

Operational Coherence as the Source of Truth¹

(version for publication, 14 February 2017)

Hasok Chang

Department of History and Philosophy of Science, University of Cambridge

Abstract

In this paper I seek to defend an epistemology that does not confine itself to the knowledge of propositions. The first section motivates this move, especially from the standpoint of the philosophy of science. The second section presents the notion of operational coherence as the key to understanding how knowledge resides in activities. The third section presents a proposal for making sense of truth on the basis of operational coherence. The final section briefly reconsiders the relation between knowledge-as-ability and knowledge-as-information.

1. Knowledge beyond propositions

The overall direction of this paper is to move beyond the propositional conception of knowledge. What I mean by that phrase is the widespread notion that knowledge (or at least the kind of knowledge that deserves the attention of epistemologists) consists in possessing the right sort of belief in the right sort of propositions.² Without denying the importance of propositional knowledge, I want to pay attention to other aspects of what we commonly call knowledge, which cannot comfortably be fitted into a propositional framework. I will not pretend that the move I am making is a novel one, especially in this journal.

¹ I would like to thank various members of the audience at the Aristotelian Society meeting of 9 January 2017 for helpful objections and suggestions, as well as encouragement. I also thank Hannah Carnegie, Guy Longworth, Tim Crane and other officers of the Society for their help in arranging the presentation and publication of this paper.

² Any doubt that such a view is taken as common sense among Anglophone analytic philosophers would be quickly dispelled by a cursory look at basic philosophical reference works such as the encyclopedia articles by Ichikawa and Steup (2016) and Truncellito (2007).

After all, it was in a Presidential Address to the Aristotelian Society that Gilbert Ryle (1900–1976) presented his distinction between “knowing that” and “knowing how” and stressed the independence and importance of the latter (Ryle 1945/46; also Ryle 1949, chapter 2); that work was discussed relatively recently in another Presidential Address by Paul Snowdon (2004), though in a critical light. And “knowing how” and “knowing that” are not all the only kinds of knowledge, either, if we pursue a line of thought that Snowdon (2004, p. 5) mentions and discards. To every kind of question “Do you know...” corresponds a different kind of knowledge. So there is not only knowing-how and knowing-that, but also knowing-why (having a causal or intentional explanation for something), knowing-what/who (recognition of another being), knowing someone or something (acquaintance), knowing-what-it’s-like (empathetic understanding), and more types besides.

It is not my intention in this paper to enter deeply into the ongoing debate as to whether “knowing how” is really only a species of “knowing that”. For now, I only beg acceptance that it is an important and meaningful thing to say that “I know how to do X” (I imagine that there are roughly equivalent expressions in most languages), and that it should not be taken for granted that knowledge as the ability to do things is inferior or subordinate to knowledge as information storage and retrieval. I believe that distinguishing between knowledge-as-ability and knowledge-as-information is a cogent way of expressing what Ryle was trying to get at by his distinction between knowledge-how and knowledge-that. It also allows us to bypass a main line of critique of Ryle, articulated by Jason Stanley and Timothy Williamson (2001), which rests on denying that knowledge-how is about ability. Knowledge-how as Stanley and Williamson take it, which consists in descriptions of how something is done, may well be a species of knowledge-that. What I want to talk about, which I believe Ryle was also concerned with, is an agent’s mental and physical functioning in bringing that something about. In this regard Snowdon’s (2004, p. 7) distinction between “knowing how” and “knowing how to” is a helpful first step, though he then rather unhelpfully lumps “how to” together with other infinitive constructions such as “when to” and “where to”.

Somehow the direction of epistemology pioneered by Ryle and his fellow-travellers such as J. L. Austin (1911–1960), not to mention the later Wittgenstein, has become sidelined in the mainstream of analytic philosophy, a tradition which they had themselves done much to establish.³ There are various good reasons for a return to this line of work, which many philosophers have attempted. My own particular motivations are rooted in the needs of the philosophy of science, rather than epistemology or other branches of philosophy. In this context it is interesting to note that Alfred Senior, the first Secretary of the Aristotelian Society and the man who initially conceived the idea of the society, was a scientist — by profession an analytical chemist, rather than an analytical philosopher. As H. Wildon Carr put it in his 50-year retrospective on the Society in 1929: “The ideal of the Aristotelian Society is the study of philosophy not as an academical subject but as the story of human thinking.”⁴ The “story of human thinking” in the modern times must surely include the story of scientific knowledge. And my sense is that knowledge-as-ability is just as important in science as in everyday life, though Ryle’s and Austin’s examples tended to be taken from everyday life.

To get our intuitions going, let’s consider briefly some examples of the sort of things that we should want to know in science, in addition to cut-and-dried facts. Here again it is instructive to take heed of Ryle’s insight (1945/46, p. 15): “The advance of knowledge does not consist only in the accumulation of discovered truths, but also and chiefly in the cumulative mastery of methods.” In advancing scientific knowledge, we should want to know how to analyze a complex organic molecule to ascertain its molecular composition and structure. And we want to know how to learn and teach such a skill. We want to know not only the trajectory of a planet, but how to compute it, which involves knowing how to solve the equations of the basic physics involved; we also want to know how to come up with such equations in the first place. We want to know how to measure temperature and humidity, the rate of inflation, and the level of well-being of a population. We want to know how to synthesize a new

³ There were also important outliers active in the early and middle parts of the 20th century, including Michael Polanyi, Marjorie Grene and Percy Bridgman.

