CHAPTER 1

Active Knowledge

1.1 Overview

How Science Goes beyond Propositional Knowledge

The most fundamental step in making an operational ideal of knowledge is to understand knowledge in terms of what people *do*. I want to focus on the sense of knowing how to do something, a sense of knowledge as ability. Philosophers normally take knowledge as the mental possession of true *propositions*, or *theories* (understood as organized sets of propositions). This conception is at the core of the dominant textbook treatment of epistemology with which mainstream analytic philosophers still start their training. Writing for the authoritative *Stanford Encyclopedia of Philosophy*, Jonathan Jenkins Ichikawa and Matthias Steup (2018) declare: 'the project of analysing knowledge is to state conditions that are individually necessary and jointly sufficient for propositional knowledge'. And then they remind us of the traditional ('tripartite') analysis of knowledge: 'justified, true belief is necessary and sufficient for knowledge.'[1] Of course, epistemologists at the cutting edge of research have moved far beyond this 'Epistemology 101' conception. But my own feeling is as with the old joke: 'How do I get to Letterfrack?' – 'I wouldn't start from here.'

My motivations for looking for a different starting-point come chiefly from the needs of the philosophy of science. The propositional view is simply not commodious enough for understanding scientific *practices*. And even for those who are not philosophers of science, paying attention to science in making a theory of knowledge makes good sense. While most aspects of human life involve knowledge in some way, in science the acquisition of knowledge is our main aim. Therefore, by observing what

[1] Similar statements are too numerous to cite in any comprehensive way. For typical examples see Dancy (1985, p. 23), and Audi (2014, p. 221).

people do in science we can learn something about what is involved in knowledge acquisition with relatively little background noise, and it is likely that we can learn some lessons applicable to knowledge in non-scientific situations as well.

Let us start by observing, without too much initial prejudice, what it is that scientists do. What kinds of things do we want to know in science? We do, of course, want to know some cut-and-dried facts. But we also want explanation and understanding, which many lovers of science prize above all. If philosophers don't want to deal with anything fuzzy like understanding, how about the knowledge of causes and mechanisms, at least? And what about having a sense of what it's like to experience nature, whether it be to smell chlorine gas or to communicate with a chimpanzee? What about knowing how to create new entities and phenomena, such as nuclear fission chain-reactions, or new 'super-heavy' chemical elements? Aren't all these things in the realm of scientific knowledge? It is not obvious at all that the main focus of scientific knowledge-making is just on propositions or theories, but we don't seem to have a philosophically rigorous way of thinking about this diverse array of types of scientific knowledge. Here we need to improve philosophy, instead of trying to dismiss or explain away the bits of science that do not fit our standard philosophical way of thinking. Before going on to the positive task, let me outline more carefully three different ways in which the philosophy of science is hampered by an exclusive commitment to the propositional view of knowledge.

(1) Much of what is prized as knowledge in science is in the realm of knowing how to *do something*, which is not easy to accommodate within the propositional view. As Ian Hacking (1983, ch. 13) has emphasized, *making* is an important mode of scientific work, regardless of any practical uses that our creations might have. More broadly, here is an important insight from Gilbert Ryle (1946, 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.' A handful of concrete examples from the rich tapestry of scientific progress in the last few centuries will give us a sense of the range of things that an epistemology of science ought to be able to handle. (I would ask you to spare one minute to take these examples in, rather than skipping through them as details unworthy of philosophical attention.) Since the eighteenth century we have known how to compute the precise trajectory of a planet, which involves

https://doi.org/10.1017/9781108635738.002 Published online by Cambridge University Press

knowing how to solve the relevant basic equations of physics. It is also important to know how to come up with such equations in the first place (and invent an appropriate kind of mathematics in which they can be framed). Nowadays we also know how to run simulations of experiments that we can't carry out physically, and how to make formal models of complex situations. We know how to predict the weather in the short term, and make good prognoses for many diseases. We know how to measure all sorts of things ranging from temperature and humidity to the rate of inflation. Feats of observation are not limited to quantitative measurement, either. We have learned how to ascertain molecular compositions and structures, and complex tasks like the sequencing of DNA molecules are now routine. We know how to image extremely distant objects with devices like the Hubble space telescope. We also know how to classify things in useful and effective ways, as with biological taxonomy or the periodic table of elements. We also have a great variety of abilities in the realm of making: synthesizing pharmaceutical agents, making high-temperature superconductors, creating and operating complex technological systems such as radio, television and the internet, and making artificial intelligence that can outperform humans in many complex mental tasks.

(2) Many thinkers who have taken a serious historical look at science have come to doubt that theories were the main units of scientific knowledge. Philosophers of science had traditionally worried about the problem of *theory*-choice, but Thomas Kuhn ([1962] 1970) showed quite convincingly that scientists' choice at the most crucial moments in history tends to be a choice concerning larger and more complex wholes, which he called *paradigms*. What exactly Kuhn meant by 'paradigm' is famously debatable, but it can be readily agreed that a paradigm contains much more than descriptive statements. That is especially evident if we take Kuhn's broader sense of paradigm as 'disciplinary matrix' (rather than as 'exemplar') (Kuhn [1962] 1970, pp. 181–7). A paradigm in this sense specifies some particular *problems* deemed worth addressing, the right *methods* of tackling those problems, and the *criteria* by which the solutions to the problems are assessed. *None of these elements consists of descriptive propositions.* Larry Laudan (1977; 1984) critiqued Kuhn on the 'holist' assumption that all elements of a paradigm must change together, but he agreed that a 'research tradition' contains diverse types of elements within it: theory, methodology and axiology.

A more recent comparison would be Paul Teller's (2021, p. S5016) rendition of the concept of 'framework', or Rachel Ankeny and Sabina Leonelli's (2016) notion of 'repertoires' in scientific practice, which pays due attention to the social organization of research. All of these thinkers have stressed that theory-choice is inextricably linked to aspects of scientific practice that are not reducible to descriptive statements.

