13

Carnap, Quine, and Williamson

Metaphysics, Semantics, and Science

Gary Kemp

For all their much-discussed disagreements over analyticity and ontology, Quine shared Carnap's more fundamental commitment to 'scientific philosophy'. The primary service of legitimate philosophy is to clarify, to make more precise, and to make more explicit the methods and deliverances of science. Sometimes the philosopher may extend those methods to novel areas – as in Quine's description of radical translation or Carnap's of semantics and modality – but decidedly there remains no room for metaphysics in some extra-scientific sense.

The mature development in Quine was what he called his 'naturalism', the view that 'it is only within science, and not in some prior philosophy, that reality is to be identified and described' (1981d, 21). For his part, Carnap wrote that the 'logic of science … is in the process of cutting itself loose from philosophy and becoming a properly scientific field, where all work is done according to strict scientific methods and not by means of "higher" or "deeper" insights' (1934/1937, 46; evidently, in a sign of the times, Carnap intends 'philosophy' pejoratively, and that his work is properly scientific). Of course, there are well-known difficulties in calling Carnap a naturalist in Quine's sense. For my purposes this issue can be avoided, but one difference that will figure later is that unlike Quine, Carnap

thinks of the proper task of the philosopher as that of designing entire symbolic languages for science, and of philosophical communication in the vernacular as but a practically necessary propaedeutic to the expression of maximally objective, interest-free theoretical truths in a symbolic language. Roughly speaking, whereas for Carnap pragmatic matters and theoretical matters are mutually exclusive, for Quine *all* matters are both pragmatic and theoretical. Nevertheless, Carnap and Quine are agreed that many notions of common sense and many questions couched in ordinary language, as well as the many untethered doctrines of the speculative thinker, are simply not suitable as they stand for serious philosophical theories or serious philosophical examination. Philosophical questions themselves must be shown to arise within science if they are genuinely to be factual ones.

In this essay I compare this shared ground between Quine's view and Carnap's with a recent estimate of the role of philosophy vis-à-vis science, which might be summed up as the re-enshrinement of metaphysics as sharing the stage with science: that advanced by Timothy Williamson (in his book *The Philosophy of Philosophy*).[1] Now Williamson rejects both Quinean naturalism (1–2) and Carnapian positivism (51), but he does not think of his view as an explicit response to Quine or Carnap, and more generally he does not think of their views as paradigmatic of the 'linguistic philosophy' he explicitly rejects – he thinks rather of such figures as A. J. Ayer, Peter Strawson, and Michael Dummett. There is nevertheless a basic disagreement between Williamson, on the one hand, and Carnap and

---

[1] There is also Williamson's less substantive 'How Did We Get Here from There? The Transformation of Analytic Philosophy' (2011). *Modal Logic as Metaphysics* (2015) contains more dimensions of his view, but space here is limited.

Quine, on the other, one that will require further understanding of the point gestured at in the last paragraph to appreciate. The disagreement is partly about science generally and partly about semantics in particular. The upshot is that from either Carnap's point of view or Quine's, this recent ascendance of metaphysics – partly abetted by Williamson – will seem an unjustified departure from science.

13.1   Carnap and Quine: Some Similarities and a Difference

All scientists properly so-called strive among other things for maximum clarity – ideally, for their products to be capable of formalisation, even if in most cases it would be pointless actually to carry it out. But neither Quine nor Carnap assume a narrow view of science: they intend 'science' in the broader sense of *Wissenschaft* (the term is partly an honorific). We have a cluster of sciences, the centre being mathematics, then physics and chemistry, then geology, biology, psychology, cognitive science, social science, linguistics, economics, history, anthropology, politics, religious studies, and the like, as we move towards the outskirts. We can form several rough gradations – which sometimes coincide and sometimes do not – of precision, objectivity, generality, and the acceptance of general laws; of freedom from being infected with practical interests, story-telling, metaphor, and rhetoric; and of methodological disputes – relatively absent at the one end and ubiquitous at the other. Nevertheless, Carnap and Quine both value the formal expression of science, at least for those with philosophical interests. Carnap famously devises new artificial

languages for the purpose, as in *The Logical Syntax of Language* (1934), whereas Quine speaks of 'regimentation', as in *Word and Object* (1960), and partly practises it.

For both Carnap and Quine, common-sense ideas, or 'folk theories' and the ordinary language used to express them, are often decidedly unsuitable. Consider Zeno's paradoxes. It took a colossal amount of time and effort until, in the nineteenth century, mathematics and physics could give satisfying answers to these ancient riddles; this required much painstaking clarification and rigorisation of certain conceptual implements – differentiation, integration, the limit, and so on – on the part of Newton, Cauchy, Weierstrass, and many others. 'Ordinary language is only loosely factual' (2008a, 285), said Quine, and no wonder; it arose amidst the traffic of ordinary people with ordinary needs, and it would be strange if it had more determinacy than such needs called for. Commenting on the conundrums of personal identity, Quine writes: 'They are questions about the concept of person, or the word "person", which, like most words, goes vague in contexts where it has not been needed. When need does arise in hitherto unneeded contexts, we adopt a convention, or receive a disguised one from the Supreme Court' (Quine 1995a, 39). Quine and Carnap agree that such conceptual engineering is not happily characterised as 'analysis' – with the implication of finding out what was there in language all along, as suggested by the analogy with chemical analysis – but instead should be characterised as 'explication' in their special sense: the replacement of existing words or phrases with more precise ones, normally homophonic with the old ones but not necessarily. Where what is in question is a not-yet-explicated φ, it is often pointless to insist that there is an essence of φ to be drawn out, either by 'conceptual analysis' of the term for φ or inquiry into the 'nature'