⁴ Carr (1928/29), p. 359. The statement about Senior occurs on p. 360.

pharmaceutical agent, and how to assess its clinical efficacy. We want to know how to make a superconductor that will operate above the temperature of liquid nitrogen. We want to know how to sequence a DNA molecule; how to run a Monte Carlo simulation of an experiment we can't carry out physically; how to model a complex situation as a causal graph; etc, etc. We should want to have an epistemology that can address these items of knowledge-as-ability directly, rather than skirting around them in an awkward and round-about way, treating them as the applications of propositions that we believe, or as inessential accompaniments to propositions.

Another source of motivation for getting away from the propositional conception of knowledge is the work of Thomas Kuhn and others who took a serious historical look at the development of scientific knowledge. Philosophers had usually conceived the task of judgement facing scientists as the problem of “*theory-choice*”, but Kuhn showed quite convincingly that scientists' choice at the most crucial moments in the history of science was between entire *paradigms*, rather than merely theories. What exactly Kuhn meant by “paradigm” was famously debatable, but at least it had to be admitted that a paradigm contained particular *methods* of work and *criteria* of judgements, as well as straightforward descriptive statements. Theory-choice came to be seen to be inextricably linked to, and essentially dependent on, choices concerning non-propositional aspects of science. This is at the heart of the Kuhnian incommensurability problem, more important than its semantic aspect. Whether or not one agrees with everything Kuhn says, it has to be admitted that the unit of analysis employed in philosophy of science must include something beyond propositions.

In previous publications I have proposed that scientific work (as well as non-scientific but knowledge-related aspects of living) can be analyzed in terms of “epistemic activities” and “systems of practice”, in conscious opposition to the more customary analysis of scientific knowledge as consisting of propositions (Chang 2011a, 2011b, 2012, 2014). I defined an *epistemic activity* as “a more-or-less coherent set of mental or physical operations that are intended to contribute to the production or improvement of knowledge in a particular way, in accordance with some discernible rules (though the rules may be unarticulated).” Epistemic activities normally do not, and should not, occur in

isolation. Rather, each one tends to be practiced in relation to others, constituting a whole system. A scientific *system of practice* is “formed by a coherent set of epistemic activities performed with a view to achieve certain aims” (Chang 2014, p. 72; Chang 2012, pp. 15–16). Let me illustrate briefly with an example: Antoine Lavoisier created a new system of chemistry whose main activities included making various chemical reactions including those involving gases, tracking chemical substances through weight-measurement, classifying compounds according to their compositions, and analyzing organic substances by combustion. The overall aims of this system included determining the composition of various substances, and explaining chemical reactions in terms of the composition of the substances.

The linchpin in this whole way of thinking turns out to be the notion of coherence, whose meaning I have left quite vague in previous publications. It is the main thing I want to elaborate on in the rest of this paper. Coherence as I intend it goes beyond consistency between propositions; rather, it consists in various actions coming together in an effective way towards the achievement of one’s aims. Coherence comes in degrees and different shapes, and it is necessarily a less precise concept than consistency, which comes well-defined through logical axioms. An important part of my proposal is to keep in mind the aims that scientists are trying to achieve in each and every situation. The presence and operation of an identifiable aim is what distinguishes actions from mere physical happenings involving human bodies, and it is also what places knowledge firmly in the realm of actions.

2. Correspondence vs. coherence

According to the standard propositional conception of knowledge, knowing something is a matter of mentally possessing propositions that correspond to the world. This picture embodies an ideal of knowledge focused on *correspondence*, which is impossible actually to approach. First of all, the idea is liable to be based on a category mistake. Otto Neurath put the point succinctly ([1931] 1983, p. 66): “*Statements are compared with statements, not with 'experiences', not with a 'world' nor with anything else.*” Hilary Putnam makes a complementary point (1995, p. 10): “To say that truth is ‘correspondence to

reality' is not false but *empty*, as long as nothing is said about what the 'correspondence' is. If the 'correspondence' is supposed to be utterly independent of the ways in which we confirm the assertions we make . . . then the 'correspondence' is an occult one, and our supposed grasp of it is also occult." Austin tried to discern a plain and untroublesome sense in which statements correspond to facts, but the nature of the correspondence remained unclear (unless one subscribed to the notion that the world actually *consisted* of facts). In the end he considered statement–fact correspondence a conventional pairing, not any kind of resemblance (Austin [1950] 1979, pp. 121–126). The Tarski disquotation scheme does not save us here: it makes sense as a matter of relationship between two languages (or language and meta-language), or tautologically as a relationship within one language; either way, it says nothing about how statement and fact (or world) might relate.