(3) The propositional view of knowledge highlights the *products* of inquiry, mostly neglecting the *processes* of inquiry. One enduring shortcoming of mainstream epistemology and philosophy of science is traceable to this issue: we distinguish the context of discovery and the context of justification, and then proceed to ignore the former. Justification is a process that can be conceived reasonably well in propositional terms, so it receives far more attention in analytic philosophy than other processes involved in inquiry, such as hypothesis-generation and concept-formation. But even if we are just dealing with justification, serious attention should be paid to the ins and outs of the cognitive processes involved in it. Most epistemologists continue to conceive justification in propositional terms, ultimately based on how well a proposition agrees with other, better-established propositions. There are mainstream epistemologists who do emphasize processes, and chief among them are the reliabilists, in a tradition reaching back to Frank P. Ramsey. However, the 'reliability' of a process, rule or agent of inquiry seems to get defined inevitably in terms of the truth of the propositions that it produces or sanctions, so we return again to the focus on products (see Goldman and Beddor 2016 and references therein). The neglect of processes creates the misleading impression that knowledge is something we passively accept or reject, rather than actively seek, create and evaluate.

We need a fresh framework for thinking about what makes a good scientific practice. In previous publications I have proposed analysing scientific work in terms of **epistemic activities**[2] (and **systems of practice**, in which various epistemic activities function jointly), in conscious opposition to the more customary analysis framed in terms of propositions (Chang 2011a; 2011b; 2012a; 2014). These concepts will be elaborated further in Section 1.3. Very briefly and intuitively for now: an activity is a programme of action designed for the achievement of a recognizable aim

[2] I will put in boldface key terms of my own devising on their first occurrence.

(or, the execution of such a programme); an *epistemic* activity is a knowledge-related activity, aimed at acquiring, assessing or using knowledge. A system of practice is a network of activities that function coherently together. Activities are not reducible to propositions in any straightforward sense; instead, I will be thinking about how propositions fit into activities.

As an initial illustration of an epistemic activity, let's actually take something that might seem very far removed from actions: the definition of a concept. Consider what one has to *do* in order to define a scientific term: formulate formal conditions for its correct uses; construct physical instruments and procedures for measurement, standard tests and other manipulations; round people up on a committee to monitor the agreed uses of the concept, and devise methods to punish people who do not adhere to the agreed uses. 'One meter' or 'one kilogram' would not and could not mean what it means without a whole variety of epistemic actions coordinated by the International Bureau of Weights and Measures in Paris. Even semantics is a matter of doing, as Percy Bridgman (1927) and Ludwig Wittgenstein (1953) taught long ago.

For an illustration of a system of practice, consider the new system of chemistry that Antoine Lavoisier and his colleagues created in the late eighteenth century, bringing about the so-called Chemical Revolution (Chang 2012a, ch. 1). The main epistemic activities in Lavoisier's 'oxygenist' system of practice included: making various chemical reactions; collecting the products of reactions, especially gases; identifying various substances through standard chemical tests; analysing organic substances through combustion; measuring the weights of the ingredients and products of reactions; tracking chemical substances through those weight-measurements; and so on. Some of these were already well-established activities in chemistry, and others were more novel. These activities were coordinated together for the purpose of achieving the overall aims of the system, such as the knowledge of the chemical composition of various substances, a good classification of all chemical substances, and the explanation of chemical reactions. Systems of practice form coherent bodies of scientific work, and they also make useful units of analysis for philosophers and historians of science.

Active Knowledge

I now want to introduce an action-based view of knowledge that can help us reach a better understanding of practices in science and other realms of

https://doi.org/10.1017/9781108635738.002 Published online by Cambridge University Press

life. But first I should acknowledge that many leading epistemologists, going at least as far back as Bertrand Russell (1912, ch. 5), have recognized that there are different kinds of knowledge, not just the propositional kind. For example, Keith Lehrer's classic text of epistemology (1990, pp. 3–4) starts by giving a list of all sorts of knowing: 'I know the way to Lugano. I know the expansion of pi to six decimal places. I know how to play the guitar. I know the city. I know John. I know about Alphonso and Elicia. I know that the neutrino has a rest mass of 0 . . .'[3] Yet, Lehrer explains that only propositional knowledge (or, knowledge 'in the information sense') is the main concern of epistemology because 'it is precisely this sense that is fundamental to human cognition and required both for theoretical speculation and practical agency'. This is an entirely typical starting point among epistemologists in the tradition of analytic philosophy. They may debate endlessly about the precise nature of justification and the exact meaning of truth, but it is rarely questioned that knowledge is a matter of *belief* in *propositions* that give *information* about the world, expressible in the form of well-formed statements.[4]

On reflection it is not clear at all that **propositional knowledge** (or **knowledge-as-information**[5]) is so much more important or fundamental than other types of knowledge that it should command the undivided attention of philosophers. I am not suggesting that the proposition-focused orthodox epistemology is *wrong*, but I do think that it is *limiting*. It obliges us to disregard many kinds of things that we readily regard as 'knowing'. One could try to argue that all these types of knowledge have less philosophical importance than propositional knowledge. As Kitcher points out, a crucial feature of propositional knowledge is that it is the most suitable form of knowledge for public communication.[6] But is that necessarily more important than the private dimension of knowing, or the tacit communication of knowledge between people in direct personal contact with one another? In the end, I am not convinced that we should expend a great deal of intellectual energy in trying to defend the primacy of propositional knowledge.

Instead, I want to endorse and develop an alternative view, which thinks about knowledge primarily in the context of action – which is to say, in the

[3] See also Snowdon (2004, p. 5).

[4] I will sometimes use 'statement' and 'proposition' interchangeably, except where it is important to distinguish a statement from its content, which is the proposition.

[5] I use 'information' here as a commonly understood idea, leaving it to those who see knowledge as information to settle on a more considered view of its nature.

[6] Philip Kitcher, personal communication, 29 January 2021.

context of life itself. This is a perspective articulated long ago by Marjorie Grene (1974, pp. 172, 158): 'knowing is something people do, an activity'; it is the 'full, concrete, historical person who is the essential agent of knowledge' (I will say more about the nature of epistemic agents in Section 1.2). In the context of action, the primary sense of knowing is knowing how to do something. 'I know how to do *X*' (distinct from 'I know how *X* is done' as a description not necessarily accompanied by actual ability) is a common, meaningful and important thing to say. I imagine there are roughly equivalent expressions in most human languages. This is a clear and well-established notion of 'knowing', and it should not be presumed to be reducible or subordinate to knowledge as the storage and retrieval of information. I propose to use the phrase '**active knowledge**' to designate this sense of **knowledge-as-ability**. Paying attention to active knowledge can help us greatly in making full sense of how we acquire and evaluate knowledge in science and various other walks of life.