of φ.2 Still, the notion may be judged to be scientifically important, governed by some principle or conjunction of principles, even if the notion is crucially vague, and typically will seem to invoke unwanted entities or entities one considers superfluous. The task then is to replace the notion, validating the principles in terms of entities which already have the seal of scientific approval.

Quine is, however, much more inclined than Carnap to stress the commonality between ordinary language and formal language. Quine avows 'my attitude toward "formal" languages is very different from Carnap's. Serious artificial notations, e.g. in mathematics or in your logic or mine, I consider supplementary but integral parts of natural language' (Quine's letter of 1943 to Church; Quine and Carnap 1991). He conceives the latter as an 'outgrowth' of the former, achieved by paraphrase of troublesome parts of the latter, not generally by outright replacement of whole languages, as Carnap favours. Carnap's task rather was to construct formal languages from the ground up, as in Languages I and II of *The Logical Syntax of Language*. This is indeed a symptom of a deep disagreement between the two, yet for these purposes we shall soon find their two paths reuniting at a still deeper level.3

---

2 In *Logical Foundations of Probability*, Carnap gives a good example: is the workaday concept of probability 'degree of confirmation' or 'frequency'? There is no real issue here, in Carnap's view (23–29).

3 Quine: 'In later years [Carnap's] views went on evolving and so did mine, in divergent ways. But *even where we disagreed he was still setting the theme*; the line of my thought was largely determined by problems that I felt his position presented' (Quine, 1970a, 41; emphasis added).

Although he formulated it comparatively late (1950), Carnap's distinction between 'external questions' and 'internal questions' can retrospectively be seen as central to his view from around 1931 onwards. An internal question is one posed within a 'linguistic framework' – an artificial language – and admits of an answer according to two types of general principles operating within that language: (1) 'transformation rules', including primitive sentences (axioms, L-rules of 1934/1937, §§10–14, 30–3) and rules specifying a logical consequence relation; and (2) – assuming it is a language set up for empirical enquiry – 'concrete sentences' (or 'protocol sentences') registering empirical data, plus 'theoretical sentences' (P-rules of 1934/1937, §51). The result is that the truth-value of an internal sentence depends on certain observations irrespective of whether it is concrete or theoretical, unless the sentence is analytic (L-true) or contradictory (L-false). By contrast, external questions are posed from outside the linguistic framework or artificial language. Such a question does not have a fully objective and precise answer, does not have a strictly rational answer, and indeed cannot genuinely have a true or correct answer. What it does have is a practical answer, an answer which involves the aims and desires of the investigators. For example, the ontological question 'Do numbers really exist?' cannot have the status that the metaphysician envisages. As an internal question, it admits of a trivial, deflationary answer, for either an artificial language will or will not have number terms and the answer will be analytic. As an external question, the question is misleading; it is more perspicuously posed as something like 'Is this language suitable for conducting mathematics?' Not only will the answer come in terms of less and more, it cannot strictly be assigned an objective status.

The cash value of the internal–external distinction of 1950 was prefigured in the so-called Principle of Tolerance, which appeared in *The Logical Syntax of Language*:

It is not our business to set up prohibitions, but to arrive at conventions … In logic there are no morals. Everyone is at liberty to build up his own logic, i.e. his own 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 syntactical rules instead of philosophical arguments. (Carnap 1934/1937, 51–2).

The basis for the Principle is, however, made more precise and explicit by the internal–external distinction. According to that distinction, because logical relations do not determinately bite except within an artificial language, one cannot subject a person to purely rational criticism for setting up whatever artificial language the person likes; one should be tolerant of such people, however little one may share their interests.

Carnap invokes the internal–external distinction in connection with ontology, but the idea can be seen to apply more generally to logic, and in particular to analyticity. For the distinction has a corresponding deflationary effect on questions of analyticity: just as answers to questions of ontology are settled by the choice of language, so it is with questions of analyticity. Carnap takes analytic truths primarily as being not discovered but stipulated, as being made true by their being decreed as such in the very setting up of the language (see Quine and Carnap 1991, 427f). Although he does maintain that the idea makes sense with respect to sentences of ordinary language, there is no precise fact about

which sentences of ordinary language have this status.4 '[T]he concept of analyticity has an exact definition only in the case of a language system, namely a system of semantical rules, not in the case of an ordinary language, because in the latter the words have no clearly defined meaning', he writes (1952/1990, 427). The ordinary concept is 'vague and ambiguous, and basically incomprehensible' (1963a, 919; see also 1956 234, 240, 241; and Quine and Carnap 1991, 430). Thus, it is ultimately futile to debate the question of the analyticity of a sentence of ordinary language, even if there are manifestly central examples.5