It must be admitted that we have no access to the “external” world, except through the statements that we regard as true. Seeking statement–world correspondence is not an operable move, unless it is a circular–tautological move: we say that we access the world through our possession of true statements, and that true statements are just those that give us access to the world.⁵ This inoperability of the propositional ideal of knowledge as correspondence gives rise to some well-known problems in the philosophy of science. At least since Kuhn’s *The Structure of Scientific Revolutions*, first published in 1962, philosophy of science has been plagued by doubts about the security of scientific knowledge — not global skeptical doubts, but practical doubts that manifest themselves concerning actual situations of scientific choice. While we uphold science as the best model of knowledge, we are also forced to admit, if we pay any attention to the history of science, that scientific change has shown no clear direction concerning the fundamental ontology of nature. Our optimism concerning scientific progress is dampened down by the “pessimistic induction” from the history of science, in which it is seen that nearly all

⁵ I believe that our notion of a statement or theory “corresponding to the world” is a metaphorical projection based on actual representational activities in which we make a depiction of something, and we have access to both the object and the depiction so that the correspondence between the two may be checked. See Chang (2016), pp. 109–111, for a preliminary exposition of this view.

previously trusted theories are later rejected (Laudan 1981). The pursuit of correspondence-truth leads us either to an epistemic dead end, or a comforting yet empty tautology.

In the alternative epistemological vision that I seek to promote, knowledge is closely related to our ability to perform successful activities. It is in order to facilitate the analysis of activities and their successes, that I introduce the notion of “coherence”. It will not do to say that an activity is successful because it somehow corresponds correctly to the world. An activity we perform resides *in* the world, the same world in which we live; they do not discernibly correspond to the “external world”, whatever one might mean by the latter phrase. A more productive perspective is to see that an activity works out because, roughly speaking, what goes into it all fits together nicely. It is important to note that coherence as I intend it is about the harmoniousness of actions, not primarily about the logical relationship between propositions. To mark that point clearly, I will use the phrase “operational coherence” whenever needed, and just “coherence” when the meaning should be clear enough from context.⁶

Somewhat more precisely: operational coherence is a harmonious fitting-together of elements and aspects of an activity, which is conducive to the successful achievement of the aims of that activity. (Note to anyone intending to comment on this paper: do not quote this formulation, as it is far from satisfactory until the terms occurring in it receive more clarification and refinement). Such coherence may consist in something as simple as the correct coordination of bodily movements needed in riding a bicycle, drinking a glass of water, or walking up the stairs (very difficult to achieve, as we have learned in contemporary robotics), or something as complex as what is involved in the integration of a range of material technologies and various abstract theories in the global positioning system (GPS). In puzzling out what coherence is, it might be helpful to think about what happens when it is lacking. If I try to drink water

⁶ In speaking of “operational” coherence, I am giving a conscious nod to Percy Bridgman’s advocacy of the operational point of view; see Chang (2009) for further details on my interpretation of Bridgman’s philosophy. As I will attempt to argue in later publications, I believe that my operational notion of “coherence” is implicit in John Dewey’s theory of knowledge.

by directing the glass to my nose, that is an incoherent activity. When we do not heed the sign that warns “mind your step”, that rare moment of stumble reminds us how carefully and how well we normally maintain the coherence of our bodily movements in everyday life without even thinking about it. Incoherence may be traceable to false beliefs (about where my mouth is, for example) or mutually contradictory beliefs, but ineptitude of belief is certainly not the only reason for incoherence. It could also be due to the lack of capability (starting with simple lack of muscular strength or failure of eyesight), the use of inappropriate materials, poor timing between different operations, and so on.

So much for intuitive illustrations: in order to make this notion of coherence workable, much more needs to be said about each part of the definition given above. First of all, what do I mean by the “elements and aspects” of activities? That was just loose talk. There are many ways of analyzing how activities are constituted, and I am not equipped to enter into a full study of the ontology of actions in this paper. For my present purposes at least, it makes sense to analyze an activity as made up of *operations*, as suggested above in my definition of epistemic activity. Take a very simple activity, for example match-lighting — so essential to the progress of chemistry, even physics, for many crucial decades! I am starting with something basic, since the analysis quickly gets very complicated when we consider even a slightly more complex activity, such as weighing-with-a-balance.⁷

Most people can probably bring up the memory of learning how to light a match, which actually takes a surprising degree of skill and coordination to do well. With one hand I hold the matchbox steady and firm, with the rough strip facing my other hand; with the other hand I hold the match tightly, just so; I pull the head of the matchstick across the rough strip on the box (no, no, the correct move is to push it), at an appropriate angle and at the right speed; I stop the movement of that hand once the flame comes on. These four operations have to come together well enough for the activity of match-lighting to be coherent. This

⁷ For a brief discussion of the latter activity, see Chang (2011b), p. 253. That example will also clarify that starting with the analysis of basic activities is not done with a reductionist implication. It is impossible to understand the simple activities without some sense of the more complex ones into which they fit, just as much as the converse.

is what was intended by “a harmonious fitting-together” in the rough-and-ready definition of operational coherence given above. It is difficult to be more precise in characterizing this quality of harmony in inter-operational interactions, or to reduce it to another, better-understood notion. We can go on listing synonyms: coordination, orchestration, concordance, back to coherence... It may be best to take “harmony” (or “harmonious”) as a primitive in its meaning, and verifiable in the end only through the achievement of the aim of the activity.