Etymology is not philosophy, but it can give some suggestive clues: the Indo-European root *gnō* is the common source of both 'know' and 'can' in English, as well as the Scots word *ken* and the German *kennen* and *können*. The old English word *cunnan* had all of the related meanings: 'to know, know how to, be able to' (Watkins 1985, pp. 23–4; also Shipley 1984, pp. 129–33). This is in line with Wittgenstein's insight (1953, 59, §150): 'The grammar of the word "knows" is evidently closely related to that of "can", "is able to".'[7] As Anthony Kenny put it: 'to know is to have the ability to modify one's behaviour in indefinite [not pre-determined] ways relevant to the pursuit of one's goals' (Kenny 1989, pp. 108–9, quoted in Hyman 1999, p. 438).

Active knowledge is at the core of scientific knowledge. We should also note that scientific abilities are not entirely distinct or disconnected from the range of abilities that support ordinary human life, or even animal life. Active knowledge in science may simply be more systematic than active knowledge in other spheres of life, in line with Paul Hoyningen-Huene's (2013) point that systematicity is what defines science. Ordinary members of human societies know how to speak languages, make inferences, tell lies, give explanations, argue with each other, and count and sort things. Most of us know how to run, swim and kick balls. Some of us even know how to sing, paint and make pottery. We know how to recognize and remember people's faces, find our destinations through complex routes, and imagine fictional things. We know how to cook food, clean house and cut

[7] He adds: 'But also closely to that of "understands".'

https://doi.org/10.1017/9781108635738.002 Published online by Cambridge University Press

fingernails. We may know how to grow crops, build bridges, perform surgery, raise children and teach skills to other people. Why should we deny that there is significant *knowledge* in all these activities, just because it may not exist in a propositional form?

We should want to have an epistemology that can handle such common items of knowledge, and starting with the notion that knowledge consists in *belief* will not get us there easily. We should try to deal with abilities directly, rather than skirting around them in an awkward and roundabout way, treating them as the applications of beliefs or inessential accompaniments to propositions. In the realm of purposive human action, knowers are active living and inquiring agents, and knowing is a state of such agents. Various leading philosophers have articulated the same kind of insight that I am trying to convey here, and I will try to extend and synthesize their ideas. I have already mentioned Grene, and her work developed in close connection with Michael Polanyi's (1958). Some others who have gone in similar directions are even regarded as founders of analytic philosophy – Wittgenstein, Ramsey, Ryle and also J. L. Austin. In current analytic philosophy, too, there are some congenial moves. Timothy Williamson's 'knowledge first' epistemology is right to resist a reduction of knowledge down to other notions such as belief, and to regard knowing as a mental state of the agent, namely the 'most general factive mental state'. But when he explains that being factive is a 'propositional attitude' that one has only to truths, it is clear that his focus remains on propositional knowledge (Williamson 2000, pp. 33–4). The tradition of virtue epistemology going back to Ernest Sosa is more promising, as it places the familiar account of propositional knowledge in the context of action, holding that 'judgment and knowledge itself are forms of intentional action' (Sosa 2017, p. 71). However, it seems to me that the potential of this approach is not fully realized. According to Sosa (ibid., p. 73) intentional action (an attempt) is apt if and only if it is successful because competent. This strikes me as modelled too closely on the textbook epistemology of propositional knowledge, according to which a belief is knowledge if and only if it is true because justified.

Instead of following such lines of development, I look back to the tradition of pragmatism. Pragmatist philosophers have clearly recognized the need to understand and assess knowledge in the context of action. I want to focus on what pragmatism can tell us about method-learning and practices of inquiry, pointing to a conception of active knowledge that can be useful in the philosophy of science. In Section 1.6 I will give a more detailed exposition of what I take pragmatist philosophy to be. For now, hear William James, who said long before Grene and others: 'The knower is an actor' (quoted in Putnam 1995, p. 17). John Dewey (1917, p. 12) went on to develop this vision fully, complete with his own memorable slogan: 'we live forward'. For Dewey experience is active, full of expectations and reactions, contrary to the impoverished view of it in traditional empiricism as the recording of information through sensory input. Experience is not just *given*, but *taken* by active agents. So is knowledge. Inquiry is pervasive in life, an essential activity of an organism coping in its environment, as emphasized recently by Joseph Rouse (2015).

Propositional Knowledge Embedded in Active Knowledge

In proposing to take knowledge primarily as ability, I am not suggesting that we ignore propositional knowledge. Active knowledge (knowledge-as-ability) and propositional knowledge (knowledge-as-information) are both in operation in science and everyday life. The urgent task is to clarify the relation between the two. We need to think about the *functions* of propositional knowledge, rather than just taking it as the be-all and end-all of human intellectual activity. Especially those of us who consider ourselves intellectuals should guard against taking it for granted that propositional knowledge is valuable in itself. Let's face it: possessing information is not an end in itself. Or rather, we should actually only count as 'information' what is *informative*, and what is informative depends on our purposes and our situations. In transmitting information we obviously want to achieve a high signal-to-noise ratio, but what is signal and what is noise depends on what we want to learn and why. In the words of Nicholas Maxwell, who has long been calling for a philosophy of wisdom instead of mere knowledge, 'the aim of science is not to discover truth *per se*, but rather . . . *valuable* truth' (Maxwell 1984, p. 91; emphasis original). On the question of the place of human values in science there is now a robust literature in the philosophy of science.[8] This book will not be a contribution to that literature, but I do think that my focus on active knowledge can help bring the considerations of values and knowledge closer together. Thinking about knowledge-as-ability brings knowledge inescapably into the context of action, where values are immediately present. Then we can have a full view of how knowledge functions in life.