Carnap tirelessly maintained and defended the analytic–synthetic distinction despite Quine's qualms over it, which culminated in Quine's having explicitly rejected it in 'Two Dogmas of Empiricism'. Quine in effect asks: when you, Carnap, are stipulating a certain class of sentences as 'analytic', what exactly is it that you are stipulating? Quine thought of Carnap's seeming failure to give a satisfactory answer to his question as showing a fundamental weakness in his view concerning the precise role of artificial languages. It is not satisfactory to explicate the notion of analyticity simply by stipulating an extension as Carnap suggested; that doesn't answer the question of what it is that is being stipulated.

---

4 For Carnap the discipline called 'descriptive' or 'empirical' semantics is well grounded. But Carnap holds that the concepts of that discipline will often have a region between 'applies' and 'the negation applies'; hence the need for artificial languages.

5 Sam Hillier writes: 'In "Testability and Meaning", Carnap also makes it clear that his syntactical stipulations are intended to function as models of the language of science … he remarks that "there is no sharp line between observable and non-observable predicates … for the sake of simplicity we will here draw a sharp distinction […]" (1936/7,455) … The same holds true for his discussion of the analytic/synthetic distinction' (2010).

Yet in later work – in *Roots of Reference* (1974a) and 'Two Dogmas in Retrospect' (1991) – Quine accepted a version of the analytic–synthetic distinction. This later definition locates it squarely in the realm of the factual, indeed of the behavioural. It comes in two stages: (i) a sentence is analytic for a person if and only if the person 'learned the truth of the sentence by learning to use one or more of its words' (1991, 395–6); and (ii) a sentence is simply analytic if and only if it is analytic for those who are competent with the relevant terms.6 This is close to what is nowadays called 'Epistemic Analyticity' – that a sentence is analytic if a linguistically competent person cannot understand the sentence without accepting it – the differences being that this notion of Epistemic Analyticity lacks the reference to 'learning' in Quine's definition, using instead the term 'understanding' roughly in place of 'learning', and that it lacks references to a particular speaker, as in Quine's definition (I set aside the thought that Epistemic Analyticity is more explicitly normative than Quine's idea).

Quine supposes further that analyticity is preserved by 'obvious' logical inference: a sentence that is reached in logically obvious steps from an analytic sentence is itself analytic (1974a, 78–80; 1991, 396). Granted that a complete set of axioms may be chosen which are analytic in his sense, he thus allows that logic is analytic. Peter Hylton puts this together with another point from 'Two Dogmas', of the falsity from Quine's point of view of the Principle of Tolerance, of the inextricability of the pragmatic from the logical: the

---

6 The 1974a treatment is not quite as forthright as this, and includes the word 'understanding' in the discussion; see 'Cognitive Meaning' (1979) for Quine's attempt to find a behavioural definition of the word.

rational justification for any sentence has both theoretic factors and pragmatic factors (2019, 14–16). The result is that decisions about whether to accept a sentence are not rendered moot by the analyticity of the sentence, which might rationally be discarded despite being analytic (as some think the principle of bivalence should be, in light of quantum mechanics). Thus, the status of a sentence as analytic is not 'epistemically significant'; it is only of academic interest that a given sentence is analytic. Quine was surely motivated, at least in part, by the desire to find a sense of 'analyticity' that would make something like the customary verbal distinctions, yet without disturbing his epistemological picture in any way.

Hylton's point also shows that despite this disagreement between Carnap and Quine, they remain largely in agreement on a key point about ordinary knowledge, concerning assertions of natural language. For both Carnap and Quine take a critical view of ordinary language; neither thinks that the analyticity of a sentence of ordinary language is of great pith and moment. Quine writes the following regarding ontology, but the point goes over easily into the status of unregimented ordinary language in its logical or 'ideological' aspect:

Begin Complex Extract

The common man's ontology is vague and untidy in two ways. It takes in many purported objects that are vaguely or inadequately defined ... [And] we cannot even tell in general which of these vague things to ascribe to a man's ontology ... where can we draw the line?

It is a wrong question; there is no line to draw. Bodies are assumed,

yes; they are the things, first and foremost. Beyond them there is a succession

of dwindling analogies ...

My point is not that ordinary language is slipshod, slipshod though it be. We must

recognize this grading off for what it is, and recognize that a fenced ontology is just not

implicit in ordinary language. The idea of a boundary between being and nonbeing is a

philosophical idea, an idea of technical science in a broad sense. Scientists and philosophers

seek a comprehensive system of the world, and one that is oriented to reference even more

squarely and utterly than ordinary language. Ontological concern is not a correction of a lay

thought and practice; it is foreign to the lay culture, though an outgrowth of it. (1981d, 9)

End Complex Extract

Many ordinary expressions are simply not suitable for the expression of 'technical science

in a broad sense'. This implies than many questions posed in ordinary language, many

curiosities, should be rejected as naïve or confused by the light of science.