This brings me to the last part of the above definition of coherence: what do I mean when I say that coherence is “conducive to” the successful achievement of the aim of the activity? (Or rather, I should simply say “conducive to the achievement...” since “successful achievement” is redundant, or perhaps just “conducive to its success” since “success” can only mean “the achievement of its aim”.) To say “conducive to” actually fails to indicate precisely my thoughts on the relation between coherence and success, which will take some spelling out. On the face of it, there are two possibilities: coherence and success just mean the same thing; or, coherence is the cause of success. One strong reason against the former possibility is that a coherent activity may well fail, due to unforeseen circumstances (and an incoherent one may succeed occasionally by accident). I may do all my operations correctly in attempting to light a match, but meet a sudden gust of wind, a mischievous friend pouring a bucket of water all over me, or any number of other possible mishaps. We could say that the activity was not coherent if it did not take precautions against all of these possibilities, but that is not a palatable option. Not only does it make it quite impossible for anyone ever to engage in a coherent activity, but it also makes coherence forever undetermined, as it depends on whether or not the innumerable “other circumstances” will come to obtain. It is better to say that my match-lighting activity is coherent, but may occasionally be unsuccessful due to circumstances. Then we would be saying that a coherent activity is successful, *ceteris paribus*. Does that amount to saying that coherence is the cause, or rather a cause, of success? There is some sense in that formulation.

It is important to keep in mind that an “activity” is not a single act, not even an actual concrete happening. Rather, it is a conceptualized thing, a *type* of action characterized by a description, not simply referred to by ostension. Such

description is always abstract in the sense of not including all the features that real concrete entities (including actions) possess. Which features to include in our description of an activity is a conventional decision, and there is no uniquely right way to identify and classify activities out of the barely differentiated stream of actions that we continually engage in, alone and with each other. So, an activity, as such, is not precisely instantiated in our actual doings, and success can only be judged through how our actual doings work out. A coherent activity makes sense in the realm of abstraction, but whether its actual execution is successful depends on all sorts of conditions. This is responsible for the sense that coherence and success are not synonymous.

The following, then, is my considered definition of operational coherence: an activity is operationally coherent if and only if there is a harmonious relationship among the operations that constitute the activity; the concrete realization of a coherent activity is successful, *ceteris paribus*; the latter condition serves as an indirect criterion for the judgement of coherence.

I must stress that this is not what epistemologists usually mean by “coherence”, as operational coherence is irreducible to logical relations between statements. In the most simple-minded version of the coherence theory of truth, coherence is taken to mean mere logical consistency within a set of statements. This is nothing short of a philosophical disaster, an invitation to vicious circularity and the most problematic kind of relativism; it eliminates any inherent link between knowledge and reality. James O. Young (2015) notes that more plausible versions of the coherence theory take the coherence relation as “some form of entailment” or “mutual explanatory support between propositions.” A similar thought to Young’s latter formulation is expressed by Richard Foley (1998, p. 157): “Coherentists deny that any beliefs are self-justifying and propose instead that beliefs are justified in so far as they belong to a system of beliefs that are mutually supportive.” But the problems of circularity and relativism remain in the idea of propositions rendering one another true by mutual support without anything else to ground any of them.

In contrast, operational coherence cannot be achieved in an arbitrary fashion by decree, wishful thinking, or mere mutual agreement. On the contrary, in order to do things successfully in the world, we need to have an understanding

and mastery of our surroundings. It is operational coherence, not the mirage of correspondence, through which the mind-independent world is actually brought to bear on our knowledge. Operational coherence carries within it the constraint by nature. In fact, having cleared away the ungrammatical illusion of a direct correspondence between proposition and reality, we can see that operational coherence is the only way in which reality can enter our practices.⁸

3. Coherence and the truth of propositions⁹

Having spelled out the notion of operational coherence, I would now like to come to a more careful consideration of propositions and their truth. We need to ask: if coherence is the property of an activity, how does it relate to truth, which is the property of a statement or a proposition? This is a significant and difficult question, to which I will attempt an initial sketch of an answer here. I want to start by building productively on the “pragmatic theory of truth” commonly attributed to William James. This theory is widely regarded as absurd, which has also contributed to the unpopularity of pragmatism among tough-minded philosophers. Here is probably the most notorious statement by James ([1907] 1978, p. 106): *“‘The true,’ to put it very briefly, is only the expedient in the way of our thinking, just as ‘the right’ is only the expedient in the way of our behaving. Expedient in almost any fashion . . .”*

I think James’s choice of the word “expedient” here was unfortunate, suggesting mere “convenience” or “usefulness”. Let’s take the spirit of his statement more sympathetically. What James is saying is that how we tell if a statement is true is by seeing if it works out in practice, and that there might not really be anything more to what it means for a statement to be true over and above how we tell that it is true. The intuition is that the truth of a statement, say “The cat is on the mat”, consists in the conjunction of a myriad of facts: that I

⁸ To help us think about operational coherence, I propose that we take Neurath’s boat, literally: “We are like sailors who have to rebuild their ship on the open sea, without ever being able to dismantle it in dry-dock and reconstruct it from the best components.” (Neurath [1932/33] 1983, p. 92) This is usually taken as a splendid coherentist metaphor for the fitting-together of propositions. But we can see boat-fixing as an activity illustrating the nature of operational coherence.