[8] See, for a few salient examples, Raskin and Bernstein (1987), Longino (1990), Douglas (2009), Kitcher (2011a), Carrier (2013), Brown (2020).

To refocus our view on the relation between active knowledge and propositional knowledge, let us first pull back to the general point that there are various types of knowledge, as already mentioned: knowing facts, knowing reasons (having a causal or intentional explanation for something), knowing someone or something (knowledge-by-acquaintance), knowing what it's like to be someone or to experience something (empathetic understanding), and more types besides. A common impulse concerning this situation is to understand all of these types of knowledge in propositional terms.[9] My own inclination is to recognize that various types of knowledge interact in a complex way with each other, rather than all being reducible to propositional knowledge. For example, consider what we really mean when we say we know someone. (In the sentences to follow, I propose a new gender-neutral third-person singular pronoun: e, er, em, ers.[10]) It involves the ability to recognize the person and distinguish em from other people – by er voice, by the feel of er hand, by the look of er face, by the shape of er body, and so on. It also involves knowing, implicitly and explicitly, some basic facts about em, including er habits and er typical reactions to situations. It also involves having a memory, often unarticulated, of some experiences that e and I have shared.[11]

Having recognized the complexity of the picture, I propose to simplify it in an important and useful way by conceiving active knowledge as a more encompassing category than the others (without implying a *reduction* of other types of knowledge to active knowledge). In this picture, all other sorts of knowledge depend on active knowledge, and they also contribute to it. To continue with the previous example: knowledge-by-acquaintance certainly depends on active knowledge as we have seen, and it also contributes crucially to active knowledge – how would we be able to do anything at all, if we did not know who was who and what was what?

[9] For sophisticated discussions in this direction, see Craig (1990), chs. 16, 17; Stanley and Williamson (2001); and Lawler (2018), ch. 2.

[10] Not only is 'e' what 'she' and 'he' have in common, but 'e' (이) just happens to be a word in Korean (my native language) indicating a 'person' gender-neutrally. It often occurs as an ending (as in 어린이 meaning 'child', and 늙은이 meaning 'old person'). I believe there is a clear disadvantage to using 'they' as both singular and plural. It is bad enough that we lost the second-person singular 'thou' long ago, later necessitating the invention of *y'all*!

[11] Knowing by acquaintance is a distinct enough kind of thing, given a different word in French and German to distinguish it from a more factual kind of knowing: *connaître*, not *savoir*, and *kennen*, not *wissen*.

From that point of view let us now return to the relation between active knowledge and propositional knowledge. It is not at all that propositional and active knowledge are opposites on the same plane, as people often seem to imply about knowledge-that vs. knowledge-how. Rather, we need to ask how propositional knowledge fits into active knowledge. The answer is twofold: propositional knowledge depends on active knowledge, and it contributes to active knowledge. Ryle gave us crucial insights on both sides of the picture. Concerning the first, he points out that knowledge-that (roughly equivalent to propositional knowledge) can be operational only when it is properly embedded in knowledge-how (similar to active knowledge). He even argues that 'knowledge-how is a concept logically prior to the concept of knowledge-that', and 'knowing-that presupposes knowing-how' (1946, pp. 4–5, 15–16).[12] In order to have any propositional knowledge one must be able to use a language. You can read this book only because you have at some point *learned to read*. The learning of one's first language requires tacit active knowledge, starting with the ability to play the pointing-game.[13] The very invention of language must have been done by people who knew how to conceptualize and communicate things in some non-verbal ways. Or consider propositional knowledge expressed in mathematical equations and formal models. What a great number of skills one needs to master in order to be able to *do* mathematics! Knowing how to multiply two large numbers is an ability, as is knowing how to solve a set of simultaneous equations, how to integrate a complicated function, or how to construct a proof. Some things in mathematics can be done by following an algorithm (if you know how to do *that*), but most mathematical work requires task-specific skills. Even if the tasks can be broken down to simpler ones, they will only bottom out in basic abilities like knowing how to count, which are not reducible to propositional knowledge. Propositional knowledge cannot stand without all sorts of active knowledge.

In the opposite direction, it is easy to recognize how propositional knowledge contributes to active knowledge. We should pay close attention to the ways in which propositional knowledge is employed by epistemic agents in their activities. Here is Ryle again (1946, p. 16): 'effective

[12] Alva Noë puts the point aptly (2005, p. 285): 'grasping propositions itself depends on know-how; but if know-how consists in the grasp of further propositions, then one might wonder whether one could ever grasp a proposition.'

[13] That ability is not to be taken for granted, as I have learned from eight years of trying in vain to point out to a friendly and very clever squirrel where I have put the nuts out for her. On the complexity of the pointing gesture in gorillas and humans, see Gómez (2004, pp. 186–90).

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'. That is to say, the very *point* of propositional knowledge lies within active knowledge. Ryle made a memorable distinction 'between the museum-possession and the workshop-possession of knowledge'.[14] The image of the workshop pushes us to ask what we *do* with propositional knowledge, and how beliefs (and other epistemic attitudes) concerning propositions fit into our epistemic activities, and contribute to the active knowledge embodied in those activities. Having belief in propositions is one particular aspect of knowledge, rather than its essence, or even its central core. Metaphorically speaking, propositional knowledge may only be occasional and localized crystallizations in the flow of activity that is the creation and use of active knowledge.

Operational Coherence and Its Improvement

How do we evaluate or assess the *quality* of active knowledge? This is a vexing and fascinating question, reminiscent of Robert Pirsig's (1974) struggle to understand 'quality' in *Zen and the Art of Motorcycle Maintenance*. For propositional knowledge there is one clear chief criterion by which its quality is evaluated: whether the statement in question is true. But the truth-criterion does not apply immediately to active knowledge. How, then, should we judge the epistemic quality of activities? Since active knowledge is a matter of ability, it may seem that success can serve as the main criterion for evaluation. If I claim to know how to do something, the most obvious test would be to see how well I can actually do it. But it is not quite so simple as that. I may do something successfully but only by some lucky accident, or through misconceptions that happen to work out. Not even a certified and reliable ability to do something, by itself, constitutes active knowledge. Paul Snowdon put the point memorably (2004, p. 18): 'No one would affirm that, because I can bleed or digest a three course meal, these are things I know how to do.' To say I *know* how to do something implies some sort of *understanding*, which is parallel to the element of justification that enters into the 'justified true belief' account of propositional knowledge.[15]

[14] This gives an overly traditional view of what happens in museums, but it was probably not far off in relation to the museums of his time.