Carnap, for his part, took Heidegger to task for being unaware of the advances in

quantification theory made by Frege and others (1932/1959, 69). Heidegger uses, for

example, the term 'nothing' as if it were name of an entity. But according to modern

logicians, 'nothing' is best construed as a quantifier, not as a referring term; there isn't any

object into whose nature one might inquire (on the standard 'syncategorematic'

explanation of quantifiers). Sentences such as Heidegger's 'What about this Nothing?' are

such that in a logically well-behaved language, no such sentences 'can even be constructed'

(70). Such are the traps of using ancestral lore. The same goes for what he terms

'metaphysical' questions in general; they are, in general, not strictly well taken if posed in terms of the unreconstructed vernacular.

Something similar goes for Sorites-inducing expressions such as 'heap', 'big', or 'dry': though such terms are fine for everyday use, Sorites-type puzzles arise for them (as well as Ship-of-Theseus puzzles for artefactual terms – not to mention the complexities of 'person', where many urgent and often competing interests converge). It is not by chance that these kinds of paradox are seldom felt in real life; for where they do arise, folk may adjust their usage of the relevant term to avoid the paradox or uncertainty. For purposes of science (but also sometimes for other purposes, such as legal purposes), the way to escape the miasma is to replace such terms, to explicate them, or simply not to use them – for example, to dispense with the monadic term 'big' and make do with the relational term 'bigger than'. Thus, Carnap thinks it's often useless work to *analyse* existing concepts or things, and instead recommends 'rational reconstruction'.

13.2   Williamson

Since I am concerned only with points of contrast between his view and the shared regions I have outlined between Carnap's and Quine's, there is much in Williamson's view that I will not discuss. For example, I am not going to discuss at any length the role (or the nature) of modal concepts – so central to Williamson's philosophy and seemingly discounted by Quine's philosophy if not by Carnap's. I shall simply pick out certain interesting parts of *The Philosophy of Philosophy* which express his view in a revealing way. It will indeed be a case of quoting him 'out of context', but I hope not unfairly. I'll

concentrate first on Williamson's denial that philosophy is primarily a linguistic activity, and second on his related estimate of philosophical thought experiments.

13.2.1 Science, Semantics, and the Linguistic Turn

Williamson inveighs against 'philosophical exceptionalism', apparently approving of Quine's disavowal of 'first philosophy' in his opposition to the idea of 'metaphilosophy' (2007, Preface). In his opinion, the methods of philosophy are not categorically different in kind from those of science, even if it goes in much more for the armchair and much less for the laboratory. And like Carnap and Quine, Williamson discounts the purported special relevance of 'intuitions' to philosophy. They are just judgements, which can be repudiated like any other judgement by evidence or argument.

For Williamson, however, there are distinctively philosophical questions which are largely left unaddressed by scientists of the lab-coat variety. Partly these are metaphysical questions – questions of essence, of necessity, of the nature of objects and properties, of ontology and existence, and so on. And partly they are epistemological questions – the nature and the extent of knowledge, belief, and justification, the role of meaning and concepts in knowledge, and so on. For Williamson, it is an extraordinary rich field of problems calling for the sharpest logical implements and minds to go with them to make genuine progress.

Carnap and Quine would squirm at this. In their estimation, these words – 'essence', 'property', 'meaning', 'belief', 'knowledge', 'concept', and so on – are just words of ordinary

usage.7 Asking point-blank questions of the form 'What is the nature of …?' may mislead one into the mere assumption that science should retain the given category, or that there must be interesting *de re* modal truths about the category that require special philosophical techniques to bring out. On the contrary, descriptions of things in these terms are often imprecise or vague, relative to our shifting and unpredictable interests, and therefore not suited for science (the recent run of results investigating the 'method of cases' in 'experimental philosophy' comport with this; Stich and Tobia 2018). Of the term 'knowledge', Quine avers that what the long-running industry begun by Gettier has shown is that 'knowledge is a bad job'; that we should simply 'make do with its separate ingredients' (1989, 109; see also 2008a, 176, 249, 3228). In normal circumstances, the inchoate theory surrounding the word makes for a well-established and profitable use, but

---

7 In his response to Schuldenfrei, Quine writes: 'My position is that the notions of thought and belief are very worthy objects of philosophical and scientific clarification and analysis, and that they are in equal measure very ill suited for use as instruments of philosophical and scientific clarification and analysis' (Quine 1981b, 185).

8 Quine: The '[d]efinition … of "knowledge" is in trouble since Gettier's challenge of the definition of knowledge as true and warranted belief' ('The Innate Foundational Endowments', 2008a, 176); 'Knowledge, nearly enough, is true belief on strong evidence. How strong? There is no significant cut off point. "Know" is like "big": useful and unobjectionable in the vernacular where we acquiesce in vagueness, but unsuited to technical use because of lacking a precise boundary. Epistemology, or the theory of knowledge, blushes for its name' (2008a, 322).

in abnormal circumstances it may fail, like a Volkswagen Beetle being asked to climb a steep and muddy slope off road. Such terms arose in the context of human beings with specific aims, such as cooperation and solving workaday practical problems, ones which seldom call for scientific precision.