⁹ This section is a revised and expanded version of the discussion in Chang (2016), pp. 114–116.

have a visual image of a cat sitting on a mat; that my friend standing next to me does, too; that when my friend goes to lift the cat off the mat she does find something furry, warm and wriggling in her hands; that a screeching meow issues from that creature; that my friend ends up with a scratch on her hands; that the vet recognizes the cat as my old cat and not some fake robot-cat; etc. That may be all there is to what we mean by the truth of the statement, if we set aside the metaphorical projection of correspondence to the inaccessible “external world” in which the “real cat” apart from all of our feline experiences maintains its ghostly existence comparable to Ryle’s “ghost in the machine”.

Preserving James’s spirit but trying to avoid the obvious pitfalls, I want to propose a different formulation, in terms of operational coherence: *A statement is true in a given circumstance if (belief in) it is needed in a coherent activity.*¹⁰ For example, take the statement that the surface area of a sphere is proportional to the square of its radius. This statement is needed in a whole array of coherent activities, ranging from figuring out the amounts of paint needed to paint balls of different sizes, to Immanuel Kant’s deduction of the inverse square law of gravitation. So the statement in question is true in a wide range of circumstances. Now I hasten to add some much-needed elaborations and qualifications concerning various parts of this definition of truth.

(1) The definition above gives a sufficient condition for truth, but not a necessary condition (only “if”, not “only if”). This may be considered a grave philosophical defect, as the usual ideal of a definition is to give a necessary and sufficient condition. On the contrary, I think it is an advantage, inviting a more open-minded kind of philosophy. Giving a sufficient condition for truth makes a concept that we can use. Not giving a necessary condition allows that there may be other ways in which the concept is meaningful. It may turn out, after much reflection, that we decide that there is no other useful sense of “truth”, but at least I am not inclined to pre-judge that issue.

(2) The activity involved in the constitution of truth does not have to be that of explicit theory-testing (which would be the nearest activity-version of the

¹⁰ A slightly different formulation was given in Chang (2016), p. 114. I am using “statement” more or less synonymously with “proposition” here, which is an imprecision that can be fixed as needed.

idea of a statement corresponding to the world). Sometimes a statement is explicitly tested and confirmed, but other times its truth is shown in its involvement in the success of other kinds of activities. The pertinent activities do not even have to fall under the rubric of “epistemic activities”, by which I mean activities that are explicitly intended to increase and improve knowledge.

(3) Why I say “coherent” rather than “successful” in my definition should be clear from the meaning of coherence discussed in the last section. Defining truth in terms of the coherence of activities removes it one step away from direct verification (which would be done in terms of success). This may be unsatisfying at first glance, but it does make the notion of truth less closely tied to accidental successes and failures determined by case-by-case variations of fringe circumstances, which should be reassuring to those who worry that the James-style pragmatic notion of truth is too capricious.

(4) Requiring that (belief in) the statement in question should be “needed” in a coherent activity is designed to remove the worry that the statement might be involved in the activity in a superfluous way reminiscent of the tacking paradox or the Gettier problem (or employed in a purely fictional or instrumentalist manner). What is involved here is not a logical necessity that we can reason out *a priori*, but a pragmatic necessity, which can only be learned empirically. In other words, in the course of checking the truth of a statement, we ask: can the coherence of the activity be maintained, if we negate the proposition in question? For example, we can perfectly well use Maxwell’s equations while denying that the ether exists, so we know that belief in the existence of the ether is not necessary in relation to the coherence of the activity of solving Maxwell’s equations (while it *was* necessary in Maxwell’s original activity of model-building, which led him to the equations in the first place). Checking for pragmatic necessity may not live up to some overblown image of a philosophical test, but it is how we get on in science, and in the rest of life, too. To the problem of suspected superfluous propositions, there is no magic solution. As Clarence Irving Lewis (1930, p. 14) put it in his review of John Dewey’s *The Quest for Certainty*: “Salvation is through work; through experimental effort, intelligently directed to an actual human future.”

(5) If we base the notion of truth on pragmatic necessity, do we run the risk of turning it into a psychologistic concept?¹¹ What if our mental make-up is such that belief in a certain fantastical proposition is psychologically necessary for us to carry out some activities? If I can only swim by believing that I am a dolphin, is it true that I am a dolphin? On the face of it, this seems like a straightforward and devastating objection to my notion of truth, or any pragmatist notion of truth. But such alleged situations would need to be examined carefully. If I can also do the swimming by believing that I am a seal, then the belief in dolphin-hood is not necessary. If I can also do it by drugging myself rather than by relying on any beliefs about my identity, then no such beliefs are necessary. Besides, if others can swim without believing themselves to be dolphins, then the belief in dolphin-hood is not something necessary for the generalized activity of swimming, but something peculiar to me. But if we can imagine the case in which most people, no matter what they try, simply cannot swim without believing that they are dolphins (and therefore doing things like dolphins do), then we will need to consider whether we aren't actually dolphins.

(6) If truth is defined in terms of coherence, it has to be a matter of degree, and I think that is right. As Austin noted ([1950] 1979, pp. 117, 130–131), “very true”, “true enough”, etc. are perfectly sensible locutions, and it is unreasonable to try to reduce ordinary judgements of truth to yes/no. Many philosophers of science, mostly in the course of trying to defend scientific realism, have already fallen into the habit of speaking about “approximate truth”; Richard Boyd (1990) has argued convincingly that it is not possible to maintain scientific realism without relying on approximate truth. In order to escape this conclusion, perhaps one could say that “approximate truth” is an imprecise way of speaking, and what we are really talking about is approximation to the truth, while truth itself remains a yes-or-no matter. But I do not see what would be gained in preserving binarity for truth in that way.