[15] Even in relation to propositional knowledge, it should be philosophical common sense that 'knowledge goes beyond the mere possession of information' and requires a kind of understanding, a 'capacity to distinguish truth from error' (Lehrer 1990, pp. 4–5).

The understanding required for active knowledge involves a purposive aspect that is also systematic, even in relation to the simplest of acts: in order to try to make something happen, the agent has to *coordinate* carefully various movements and thoughts with each other and with external circumstances, towards the achievement of an *aim*. (This coordination may not be explicit, and the aims may not be conscious. And understanding may be something we attribute to others, as well as a matter of self-recognition.) I propose to use the term **operational coherence** to refer to such a state of **aim-oriented coordination**. Operational coherence is a key concept that I will use throughout the rest of this book, and I will spell out its meaning more fully in Section 1.4. The rough idea, metaphorically, is that a coherent activity makes sense because what goes into it all fits together nicely; the coherence consists in various aspects of the activity coming together in a harmonious way towards the achievement of its aims. In using the term 'operational' here I am giving a conscious nod to Bridgman's philosophy as noted in the Introduction. Within recent philosophical literature, the closest point of contact I have found is the work of Paul Thagard in his aptly titled book *Coherence in Thought and Action* (2000). I must stress that operational coherence is not primarily about the logical relationship between propositions. As a quality pertaining to activities, it is also not meaningful in the absence of agents who carry out purposive actions. There is a strong hermeneutic aspect to operational coherence, as it is based on how actions make *pragmatic sense* within a purposeful activity. It is a matter of *doing what makes sense to do*. Operational coherence does not reside in the 'mind-independent world', yet it expresses the empirical ('external') constraints on our thought, because the design of a coherent activity incorporates what we have learned from experience about what tends to make sense to do and what does not.

To understand the nature of active knowledge fully, we must also pay attention to the processes by which it is acquired and improved – in other words, to the epistemic activities that constitute *inquiry*. The need for attention to inquiry becomes even clearer if we ask why we want to have a theory of knowledge at all. Echoing Kitcher (2011b, p. 508), I contend that a fundamental purpose of epistemology should be to help us *get* more and better knowledge. If so, epistemology needs to tell us something instructive about the processes through which knowledge is gained, improved and evaluated, so that we can manage them better. This is the direction in which I have attempted to steer my own epistemological thinking (see, e.g., Chang 2011a).

Let's start with Charles Sanders Peirce's view that inquiry begins with a disturbed and unsettled state of doubt.[16] Dewey took this view and developed it to the full, from a more clearly action-oriented view that took inquiry as something done by an organism in its environment: '*Inquiry is the controlled or directed transformation of an indeterminate situation into one that is so determinate in its constituent distinctions and relations as to convert the elements of the original situation into a unified whole.*'[17] What did he mean by converting the elements of a situation into 'a unified whole'? I like to imagine that he was pointing to something like my notion of operational coherence, and a pragmatic sense of understanding that comes from operational coherence. Here is a very suggestive remark from Dewey (1925, p. 50): 'The striving to make stability of meaning prevail over the instability of events is the main task of intelligent human effort.'[18] So some sort of *sense-making* is clearly seen by Dewey as an important aim of inquiry, to overcome the initial puzzlement that sets inquiry off.

Inquiry is a striving towards greater operational coherence. In order to escape the state of disorder that prompts inquiry, we must change something about the situation in order to create more operational coherence in the activities that we can perform in the situation. This adaptation takes the form of **aim-oriented adjustment**, a concept that I will explain fully in Section 1.5. This adaptive process is driven at each moment by the relief and satisfaction provided by increasing coherence, without a fixed final destination. Anything and everything that is within our power to change may be changed in the process of increasing operational coherence. There is nothing in our knowledge that is fixed and validated forever and unconditionally, and inquiry does not follow pre-determined and eternal methods. Inquiry is a process of method-learning as well as content-learning. We are truly floating in Neurath's boat, '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 [1932/3] 1983, p. 92).

[16] Peirce (1877, p. 5). See also Peirce (1986), W3:248. According to Cheryl Misak (2013, pp. 32–3), 'in Peirce's view, what is wrong with the state of doubt is . . . that it leads to a paralysis of action'.

[17] Dewey (1938, pp. 104–5); emphasis original. Rouse (2015) is an exemplary inheritor of this view of inquiry among our contemporary philosophers. Taking inquiry explicitly as a process, Dewey laid out the following steps of it (1938, pp. 105–12): antecedent conditions; institution of a problem; the determination of a solution; reasoning and the examination of meaning; and the development of an idea 'in terms of the constellation of meanings to which it belongs' (ibid., p. 112).

[18] I thank Céline Henne for pointing out this passage to me, and for numerous other pointers and insights concerning Dewey's work.

All attempts to locate certainty at the foundation of our knowledge having failed, various philosophers in the early twentieth century articulated the realization that the traditional ideal of certainty must be removed if we are to forge a realistic ideal of knowledge. These included Wittgenstein, especially in the notes published posthumously as *On Certainty* (1969), and Dewey in *The Quest for Certainty* (1929). If the struggle to settle an unsettled situation is sufficiently successful, it will result in a situation that is sufficiently settled for helping us launch further inquiries that are more restricted. Such restricted forms of inquiry, including fact-gathering and learning how to perform certain well-defined tasks, are the types of work that Kuhn put under the rubric of 'normal science', which can happen once a paradigm is well-entrenched. These may be the more readily recognized sort of inquiry, but they can only function after some successful unrestricted inquiry has been carried out in order to fix the framework within which they are carried out, as Céline Henne (2022) discusses in making her distinction between 'framing' and 'framed' inquiry.