Let us now compare Williamson's attitude towards analyticity with Quine's and then with Carnap's. Williamson finds that neither analyticity in the 'epistemological' sense nor analyticity in the 'metaphysical' sense can make good on the traditional idea that analytic truths are categorically unlike other truths, insubstantial or 'trifling' in Locke's sense. '[P]hilosophical truths are analytic at most in senses too weak [not sufficiently distinctive] to be of much explanatory value or to justify conceiving contemporary philosophy in terms of a linguistic or conceptual turn' (53), he concludes. But this also is true of Quinean analyticity, which as we have seen includes logic and perhaps a sizeable proportion of mathematics. As recently noted, Quine recognises a sense of 'analyticity' which is very close to the epistemic sense, but he does not claim that that sense of analyticity is therefore epistemologically significant. Quine therefore more or less agrees with Williamson on this point.

The situation is more delicate with respect to Carnapian analyticity. Carnap, as explained above, does not think that questions of whether or not a sentence of ordinary language is analytic are worth agonising about. Not only can one quibble until the cows come home about 'whatever has shape has size', 'force = mass times acceleration', or even 'bachelors are unmarried', philosophical semantics since the 1970s have made it seem that 'dogs are animals' and the like – which many people had assumed to be analytic – are *not* true-by-meaning, so not analytic, for a term such as 'dog' is held to be a natural kind term,

and it is conceivable, though unlikely, that we are mistaken about the nature of their referents. For Carnap these cases cry out for explication. In setting up an artificial language, these are all questions that one is free to make decisions about. Thus, if what is in question is Carnap's attitude towards the concept of analyticity with respect to natural language, then he assuredly does not propose to conduct epistemology armed with that concept. If on the contrary what is in question is Carnap's explicated, artificial-language version of analyticity, then he certainly *does* think that the concept is epistemologically significant, but it is unclear that Williamson has a criticism of Carnap's plan for artificial languages, even if the plan were judged not to be, as Carnap believed, philosophically central.

So analyticity is not the place to locate a substantial difference between Quine and Carnap, on the one hand, and Williamson, on the other. Nevertheless, they do substantially differ from Williamson over empirical semantics, with extensive repercussions for other topics. For Williamson's attitude towards certain other key words of semantics itself is apt to be puzzling from the Quine–Carnap explicative point of view. Of the term 'synonymous', he writes:

> Although [Quine] may succeed in showing that "analytic" is caught in a circle
> with other semantic terms, such as "synonymous," he does not adequately
> motivate his jump from that point to the conclusion that the terms in the
> circle all lack scientific respectability, as opposed to the contrary conclusion
> that they all have it. Given any science, someone may insist that it define its
> terms, and the terms used to define them, and so on until it is driven round in
> a circle. By itself, that hardly demonstrates the illegitimacy of the science.

Every discipline must use undefined terms somewhere or other ... After all,

semantics is now a thriving branch of empirical linguistics. (50)

And: 'Pairs such as "furze" and "gorse" are pre-theoretically plausible cases of synonymous

expressions that speakers can understand in the ordinary way without being in a position

to know them to be synonymous' (67).

We can grant that these are 'pre-theoretically plausible cases of synonymous

expressions', but are they *really* synonymous? The uncertainty may be with the term

'synonymous'. The common-sense use of the term is interesting, but a well-grounded

verdict depends on how the word comes to be satisfactorily defined or whether indeed

such a word is retained in a mature theory. Perhaps 'semantics is now a thriving branch of

empirical linguistics', but the standing of the field does not depend on whether there is a

well-grounded and fruitful definition of 'synonymy' to be had, as Quine repeatedly

emphasised.

The circle of semantical terms of which Williamson speaks is by no means

sufficiently stable to be generally accepted as is by working scientists. One can imagine

various senses of what it means to be 'thriving', but at any rate the sub-region of linguistics

called 'semantics' is more towards the ubiquitous end of the rough scale of sciences

described at the outset of Section 13.1. There are just too many views about what are its

proper methods, aims, and concepts to declare the existence of anything like a dominant

paradigm – there is teleosemantics, truth-maker semantics, inferentialism, conceptual role

semantics, situation semantics, game-theoretic semantics, dynamic semantics, intention-

based semantics, semantical minimalism, expressivist semantics, radical contextualism,

and others. And as Chomsky himself warns, many of these might well incorporate philosophical speculation masquerading as science (2000, 40–45).

Williamson says, 'When clarification is needed in some specific respect, it can be achieved by stipulation or otherwise, as elsewhere in science' (51), but he does not actually propose what for Carnap and Quine would be a viable explication of the term 'synonymy' (as applied to natural language). This verges on the view that the ordinary concept of synonymy is sufficiently shipshape just as it stands. But from Carnap's or Quine's point of view, it is unscientific to presuppose a pre-existing distinction between problems of meaning and problems of fact, implying the possibility of a sharp and well-grounded notion of synonymy. It is not mere scepticism to suspect that the term is no more than a practical device of ordinary language, not a term to rest with in science itself.

