should make a distinction between two kinds of readjustment in the case of a conflict with experience, namely, between a change in the language, and a mere change in or addition of, a truth-value ascribed to an indeterminate statement . . . A change of the first kind constitutes a radical alteration, sometimes a revolution, and it occurs only at certain historically decisive points in the development of science. On the other hand, changes of the second kind occur every minute. (Carnap 1963: 921)

If this distinction can be drawn, then even the one-level naturalist can say that an analytic statement is unrevisable *short of a change in language*, and in that sense, might be considered a priori.

Quine's reaction to this possibility, famously, is to reject the distinction between analytic and synthetic. Whatever we think of his arguments, even if a scientific inquiry into the semantics of natural language, within linguistic theory, were to conclude that 'Bachelors are unmarried' is true by virtue of the meanings of the terms involved, I sincerely doubt that any such inquiry would generate the outcome that 'continuous manifold' is part of the meaning of the term 'space', or that the axiom of replacement is part of the meaning of the term 'set'. So I will follow Quine in rejecting this as a viable account of the a priori.[43] But our methodological inquiries—part of our scientific study of science itself—suggest that the conventional/pragmatic vs. theoretical/empirical distinction, which Quine also rejects, might be revived without appeal to an external level of analysis. So perhaps there is room for the a priori on a one-level view, after all, making use of this methodological distinction.[44]

To see the trouble with this approach, consider the notion of knowledge. On the two-level model, the conventional/pragmatic considerations are raised externally, where knowledge is not an issue; once we adopt the linguistic framework, we have a scientific account of what counts as knowledge, and the elements of the

---

[43] I mean 'viable' in the Kantian tradition of delivering some (seemingly) contentful scientific or mathematical claims (like the examples in the text, the sort of thing Kant might count as synthetic) as a priori. (Notice that none of this discussion concerns the status of elementary logic.)

[44] See n. 41.

linguistic framework count as a priori knowledge by that standard. On the one-level approach, the conventional/pragmatic considerations and the theoretical/empirical considerations are both raised internally; we recognize internally that some hypotheses are adopted as a result of full empirical testing and confirmation, while others are adopted for other reasons. Given our scientific account of knowledge, it seems unavoidable that the former hypotheses will count as knowledge, while the latter will not. Our discussion of the continuity of space-time brings this point into high relief: it would seem odd to say that scientists know space-time to be continuous. Space-time is so represented, but the question of its actual continuity remains open. Finally, if we don't know something, we can't be said to know it a priori. So it seems a neo-Carnapian defence of the distinction between changes of meaning and changes of theory that rests on the methodological distinction between the conventional/pragmatic and the theoretical/empirical elements of our theory will not support a viable notion of a priori knowledge.

This leaves the one-level theorist with the brute fact of this methodological distinction. As naturalistic philosophers of science, we will try to understand and explain this phenomenon of scientific practice, even if we can't use it to construct a notion of a priority: to say that some hypotheses are adopted for conventional/pragmatic reasons is only the barest beginning of an account of how such hypotheses function. In particular cases, a range of deep and important methodological questions arise. Sometimes, we are unsure whether or not a given hypothesis is, in fact, conventional/pragmatic (e.g. physical geometry, and careful analysis is required.[45] Sometimes, conventional/pragmatic hypotheses are adopted even when they are known to be false, like the hypothesis that matter is continuous in fluid dynamics; in such cases, we need to ask what it is about the world that is reflected in that idealization, what it is about the world that supports the effectiveness of that idealization.[46] Sometimes, the status of a conventional/pragmatic hypothesis is unknown, like the hypothesis of continuous space-time. (In fact, I suspect we are willing to adopt hypotheses like continuous space-time in the spirit we do because we are confident that even if they should turn out to be false, they will remain good idealizations for many purposes.) In such cases, we need to ask whether the hypothesis is amenable to further testing: if not, why not?, and if so, how? These are among the important naturalistic

---

[45] e.g. consider Friedman's debate with the proto-naturalist Reichenbach on the status of geometry: Friedman aims to show that empirically equivalent theories with alternative geometries are not equally valid, as Reichenbach claims, but rather, that 'there are methodological principles or criteria . . . that are capable of singling out one theory from a class of empirically equivalent theories' (Friedman 1983: 268). Friedman insists that 'These criteria can be justified only by showing that they tend to produce true hypotheses in actual scientific practice in the real world . . . Hence, any justification of our actual inductive methods must itself be empirical' (ibid. 273). For these two, the debate over the conventional or empirical status of geometry takes place within a scientific (i.e. naturalistic) study of science.