Plan of the Rest of the Book

What next? I hope you will want to read on, and there are different options for how to proceed. If you want to get a full and detailed treatment of the themes discussed so far, please continue with the remaining sections in this chapter (or at least some selected sections), which will discuss various specific aspects of active knowledge in more depth and detail: epistemic agents (Section 1.2), epistemic activities (Section 1.3), operational coherence (Section 1.4), and inquiry (Section 1.5). I will also explain the affinity of my ideas to pragmatism (Section 1.6).

If you are already quite persuaded by the programme laid out so far and anxious to find out the rest of the story, you might want to go directly on to Chapters 3, 4 and 5. Just take in the overview (Section 1) of each chapter if you have limited time and interest, but I hope you will get into some of the further sections as well. Selective reading can be guided by the summary at the beginning of each section. (In order to distinguish visually the friendlier surface-level summaries from the more specialist and detailed discussions, I have put the Introduction, the overview of each chapter, and the summary of all the other sections in a different typeface.) Chapters 3 and 4 will show how I use the notion of operational coherence in order to reconceive the very notions of reality and truth. Building on these views of truth and reality, Chapter 5 will present a doctrine of realism that is suitable for realistic people.

But before I get to all of that, a *clearing* needs to be made, where these ideas can have space to grow. This will be the task of Chapter 2, which will try to show how we can and should get away from the well-entrenched notion that our knowledge of nature is a matter of *correspondence* between our theories and the ultimate reality 'out there'. That will be the only primarily negative and critical chapter in the book, the other chapters being mainly devoted to the positive exposition of my own ideas. If you feel that the unorthodox approach to knowledge that I have laid out so far is unnecessary because more orthodox approaches will do the job equally well or better, please do read Chapter 2 before other chapters (or before abandoning the book altogether). On the other hand, if you are in agreement with my general approach and would like to see how I defend it against orthodox views, you may also find Chapter 2 interesting and instructive for that purpose.

1.2 Epistemic Agents

To understand the nature of active knowledge in a philosophically rigorous way, we need to have a clear and precise characterization of epistemic activities. That task begins with understanding knowers as full-fledged agents. Epistemic agents do not simply possess beliefs and desires, which are the chief notions employed in mainstream discussions in the philosophy of action. They are beings also endowed with certain physical and mental capacities, who engage in purposive actions, and make genuine choices and judgements. Epistemic agents are embedded in social communities that embody and enforce certain normative standards. There is an iterative process of emergence through which individuals arise from society and society is constituted through interactions between individuals, ascending to ever-higher levels of sophistication.

Basic Properties of Epistemic Agents

A full account of active knowledge requires a good *ontology* of epistemology. There is a tendency in the philosophy of science to present the scientist as a ghostly being that just has degrees of belief in various descriptive statements, which are adjusted according to some rules of rational thinking (e.g., Bayes's theorem) that remove any need for real judgement. Whatever does not fit into this bizarre and impoverished picture, we tend to denigrate as matters of 'mere' psychology or sociology. We need a more serious understanding of scientists as agents, not as passive receivers of information or algorithmic

https://doi.org/10.1017/9781108635738.002 Published online by Cambridge University Press

processors of propositions. With a brief sketch of epistemic agents here, I want to at least express a recognition of the issues that we need to think about.

Start again with the motto from James and Grene: the knower is an actor. The first step is to recognize purposive behaviour in the knower, at least in terms of instrumental rationality. The most basic thing about an agent is that e acts purposefully (or can be understood as acting purposefully), striving to achieve certain aims that are formulated on the basis of er desires and beliefs. The agent takes the kind of actions that e believes will contribute towards the satisfaction of er desires. That is the 'standard story of action' in philosophy, as Jennifer Hornsby calls it, according to which actions are 'belief-desire caused bodily movements' (2007, p. 180, also p. 165). Hornsby (2004, p. 2) has criticized this account strongly as not giving a truly active role to the agent, 'not a story of agency at all'. But actually, in philosophy of science even just a proper recognition that the epistemic agent has desires, rather than just beliefs, would be a significant advance. There are many types of pleasure that motivate scientists (and other human beings), including physical comfort and sensual well-being, abstract and concrete understanding, love and conviviality, self-esteem, security, legacy, and a sense of beauty, order and coherence. And on top of that, of course, we have to think about the things that various people have regarded as 'the aim(s) of science': truth, empirical adequacy, economy of thought, etc., which are best understood in terms of values. In order to understand concrete practices, we need to see how such desires and values shape the specific proximate aims that drive particular epistemic activities, as I will discuss further in Section 1.3.

Aside from desires and values, epistemic agents have beliefs indeed. But we need to think about much more than just explicit and articulated beliefs. There are also things we take for granted without examination, and such presumptions are necessary for enabling any kind of action. Of particular importance are expectations concerning the future. Expectations are often not beliefs at all, if by belief we mean a conscious assent to an articulated proposition. Expectations often exist on the 'horizon' as Edmund Husserl would have it (see Føllesdal 1990), or, in the tacit dimension in Polanyi's way of thinking. They can even consist in *not entertaining* certain possibilities. When I am walking along as normal, my expectations involved in that activity will not be exposed or even formulated until they are met with something incoherent with them, such as the tremors of an earthquake, or the left arm grabbed by an excited old friend, or a gaping hole in the pavement. Scientific practice is also full of expectations, sometimes guiding our activities smoothly, sometimes preventing certain activities, sometimes making us attempt something repeatedly without a clear sense of why.

Something often neglected in philosophical accounts of science is epistemic agents' capabilities, or capacities, which are considered mostly in the discussions of ethics and human flourishing. In a related way, capacities are discussed in terms of legal responsibilities, or cognitive capacities in the philosophy of mind. Hornsby criticizes the standard story in the philosophy of action for leaving out agents' capacities, too:

> When abilities are allowed a place in the explanation of action, it becomes clear how narrowly focused are the explanations from agents' reasons given in the standard story. I said that it would be a sort of magic if someone's intentionally doing something were consequential merely on their having a desire and a belief. (Hornsby 2007, p. 170; see also 2004, pp. 20–2)

It is important to recognize both mental and physical capabilities here, and also their mutual entwinement, remembering Polanyi's (1958, ch. 4) emphasis on the role of skills in scientific work, with due attention to embodiment. Most of the specific capabilities have to be learned, so we need to consider the process of learning and training. The consideration of capabilities needs to enter the discussion of a wide range of issues in the philosophy of science, such as observability, testability, simplicity and incommensurability – and therefore also realism, demarcation, confirmation, theory-choice and so on. Considerations of scientific rationality are greatly hampered if we do not consider what the capabilities of the agents involved are, because the judgement of what is rational for them to attempt depends greatly on what they are in fact able to do (remember the old lesson in ethics: 'ought implies can').