Another telling aspect of Williamson's view emerges in connection with the problem of vagueness:

> Some philosophers … regard the attempt to give a systematic statement of the truth conditions of English sentences in terms of the meanings of their constituents as vain. For them, the formalization of "Mars was always either dry or not dry" as $(\forall t)\,(\text{Dry}(m, t) \vee \neg\text{Dry}(m, t))$ is already a mistake. This attitude suggests a premature and slightly facile pessimism. No doubt formal semantics has not described any natural language with perfect accuracy; what has not been made plausible is that it provides no deep insights into natural languages. (37)

A reasonable move from the Carnap–Quine point of view would indeed be to 'deny the relevance of formal semantic theories to vague natural languages' (37; not to dispute

Williamson's criticisms of the other formal approaches to vagueness, such as supervaluation or degrees of truth). Lay people may well regard the setting up of premises for a Sorites-style reduction to absurdity as setting them up for a trick, a trick moreover that does not strike them as profound, the reaction being a shrug of the shoulders or an eye roll. When confronted explicitly with the induction premise – for example, that if x is not dry then x with a single molecule of liquid removed is not dry – they will perhaps agree with Peter Strawson that 'ordinary language has no exact logic'(1950, 344), even though it may be unclear to them exactly what they are denying. Words such as 'dry' are only as precise as ordinary purposes demand. For extraordinary purposes, for scientific purposes, we can replace it with 'dryer than' (as Carnap replaces 'warmth' by 'temperature'; Carnap 1950/1956, 29). One simply refuses to play the trickster's game.

Neither Carnap nor Quine took the linguistic turn in the sense Williamson intends, yet their views about semantics and science render their general philosophical views manifestly antithetical to Williamson's. Carnap's interest in rational reconstruction and Quine's in regimentation place the language of science at the centre of philosophical concern, but for this purpose they avoid using the intuitive semantical notions in effect favoured by Williamson and others to describe natural or ordinary language. They use only at most their corresponding explicated versions, and strictly only with respect to artificial or regimented languages.9 Neither Carnap nor Quine thinks that there is a special realm of philosophical truth to be divined by reflecting on our ongoing and sometimes unrestrained

---

9 Tarski is the source of both the later Carnap's use of semantical concepts and of the post-youth Quine's. But Tarski is clear that the notions he employs are explicated versions of the (inconsistent) ordinary notions.

thought and talk about the world, as opposed to clarifying, streamlining, and marginally adding to our best scientific theories of the world.

13.2.2 Thought Experiments and Science

I now shift to my other point of comparison, that of their respective attitudes towards 'thought experiments'. I shall argue that – from the Carnap–Quine point of view – Williamson is wrong to deny that there is a fundamental difference between thought experiments in philosophy and those in the empirical sciences. I shall begin with some points about the subject in general.10

There is no doubt that the rhetorical, heuristic, and pedagogical value of thought experiments is often striking – no less than Galileo, Newton, and Einstein provide examples. They can serve to dramatise aspects of the world, to isolate and highlight them by descending from the general to the particular level, even if the particulars are hypothetical and often ideal. It seems natural to divide thought experiments into two kinds, which naïvely we might call 'empirical' and 'philosophical' (or 'conceptual'); Williamson questions that division, at least that there is an epistemologically relevant difference.

A typical laboratory experiment involves a theory that implies that certain things will happen under certain circumstances. One brings about those circumstances and checks for those things. The motive is that one is not sufficiently certain what result will be

---

10 Carnap doesn't explicitly address thought experiments I far as I'm aware; for Quine, see (1981b, 185; 2008a, 284, 341, 345, 437, 448).

obtained, or that one needs the result to persuade others, or both; for these reasons one must get out of one's armchair – a mere thought experiment is unavailing.

Where one does in some sense already 'possess' the result, a thought experiment seems to be both appropriate and valuable. Perhaps the result is merely implicit in principles which one accepts, or one simply has not thought through the relevant application to reality of the principles. For example, Galileo considered Aristotle's claim that the speed at which a body falls is proportional to its weight, and contested it by reasoning along the following lines: if heavy things fall faster than light things, then if a small lead shot were tied to a cannonball with a length of twine and they were dropped, the cannonball's flight must be slowed as it pulls the small lead shot downwards by the twine; but on the other hand since the cannonball-and-lead-shot-assembly is heavier than the cannonball itself, it must fall faster than the cannonball falls alone; the contradiction shows the hypothesis to be incorrect, so it is not the case that heavy things fall faster than light things. We want to say that this sort of thought experiment is one of physics – is 'empirical' – despite its having been conducted from the armchair.

A well-known thought experiment which we are naïvely inclined to say is of another type, the philosophical type, was described by Gettier and is stressed by Williamson, a variant of which might run as follows: if knowledge is justified true belief, then someone who justifiably believes that there is a sheep in the field – but only because they see in the field a dog disguised as a sheep – when in fact there is sheep in the field which, however, the person does not see, would have a justified true belief that there is a sheep in the field, and therefore would know that there a sheep in the field; but such a person would not

know that there is sheep in the field, so it is not the case that knowledge is justified true belief.