[46] I am grateful to Mark Wilson for making the importance of these issues clear to me.

inquiries that emerge out of our foiled pursuit of a one-level account of the a priori.

Finally, let me double back to the two-level neo-Carnapian account of the a priori that we left hanging a moment ago. At the time, I suggested that the interplay between conventional/pragmatic and the theoretical/empirical is too involved and too important to permit the two to be segregated as the two-level picture requires. We can now see more clearly what I had in mind. If the conventional/pragmatic is forcibly separated from the theoretical/empirical, if the decision on scientific language must be made before scientific inquiry begins, then the differences between hypotheses of these two sorts cannot be studied side-by-side, using scientific methods. This means that the important naturalistic inquiries just sketched cannot find a home: from the external perspective, we are not doing science, but pragmatics, so questions of literal truth or falsehood cannot be raised; from the internal perspective, the linguistic framework is a given, known a priori, not subject to further debate. In other words, because the pragmatic considerations arise at the higher level, their interactions with internal facts cannot be discussed internally. So the trade-off is between a two-level view that provides an account of the a priori, but blocks a certain style of scientific inquiry into science, and a one-level view that fails to underwrite the a priori–a posteriori distinction, but encourages this style of inquiry into science. For my money, there is no contest.

Let me sum up. I have traced the role of an extra-scientific level of analysis in the Kantian and the Carnapian accounts of a priori knowledge, and I have suggested that such accounts suffer from two serious difficulties: how to explain the status of this extra-scientific analysis, and how to differentiate revisions in the purportedly a priori claims from ordinary scientific progress. In my opinion, these difficulties are symptoms of a more basic error, an error to be avoided by adopting a one-level, naturalistic approach, that is, by adopting the methods and results of science and working from within to understand, clarify, and improve them. Any discussion of naturalism these days is—overtly or covertly—an attempt to define the term; I've tried to locate the fundamental naturalistic impulse in a stubborn scepticism about any of the recurring two-level philosophies, about any philosophy that posits an extra-scientific perspective from which to view science. I've suggested that naturalism of this sort turns up a methodological contrast within science between hypotheses subjected to the full range of empirical testing and potentially confirmed thereby and hypotheses adopted for conventional, pragmatic reasons. A neo-Carnapian might hope to revive a two-level or a one-level account of the a priori based on this distinction, and I've explained why I think the one is not advisable and the other not viable. But while naturalism may doom the time-honoured notion of a priori knowledge, it also highlights new methodological questions about the detail of scientific practice, a trade-off that I for one am quite willing to accept.

References

Allison, Henry (1976), 'The non-spatiality of things in themselves for Kant', *Journal of the History of Philosophy* 14: 313–21.

——— (1983), *Kant's Transcendental Idealism* (New Haven, Conn.: Yale University Press).

Ameriks, Karl (1982), 'Recent work on Kant's theoretical philosophy', *American Philosophical Quarterly* 19: 1–24.

Beck, Lewis White (1976), 'Toward a meta-critique of pure reason', repr. in his *Essays on Kant and Hume* (New Haven, Conn.: Yale University Press, 1978): 20–37.

Bennett, Jonathan (1971), *Locke, Berkeley, Hume* (Oxford: Oxford University Press).

Berkeley, George (1710), *Principles of Human Knowledge*, repr. in D. Armstrong (ed.), *Berkeley's Philosophical Writings* (New York: Macmillan, 1965).

Bird, Graham (1962), *Kant's Theory of Knowledge* (New York: Routledge & Kegan Paul).

——— (1982), 'Kant's transcendental idealism', in G. Vesey (ed.), *Idealism Past and Present* (Cambridge: Cambridge University Press): 71-92.

Carnap, Rudolf (1934), *Logical Syntax of Language*, trans. A. Smeaton (London: Routledge & Kegan Paul, 1937).

——— (1950), 'Empiricism, semantics and ontology', repr. in P. Benacerraf and H. Putnam (eds.), *Philosophy of Mathematics*, 2nd edn. (Cambridge: Cambridge University Press, 1983): 241–57.

——— (1963), 'Replies and systematic expositions', in P. Schilpp (ed.), *The Philosophy of Rudolf Carnap* (La Salle, Ill.: Open Court): 859–1013.

Coffa, Alberto (1991), *The Semantic Tradition from Kant to Carnap: to the Vienna Station*, L. Wessels (ed.) (Cambridge: Cambridge University Press).