The last aspect I want to stress in the ontology of epistemic agents is the fact that they make choices and judgements. The standard story of action does not seem to leave any room for judgement, regardless of the types of cause that the explanations appeal to: utility-seeking, cognitive-psychological, neurological, sociological, what have you. For example, in an interest-based sociological explanatory schema, the picture of the individual is just as impoverished as that of the utility-maximizer in the individualist rational-choice theories. Instead, we need to find ways of giving some substantive meaning to words like 'choice' in the phrase 'theory choice', and 'decision' in 'decision theory' (cf. Bradley 2017). Even if the correct account of actions would be ultimately deterministic, decisions to act in specific ways are determined by a distinct combination of aims, beliefs and capabilities for each individual and each community. This will create at least an appearance of freedom on the part of the agents, and make the reality of scientific practice pluralistic.

The Social Dimension: Iterative Emergence

It is crucial to take into account the social nature of epistemic agents. This is no place to attempt a full social ontology, social epistemology or sociology of knowledge, but I do want to offer some pertinent reflections that I have found useful in framing the discussions to follow in the rest of this book. It is an urgent future task to improve my idiosyncratic home-made social theory with better attention to the considerable existing literature (for a notable recent synthesis see Epstein 2015). As Helen Longino (forthcoming) stresses, it is crucial for social epistemology to pay attention to interactive processes, and individuals only become epistemic agents through deliberative interactions.

Here is an unlikely source of inspiration. In the chapter on 'conviviality' in his classic *Personal Knowledge,* Polanyi laid out the necessity of the social in the epistemic. He stressed the crucial role of the 'civic coefficients of our intellectual passions' (1958, p. 203), such as the 'sharing of convictions', the 'sharing of a fellowship', 'co-operation' and the 'exercise of authority or coercion' (ibid., p. 212). These operate at a tacit level in the first instance, effective even in various non-human animals. This is an often neglected aspect of Polanyi's thinking: the tacit is embodied and individual, but also inherently social.[19] Wittgenstein and Polanyi concurred that knowledge must be founded on the trust we place in others, and the store of facts we rely on necessarily rests on the testimony of others (see Daly 1968). Martin Kusch (2002) and others have built on this tradition greatly.

The recognition of the necessity and priority of the social should not be flattened into the thought that 'everything is social'. For something like scientific knowledge to arise, we must have independent individuals as well as unindividuated collectives. The society–individual interaction needs to be conceived in an iterative way, avoiding reductionism in both directions. Individual action and cognition are grounded in society, but they are not *merely* social, and we should not presume that they are explainable or even fully describable by means of collective factors alone. The rational individual agent arises from the social matrix, but with a capacity for independent thought and dissent. And from the association of such individuals emerges a higher level of sociality that forms an integral aspect of life as we know it, and also forms the basis of advanced systems of knowledge. This iterative emergence of the individual from the social and the social from the individual continues indefinitely.

[19] And there are remarkable abilities that non-individualist and mostly inarticulate society can achieve, as in the case of bees and ants.

The first thing to understand in this whole picture is the process of individuation.[20] How does the distinctive individual come to exist? Each developing individual has a particular physical and mental make-up that is different from others, and a life trajectory that is inevitably different from anyone else's. The individual also has a capacity for seeking operational coherence, for *pragmatic sense-making* that harmonizes er thoughts and actions with er make-up and circumstances. What makes sense for each person is different, and this prompts *dissent*, an inner voice that says 'no' to what others say, an inner revulsion against social expectations. This is as real as the submission-instinct shown in the Milgram experiment that Barry Barnes (2000) makes so much of. If the social is entirely satisfactory, there is no need for the individual. The individual defines erself *against* the social.

Individuation is especially important when we consider science. I would even say that the inquiring attitude is essentially linked to the emergence of the individual. In an important sense, science begins with dissent and critique. This is tied with Dewey's (1938) notion that inquiry begins with an unsettled situation. Inquiry wouldn't start at all without being motivated by some individual's dissatisfaction. There is also a deep connection here with the fundamental empiricism in science, which is founded on the authority of observation, which has an irreducibly individual, even private, dimension. Saying 'let me check for myself' is the beginning of the scientific attitude, and it is a fundamentally *individualist* declaration (Bridgman 1940; Pritchard 2016). This is not the old methodological individualism, which ignores the social grounding of individuals. Rather, it is an activist individualism promoting and celebrating the rise of the autonomous individual standing up on the social ground.

But this is certainly not the end of the story. From the association of self-actualizing individuals emerges a higher level of sociality. In this process there is one aspect that deserves particular attention: *second-person interactions*, in which I treat a fellow member of society as an individual like myself. From such interactions arises something truly irreducible to individuals. The importance of the second-person standpoint in ethical life has been emphasized by Stephen Darwall (2006), and of course long ago by Martin Buber (1923), but it is also crucial in epistemology and the philosophy of science. There is a deep presumption underlying any second-person interaction: you are an individual, with the basic rationality that consists in the

[20] This is a term I once absorbed from Jungian psychology, but it does not need to be placed in that precise context. It is related to what other psychologists have called self-actualization or ego development.

https://doi.org/10.1017/9781108635738.002 Published online by Cambridge University Press

coherence of belief and action, and the basic cognitive capacity required for communication, also with at least a minimum degree of good will and conviviality. We often lose sight of second-person interactions in the usual clash between individualistic and social perspectives, but so many of the common speech acts in life are of the second-person variety: commands, questions-and-answers, assertions-and-(dis)agreements, arguments, and explanations. Many key epistemic activities, too, are inherently in the second person. Describing or explaining is meaningless and pointless unless there is at least some imagined 'you' to whom it is directed. In philosophy and science we talk too often about questions and answers without thinking about the second-person interaction of asking-and-answering. Even more problematic is the divorcing of 'arguments' from *arguing*, which is something you and I engage in, and the removal of 'justification' from the persuasion of me by you. It will not do to reduce all of these things to the flatland of propositions and their logical relations.