(7) My definition only states what is true “in a given circumstance”, so truth comes attached with a specific scope (as well as coming in degrees). But

¹¹ I thank Mike Martin for raising a version of this worry when the spoken version of this paper was given at the Aristotelian Society. The example he used was about overconfidence on the part of scientists.

this is also to say that we may attempt to extend the scope of the truth of a statement constituted in one activity, through the use of the same statement in other coherent activities. A universal truth would be a statement that is true in all circumstances in which it can be applied.¹² When we say that mere convenience should not be mistaken for truth, that is normally because “the truth will out”; that is, we should not say that a statement is “true” without qualification, if we expect that it might be shown not to be true in some other circumstances. When we say “It may seem as if *P* were true in these circumstances, but *P* is actually not true”, what else can we be meaningfully asserting, other than that *P* is, or will be, shown to be false in some other circumstances? James’s notorious definition of truth quoted above actually continues as follows, which tends to confirm my reading of him: “. . . and expedient in the long run and on the whole of course; for what meets expediently all the experience in sight won’t necessarily meet all farther experiences equally satisfactorily. Experience, as we know, has ways of *boiling over*, and making us correct our present formulas.” The last bit of James’s statement actually fits very well with my notion that operational coherence is the only way in which reality gives input into our knowledge. And this gives operational-coherentist truth the mind-independence that realists value most in correspondence truth, while it is an “internal” notion meaningful within a system of practice, not without it.¹³

A slightly extended example may usefully illustrate further how the operational-coherentist notion of truth works. Take what was perhaps the single most important proposition in the history of organic structural chemistry in the 19th century: “Carbon has valency 4”, meaning that it is capable of forming 4 bonds with other chemical units (atoms or radicals). This statement was needed in the successful working-out of numerous molecular structures. It was also needed in the understanding and execution of substitution reactions. For example, a body of methane gas (CH₄) could be made to absorb a volume of chlorine gas and emit an equal volume of hydrogen gas, turning the methane into

¹² This is close to Peirce’s view that a belief is true if it “would never lead to disappointment, no matter how far we were to pursue our inquiries” (Misak 2007, p. 68).

¹³ Here I draw inspiration from Hilary Putnam’s “internal” or “pragmatic” realism; see Sosa (1993) for a convenient and insightful exposition.

chloromethane (CH_3Cl); such a substitution could be made four times in total, in the end yielding carbon tetrachloride (CCl_4). Such successful instances, it is fair to say, indicate the truth of “Carbon has valency 4.” But this truth was a limited one. We know, for example, that the structure of carbon monoxide remained a mystery for a long time. Even carbon dioxide was not trivial to understand, but it could be accommodated by saying that the carbon atom formed a *double* bond with each of the two atoms of oxygen (valency 2), thereby using up all of its 4 bonding-potentials, as indicated by the graphic formula $\text{O}=\text{C}=\text{O}$. But it was not clear at all how carbon monoxide (CO) could be understood.

As far as I can see now, my operational-coherence theory of truth does not differ substantially from James’s pragmatist theory of truth freed from misunderstandings, or from Dewey’s notion of “warranted assertability”, or from Charles Sanders Peirce’s account of truth as presented by Cheryl Misak (2007).¹⁴ According to all of these conceptions, if our use of a theory has led to successful outcomes and not as a result of any strange accident or coincidence as far as we can see, then we can and should say, modestly and provisionally, that the relevant statements made in this theory are “true” — in the same sense as we say that it is true that rabbits have whiskers and live in underground burrows. This “truth” is operational and verifiable. It is the same thing as empirical confirmation, taken in a broad sense.¹⁵ It is achievable, to various degrees, and its pursuit is clearly useful. A statement being true will mean that it passes all the tests of correctness that we can apply. “Is it true that there is an airport in Cambridge (the one in southeast England)?” (There is.) We know exactly how to answer such a question, and how to double-check and triple-check the answer as needed, and under which circumstances to start doubting the statement. As Putnam (1995, p. 10) put it succinctly, paraphrasing James: “Truth . . . must be such that we can say how it is possible for us to grasp what it is.”

When my work is more extended and developed, it will also include a similar pragmatist characterization of the notion of “reality”, of which I will only

¹⁴ In a separate paper I intend to give a detailed discussion of how my theory of truth relates to extant theories of truth, including pragmatist ones.

¹⁵ It is closely related to what I had called “truth₅” in Chang (2012), p. 242.