Field, Hartry (1972), 'Tarski's theory of truth', *Journal of Philosophy* 69, pp. 347–75.

Fine, Arthur (1986), *The Shaky Game* (Chicago: University of Chicago Press).

——— (1996a), 'Afterword', in 2nd edn. (1986): 173–201.

——— (1996b), 'Science made up: constructivist sociology of scientific knowledge', in P. Galison and D. Stump (eds.), *The Disunity of Science* (Stanford: Stanford University Press): 231–54.

Friedman, Michael (1983), *Foundations of Space-time Theories* (Princeton: Princeton University Press).

Goldfarb, Warren (OAC), 'On Alberto Coffa's *The Semantic Tradition from Kant to Carnap*' unpublished.

Gödel, Kurt (1946/9), 'Some observations about the relationship between theory of relativity and Kantian philosophy', in his *Collected Works*, iii, S. Feferman *et al.* (eds.) (Oxford: Oxford University Press, 1995): 230–59.

Guyer, Paul (1987), *Kant and the Claims of Knowledge* (Cambridge: Cambridge University Press).

Horwich, Paul (1990), *Truth* (Oxford: Blackwell).

Kant, Immanuel (1781/7), *Critique of Pure Reason*, trans. Norman Kemp Smith (New York: St Martin's Press, 1965). (References to the Critique in the text are given in the customary style: the page number in the first (A) edition followed by the page number in the second (B) edition.)

——— (1783), *Prolegomena to Any Future Metaphysics*, trans. L. W. Beck (Indianapolis: Bobbs-Merrill, 1950).

Körner, Stephan (1955), *Kant* (London: Penguin Books).

Körner, Stephan (1970), *Categorial Frameworks* (Oxford: Blackwell).

Leeds, Stephen (1978), 'Theories of reference and truth', *Erkenntnis* 13: 111–29.

Locke, John (1690), *An Essay Concerning Human Understanding* (New York: Dover, 1959).

Maddy, Penelope (1997), *Naturalism in Mathematics* (Oxford: Oxford University Press).

Matthews, H. E. (1969), 'Strawson on transcendental idealism', *Philosophical Quarterly* 19: 204–20.

Pitcher, George (1977), *Berkeley* (London: Routledge & Kegan Paul).

Prichard, H. A. (1909), *Kant's Theory of Knowledge* (Oxford: Oxford University Press).

Putnam, Hilary (1971), 'Philosophy of Logic', repr. in his *Philosophical Papers*, ii, 2nd edn. (Cambridge: Cambridge University Press, 1979): 323–57.

—— (1981), *Reason, Truth and History* (Cambridge: Cambridge University Press).

Quine, W. V. (1948), 'On what there is', repr. in his (1980): 1–19.

—— (1951a), 'Two dogmas of empiricism', repr. in his (1980): 20–46.

—— (1951b), 'Carnap's view on ontology', repr. in his (1976): 203–11.

—— (1955), 'Posits and reality', repr. in his (1976): 246–57.

—— (1960), *Word and Object* (Cambridge, Mass.: MIT Press).

—— (1970), *Philosophy of Logic* (Englewood Cliffs, NJ: Prentice-Hall).

—— (1975), 'Five milestones of empiricism', repr. in *Theories and Things* (Cambridge, Mass.: Harvard University Press): 67–72.

—— (1976), *The Ways of Paradox*, revised edn. (Cambridge, Mass.: Harvard University Press).

—— (1980), *From a Logical Point of View*, 2nd edn. (Cambridge, Mass: Harvard University Press).

—— (1981a), 'Things and their place in theories', in (1981b) 1–23.

—— (1981b), *Theories and Things* (Cambridge, Mass: Harvard University Press).

Reichenbach, Hans (1920), *The Theory of Relativity and A Priori Knowledge*, trans. M. Reichenbach (Berkeley and Los Angeles: University of California Press, 1965).

—— (1949), 'The Philosophical significance of the theory of relativity', in P. A. Schilpp (ed.), *Albert Einstein: Philosopher-Scientist* (La Salle, Ill.: Open Court): 287–311.

Reichenbach, Maria (1965), 'Introduction to Reichenbach' (1920): pp. xi–xliv.

Strawson, P. F. (1966), *The Bounds of Sense* (London: Methuen).

Van Fraassen, Bas (1980), *The Scientific Image* (Oxford: Oxford University Press).