What is the distinction between these kinds of thought experiment? One might think it crucial that the first experiment presumably involves the physical-spatial imagination, whereas the second one does not. Williamson, however, is sceptical that the second one really lacks this feature; discussing knowledge-involving scenarios like the second thought experiment above, he writes:

individual differences in the skill with which concepts are applied depend constitutively, not just causally, on past experience …

In a similar way, past experiences of spatial and temporal properties may play a role in skilful mathematical "intuition" that is not directly evidential but far exceeds what is needed to acquire the relevant mathematical concepts. (168–9)

Thus, Williamson denies that the distinction between a priori and a posteriori knowledge captures a distinction between our two examples in 'a deeper theoretical analysis' (169).

All the same, even if it is not precisely an exercise of the a priori faculties, one might think of Gettier as making a point about the *concept* of knowledge, namely that the concept is not that of justified true belief, whereas Galileo's lesson affects the *theory* of gravity, without disturbing the concept of gravity itself. Call this reaction (a). One can also feel a tug

in the opposite direction, either supposing (b) that despite appearances, and as in the case of Galileo and the concept of gravity, Gettier's conclusion does not teach us anything about the concept of knowledge so much as having implications for theories in which the concept figures; or (c) that in both cases the lesson affects the concepts themselves, for the concepts depend on the theories in which they figure.

Williamson chooses (b). Together with the theory and conditions, in both cases the result is arrived at via some subtle use of the imagination and perhaps knowledge of other matters. He argues impressively against (a) – in sum, the attempt to make out exactly why one experiment is conceptual and the other is not simply runs into the weeds (I will not go into the detail of the discussion); (c) is ruled out for similar reasons, and a further reason against it comes by considering communication (in their 1992 Fodor and LePore make a similar point):

> Such variation ... [differing results from thought experiments] may result from cross-cultural variation in the meaning of "know" or other epistemological terms, but it need not. It may occur between sub-communities of English speakers who all use the words as part of a single common vocabulary, but disagree in their applications of them, just as different communities may disagree in their applications of the word "justice" while still using it with a single shared meaning. (190)

The conclusion is that the two thought experiments are not fundamentally different.

I'll restrict myself to Carnap in this paragraph. It is doubtful that Carnap would have set much store by the idea that despite disagreement over cases, we communicate with 'the word "justice" while still using it with a single shared meaning'. But more generally I think

that Carnap's view would be that the distinctions on which Williamson's discussion

depends – especially that between the conceptual and the factual as applied to ordinary if

learned speech but in other applications as well – are not sufficiently precise to ground any

such sweeping conclusions. If the distinction coincides with the a priori–a posteriori

distinction, then the a priori does not have anything like a determinate extension if the field

is ordinary assertions, sentences, or beliefs. One has to explicate it in terms of analyticity,

and as we've seen analyticity is only genuinely fruitful with respect to artificial languages,

in Carnap's view. A question such as 'Is the reasoning [in a given thought experiment] a

priori?' is often a 'pseudo-question' without cognitive content or at any rate with

insufficient cognitive content, leading only to 'sterile debates'.

In Gettier cases, Putnam's twin-earth cases, and so on, there is scant possibility of

the thought experiment conflicting with an actual experiment. Perhaps for that very reason,

one supposes that these thought experiments have a type of authoritative standing that the

physical ones lack. Yet the results are disputable. Williamson writes:

> [N]ative English speakers sometimes dispute the Gettier verdict ... In doing
>
> so, they show poor epistemological judgment but not linguistic incompetence
>
> ... We assent to (3*) ['If there were an instance of the Gettier case, it would be
>
> an instance of justified true belief without knowledge'] on the basis of an
>
> offline application of our ability to classify people around us as knowing
>
> various truths or as ignorant of them ... That classificatory ability goes far
>
> beyond mere linguistic understanding of "know". (188)

This reference to a fine-honed skill, this 'classificatory ability', is important to Williamson's

case. He goes on: '[S]omeone with a distorted epistemological outlook may reject (3*) yet

still possess the relevant concepts: they genuinely believe that the subject of the Gettier case would not have justified true belief without knowledge' (189).

And further:

Much of the evidence for cross-cultural variation in judgments on thought experiments concerns verdicts by people without philosophical training. Yet philosophy students have to learn how to apply general concepts to specific examples with careful attention to the relevant subtleties, just as law students have to learn how to analyse hypothetical cases. Levels of disagreement over thought experiments seem to be significantly lower among fully trained philosophers than among novices. (191)

The reliance on this 'classificatory ability', this skill for 'applying epistemological concepts', indeed likens the epistemologist to a legal judge if not a seer, called upon to judge even in the absence of intersubjectively decisive reasons. Yet a principal reason that judges are essential to the legal system, and cannot typically be replaced by computers, is that the fact that *they judged* as they did is *itself* a central fact of interest. A proper scientist, on the contrary, tries in principle to eliminate any reliance on the judgements and skills of human beings, and in principle reckons a failure to do this a failure of the experimental design. To claim that sort of status for 'applying epistemological concepts' would require the formulation of criteria for their application – rules or methods of measurement that a computer could follow – which is exactly what Williamson's reliance on this 'classificatory ability', the idea of knowledge as primitive, rules out.