give a telegraphic summary here.¹⁶ The easiest way to see how operational coherence can also ground an operative notion of reality is to start with Ian Hacking's "entity realism": "If you can spray them, then they're real." (Hacking, 1983, p. 23) Concepts enabling successful activities deserve our realist confidence. I propose a coherence theory of reality: a putative entity should be considered real if it is employed in a coherent activity that relies on its existence and its basic properties (by which we identify it). This notion of reality (or realness) might be written with a lowercase "r" in order to distinguish it from the idea of "Reality" that denotes the whole "world" as some transcendent existence. Like truth in my operational-coherentist conception, "small-r" reality comes in different degrees, and is defeasible, as it is based on coherence.¹⁷ According to this notion of reality, phlogiston or caloric or ether, within its own domain of successful use, is as real as tables and chairs and cats and dogs are in our daily lives. When Hacking says that positrons are real, or when I say phlogiston is real (Chang 2012, chapters 1 and 4), the sense of it is that a specific *part* or aspect of the overall Reality is somehow being captured in our conception. And this parsing-out of Reality into various real entities is crucial in any kind of cognitive activity. If we cannot identify sensible parts (or aspects) of nature, we cannot say anything intelligible, make any kind of analysis, or engage intelligently with nature in any specific and directed way. So we have no choice but to worry about whether we are able to do the parsing well, and a kind of entity realism is prior to any truth realism one might hope for. But how can we ever tell whether we have done the parsing correctly? Again, "salvation is through work" — we can never be absolutely sure, but we check, double-check, and try checking continually in new domains of phenomena.

4. Words, deeds, and knowledge

Why do we want a theory of knowledge? I contend that it is in order to help us have more and better knowledge. Then our theory of knowledge needs

¹⁶ A slightly more extended account can already be found in Chang (2016), pp. 116–118, and in Chang (forthcoming), section 4.

¹⁷ C. I. Lewis (1929, chapter 7) employed this device of distinguishing "Reality" and "reality" for similar purposes.

to tell us something instructive about the processes through which knowledge is gained and improved. This is the direction in which I have attempted to steer my own epistemological thinking (see especially Chang 2011a). So, what have we learned in that regard from the consideration of operational coherence, and what do we still need to learn from it?

One important issue that deserves fuller consideration is the precise relation between knowledge-as-information and knowledge-as-ability. There are two distinct senses in which knowledge-as-ability is *larger* than knowledge-as-information. Firstly, I believe that Ryle was correct in saying: “knowledge-how is a concept logically prior to the concept of knowledge-that”, and “knowing-that presupposes knowing-how” (1945/46, pp. 4-5, and pp. 15-16); this notion is further articulated by Jennifer Hornsby (2007). In the other direction, knowledge-as-information enters as an important contributing element in knowledge-as-ability. Here is Ryle again (1945/46, 16): “effective possession of a piece of knowledge-that involves knowing how to use that knowledge, when required, for the solution of other theoretical or practical problems. There is a distinction between the museum-possession and the workshop-possession of knowledge.” Put together these two aspects of the embedding of knowledge-as-information into knowledge-as-ability, and we can begin to see belief in propositions as one particular aspect of knowledge, rather than its core or essence. Knowledge-as-information may only be flickering moments in the continual creation and use of knowledge-as-ability, and propositional belief only occasional crystallizations in that flow of activity.

These thoughts also point to a larger project of considering how verbal articulations aid life (see Polanyi 1958, chapter 5). The consideration of belief and truth does not exhaust the role of articulation in knowledge and in intelligent life. As Ryle (1945/26, 12) pointed out, the verbalizations that occur when we try to articulate the principles guiding our activities are in the imperative mood, not in the declarative/indicative. Look to imperatives for the most obvious occasions for the correspondence between the verbal and the non-verbal in life; the correspondence between what we say and what we do is what we really ought to concern ourselves with, rather than the imagined correspondence between the verbal and the transcendental. The philosophical

grammar¹⁸ of imperatives is an urgent task for philosophers of science trying to pay attention to scientific practice, and to pragmatists more generally. And don't we also need to pay similar attention to the philosophical grammar of interrogatives? In making these considerations, we would do well to remember Austin's caution ([1950] 1979, p. 131): "many utterances which have been taken to be statements ... are not in fact descriptive It is simply not the business of such utterances to 'correspond to the facts' (and even genuine statements have other businesses besides that of so corresponding)."

So, there is a great deal to do. But for now, I hope I have shown that moving away from the narrowly propositional view of knowledge allows us to retool the notion of truth so that it becomes operable, and similarly with the notion of reality. Thereby we can reclaim these key concepts for the use of people who are actually engaged in the production and improvement of empirical knowledge. We live in the world, and knowledge is only meaningful from that perspective within the world. It is a futile and pernicious philosophical dream to seek the God's-eye view, to hope to find an "external" perspective from which we can tell the "real" shape of the world. Roberto Torretti (2000, p. 114) blasts the "scientific realists" who believe "that reality is well-defined, once and for all, independently of human action and human thought," yet "in a way that can be adequately articulated in human discourse." They hold that science aims to develop "just the sort of discourse which adequately articulates reality — which, as Plato said, 'cuts it at its joints' —, and that modern science is visibly approaching the fulfilment of this aim." Torretti confesses that he finds it difficult "to accept any of these statements or even to make sense of them." The notions of truth and reality are in fact perfectly meaningful in the phenomenal realm of representing and intervening, and they should stay in that realm.

List of references

Austin, J. L. [1950] 1979: 'Truth'. In *Philosophical Papers*, 3rd ed., edited by J. O. Urmson and G. J. Warnock, pp. 117–133. Oxford: Oxford University Press.

¹⁸ Here I am, of course, harking back to Wittgenstein; see Baker and Hacker (1985).