Galileo could rely on his thought experiment because he was confident of what the relevant observations of an actual experiment would be: the objects, aside from

considerations of air friction, would fall at the same rate. It is not a logical contradiction that drives the thought experiment; perhaps the cannonball-and-lead-shot system behaves as a single object the more rigid the connection between them, and only minimally when held together with a piece of twine. If there were any iota of doubt, then in principle he could have rigged up an actual experiment and measured. There is no such contrast in the case of the Gettier thought experiment – there is nothing to measure. The same point can be made by observing that widespread and well-informed disagreement over the result of a Gettier thought experiment needn't be rationally resolvable, whereas widespread disagreement over a Galileo-style thought experiment would be rationally resolvable by conducting an actual experiment. True, it is sometimes practically impossible to perform an actual experiment of the requisite sort – Einstein's scenario of the observations one would make travelling at the speed of light might be cited – but still the facts being pointed out have nothing to do with the observer.

Carnap would not want to join Williamson in his evident confidence in making the distinction between cases of factual disagreement and cases of linguistic or conceptual disagreement. He would take note of the lack of unanimity of response, the sheer range and diversity of theories of knowledge, the perplexing chaos of competing accounts of what it is to know, and take this as a sign of a lacuna in our theoretical landscape. *If* there is some theoretical need to do so, then one can propose explications incorporating, for example, the causal theory of knowledge, or reliablism, or something else. From Carnap's point of view, these are simply alternative explications of the familiar word 'knowledge'. Perhaps they can *all* be accepted, with differing subscripts on the term 'knowledge'. But there needn't be

a fact which can be discovered that such a response brings to light, one which was already there, hidden somehow from sight.

The situation is similar for such questions as 'Do cats understand?', 'Do chess-robots think?', or 'Why do things happen?', to borrow examples from Chomsky (2000, 19, 45, 114). The feeling that such questions are well formed and have answers just as they stand is sometimes undeniable – and is encouraged no doubt by the thought, so popular since Putnam and Kripke discussed the topic, that ordinary words stand for natural kinds – but one can join with Chomsky (and Wittgenstein) in suggesting that these, like the question of whether submarines swim or whether Pluto is a planet, are at best matters of decision, not fact (even if some decisions might fit emerging science better than others, such as the recent decision to exclude Pluto from planethood).

I have so far discussed Williamson's view of thought experiments without discussing its most significant feature, that at their centre is counterfactual reasoning (involving conditionals of the form 'If P were so, then Q would be'; 141ff). Space permits to make only a point concerning that arch anti-modalist Quine.

Quine observed that many counterfactual conditionals have a hidden egocentric element, as in the pair of examples that he cites from Nelson Goodman: 'If Caesar were in command, he would use the atom bomb; If Caesar were in command, he would use catapults' (1960, 222). It is only a specific context, in particular a speaker intention, that determines which of these sentences is appropriate (a supposition that has been deemed harmless, most famously by Lewis, at (1974a), 66–67). Nevertheless, for Quine, counterfactual conditionals do have a vital scientific use in connection with dispositions – at least those expressing dispositions narrowly so-called are relatively immune to changes

in context, and are often indispensable to science. Solubility in water, for example, can be expressed as 'if x were immersed in water, x would dissolve'. Such forms are not extensional, and therefore according to Quine must be regarded as second best, not to be used in a reality-limning, regimented statement of a theory (1960, 222–5). But a disposition, for Quine, is in turn identical with an underlying structural trait which can, in principle, be expressed in extensional language as a universal generalisation – even if we cannot, given the present state of knowledge, specify which (1974a, 8–15; 1995, 21; in the case of solubility, the relevant generalisation pertains to chemical structure).

The key is the role of universal generalisations, not exactly dispositions as such. The idea is that legitimate science has no 'ungrounded' true counterfactual conditionals, none that cannot in principle be explained in terms of extensional, universal generalisations. True, this is still not as rich an apparatus as the full gamut of possibility made available on Williamson's conception, but I put it that it is sufficient for such uses as Galileo's in his thought experiment.

13.3 Conclusion

Williamson is not alone in assuming that certain notions of everyday life – knowledge, belief, reference, assertion, causality, essence, possibility, conceivability, and so on – are fit notions in terms of which to frame philosophical questions and answers. Mustering a genuinely persuasive argument against Williamson's instinct would be extremely difficult, and would amount to something on the order of Kant's *Critique of Pure Reason* or Wittgenstein's *Investigations*, revised and augmented. But what I want to point

out, and emphasise, is that neither Quine nor Carnap make such an assumption. In fact, their not making it is the motive arguably most distinctive and central to their philosophies. In a word, they regard almost every such notion as standing at least in need of what they call explication, and many of them, for Quine most visibly with his penchant for desert landscapes but Carnap too, represent curiosities which are perhaps natural but have no place in science.11