- Reprinted from *Proceedings of the Aristotelian Society*, Supplementary volume 24 (1950).
- Baker, G. P., and P. M. S. Hacker 1985: *Wittgenstein: Rules, Grammar and Necessity*. Oxford: Basil Blackwell.
- Boyd, Richard 1990: 'Realism, Approximate Truth, and Philosophical Method'. In C. Wade Savage (ed.), *Scientific Theories*, pp. 355–391. Minneapolis: University of Minnesota Press.
- Carr, H. Wildon 1928/29: 'The Fiftieth Session: A Retrospect'. *Proceedings of the Aristotelian Society*, new series, 29, pp. 359–386.
- Chang, Hasok 2009: 'Operationalism'. In Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Fall 2009 Edition).
<https://plato.stanford.edu/archives/fall2009/entries/operationalism/>
- 2011a: 'The Philosophical Grammar of Scientific Practice'. *International Studies in the Philosophy of Science*, 25, pp. 205–221.
- 2011b: 'Compositionism as a Dominant Way of Knowing in Modern Chemistry'. *History of Science*, 49, pp. 247–268.
- 2012: *Is Water H₂O? Evidence, Realism and Pluralism*. Dordrecht: Springer.
- 2014: 'Epistemic Activities and Systems of Practice: Units of Analysis in Philosophy of Science after the Practice Turn'. In L. Soler, S. Zwart, M. Lynch, and V. Israel-Jost (eds.), *Science after the Practice Turn in the Philosophy, History and Social Studies of Science*, pp. 67–79. London and Abingdon: Routledge.
- 2016: 'Pragmatic Realism', *Revista de Humanidades de Valparaíso*, 4, pp. 107–122.
- 2017 (forthcoming): 'Is Pluralism Compatible with Scientific Realism?' In Juha Saatsi (ed.), *The Routledge Handbook on Scientific Realism*. London and Abingdon: Routledge.
- Foley, Richard 1998: 'Justification, Epistemic'. In E. Craig (ed.), *Routledge Encyclopedia of Philosophy*, vol. 5, pp. 157–165. London: Routledge.
- Hacking, Ian 1983: *Representing and Intervening*. Cambridge: Cambridge University Press.
- Hornsby, Jennifer 2007: 'Knowledge and Abilities in Action'. In Christian Kanzian and Edmund Runggaldier (eds.), *Cultures. Conflict - Analysis – Dialogue*

- (Publications of the Austrian Ludwig Wittgenstein Society, new series, vol. 3), pp. 165–180. Frankfurt: Ontos Verlag.
- Ichikawa, Jonathan Jenkins, and Matthias Steup 2016: 'The Analysis of Knowledge'. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition),
<<https://plato.stanford.edu/archives/win2016/entries/knowledge-analysis/>>.
- James, William [1907] 1978: *Pragmatism and The Meaning of Truth*. Cambridge, MA: Harvard University Press.
- Kuhn, Thomas S. 1962: *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Laudan, Larry 1981: A Confutation of Convergent Realism. *Philosophy of Science*, 48, pp. 19–49.
- Lewis, Clarence Irving 1929: *Mind and the World-order: Outline of a Theory of Knowledge*. New York: Dover.
- 1930: '[Review of] The Quest for Certainty: A Study of the Relation of Knowledge and Action [by] John Dewey'. *The Journal of Philosophy*, 27, pp. 14–25.
- Misak, Cheryl 2007: 'Pragmatism and Deflationism'. In Cheryl Misak (ed.), *New Pragmatists*, pp. 68–90. Oxford: Clarendon Press.
- Neurath, Otto [1931] 1983: 'Sociology in the Framework of Physicalism'. In Robert S. Cohen and Marie Neurath (eds.), *Philosophical papers 1913–1946*, pp. 58–90. Dordrecht: Reidel.
- [1932/33] 1983: 'Protocol Statements'. In Robert S. Cohen and Marie Neurath (eds.), *Philosophical papers 1913–1946*, pp. 91–99. Dordrecht: Reidel.
- Polanyi, Michael 1958: *Personal Knowledge: Towards a Post-critical Philosophy*. Chicago: University of Chicago Press.
- Putnam, Hilary 1995: *Pragmatism: An Open Question*. Oxford: Blackwell.
- Ryle, Gilbert 1945/46: 'Knowing How and Knowing That: The Presidential Address'. *Proceedings of the Aristotelian Society*, new series, 46, pp. 1–16.
- 1949: *The Concept of Mind*. London: Hutchinson.

- Snowdon, Paul 2004: 'The Presidential Address. Knowing How and Knowing That: A Distinction Reconsidered'. *Proceedings of the Aristotelian Society*, 104, pp. 1–32.
- Sosa, Ernest 1993: 'Putnam's Pragmatic Realism'. *The Journal of Philosophy*, 90, pp. 605–626.
- Stanley, Jason, and Timothy Williamson 2001: 'Knowing How'. *The Journal of Philosophy*, 98, pp. 411–444.
- Torretti, Roberto 2000: 'Scientific Realism' and Scientific Practice. In E. Agazzi and M. Pauri (eds.), *The Reality of the Unobservable*, pp. 113–122. Dordrecht: Kluwer.
- Truncellito, David A. 2007: 'Epistemology'. *Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/epistemo/>
- Young, James O. 2015: 'The Coherence Theory of Truth'. In Edward N. Zalta (ed.), *Stanford Encyclopedia of Philosophy* (Fall 2015 edition). <https://plato.stanford.edu/archives/fall2015/entries/truth-coherence/>