A social group develops a thicket of second-person interactions. What Wilfrid Sellars called 'we-intentions', adapted later in different ways by Raimo Tuomela and John Searle, must emerge from you-and-I interactions.[21] In this thicket of interactions people establish customs, routines and institutions. We now have a higher level of sociality, composed of individuals who consciously connect with each other in order to live better together. Society, so formed, shapes individuals in turn, in a process similar to the earlier-stage shaping of individuals through inarticulate socialization. When patterns of social interactions become settled, the social norms, routines and expectations may become so internalized as to no longer require explicit reinforcement. If so, the individual becomes submerged into society again – as a willing and comfortable master of unspoken rules, or a timid conformist, or a half-comprehending misfit. What was once consciously negotiated becomes *sedimented* (Husserl [1954] 1970; Føllesdal 2010) – that is, added to the stock of shared unarticulated culture. Such a process happens in science, too: think of people who use thermometers, clocks or pH meters with no thought as to how the standards are established and maintained.

This formation of higher-level society by conscious agreement is still not the end of the story. A higher level of individual struggles to emerge out of that newly sedimented sociality. This cycle of socialization and individuation can continue indefinitely as societies and individuals become ever more complex and sophisticated. And a backwards glance also shows that the initial

[21] See Schweikard and Schmid (2021), esp. sec. 2.3, on Sellars's and other related ideas.

social dimension with which I started my discussion was not the very beginning of the process. There were individual animals before then, each of which was shaped in a proto-social environment; and so on.

1.3 Epistemic Activities and Systems of Practice

To facilitate the philosophical understanding of active knowledge, we need to craft good *units of analysis* that can accommodate aspects of knowledge and inquiry that are not captured by propositions. For this purpose I propose the notion of **epistemic activity**. An activity is not a one-off act, but a programme of action designed for the achievement of a recognizable aim. An activity has an inherent aim that partly defines it, and also various external functions. Various activities can be pulled together in order to form an integrated activity, with a new overall aim. However, the relationship in such integration is not reductive, and the contributing activities are not necessarily simpler than the overall integrated activity. When we are considering extensive and complex fields of work such as science, the most important unit of analysis is a **system of practice**, which is a coherent network of activities comparable to a Kuhnian paradigm. A system of practice, unlike an integrated activity, has multiple aims. Activities and systems both exhibit operational coherence, a concept that will be characterized in more detail in Section 1.4.

Activities and Practices

The usual unit of analysis in epistemology is a proposition, or a theory conceived as a collection of propositions, but these are not appropriate for active knowledge. In order to make full sense of the notion of knowledge as ability, we need suitable notions of action and practice. What exactly is doing, and what are those somethings that we do? Before I try to explain my own concepts of epistemic activities and systems of practice, I should comment briefly on some relevant bodies of work in the philosophy of action, and in theories of practice in the social sciences and in science and technology studies.

Within the literature in the philosophy of action I look particularly to the work of Jennifer Hornsby, who emphasizes two points that are fundamental to my thinking. First, she stresses the importance of attending to *activities* (Hornsby 2007a, p. 170): 'Human agents participate not only in once-off [sic] actions, but also in activities.' An activity is a complex, organized and regulated series of doings. While an individual act may be performed in a haphazard way, to call something an 'activity' implies routinized and repeated doings directed at an aim, following a reasonably stable set of rules and norms. Second, Hornsby focuses on capacities and on knowledge-as-ability: 'a person's knowledge of how to do things informs more than token actions', and 'much human agency is made possible by people's possession of capacities which ensure that they have standing abilities to engage in one or another activity' (ibid., p. 179).

There is a large and diverse body of literature theorizing about the nature of practices.[22] One question to be addressed before I go on to elaborate my notions of 'epistemic activity' and 'system of practice' is whether I shouldn't just speak about practices, which would also help connect my thoughts more straightforwardly to existing discourses. One difficulty is that people have meant all sorts of different things by 'practice', so it is impossible to be precise about it without further specification. The best attempt at a concise and fair summing-up that I have found is by Theodore Schatzki (2001a, p. 11), who states that practices are commonly conceived as 'arrays of activity', or more specifically, as 'embodied, materially mediated arrays of human activity centrally organized around shared practical understanding'. In that sense a practice is akin to what I call a system of practice, but in other uses a practice designates something more akin to an epistemic activity.

Rouse (2001, p. 199) makes an important distinction between mere regularities in people's behaviours and normative structures that govern people's behaviours. Only the latter deserve to be called 'practices': 'actors share a practice if their actions are appropriately regarded as answerable to norms of correct or incorrect practice'. I follow Rouse in adopting the normative view of practice. According to Schatzki (2001b, pp. 60–1), the normativity of practice is a 'teleoaffective structure', which combines a specification of ends and values. I also follow Rouse, Schatzki and others in emphasizing the place of understanding in practices, of which I will speak further in Section 1.4. Another important dimension of practice to note, especially for the purpose of philosophy of science, is the fact that most of our practices involve engagement with objects other than ourselves. As Rouse (2001) and Karin Knorr Cetina (2001) both stress, this objectual engagement brings open-endedness, contingency

[22] As my principal guides I follow two collections: *The Practice Turn in Contemporary Theory* (Schatzki et al. 2001), and *Science after the Practice Turn in the Philosophy, History, and Social Studies of Science* (Soler et al. 2014). The editors' introduction to the latter volume provides a masterful and comprehensive survey of the history and the main characteristics of the 'practice turn' especially as it pertains to science studies. Also very instructive are the accounts by Michael Thompson (2008) and David G. Stern (2003). Within the philosophy of science, a very important yet sadly neglected work in this direction is Harré and Llored (2019).

and creativity to practices. All of these are lessons that inform my conception of epistemic activities and systems of practice.

Epistemic Activities

In my work in the history and philosophy of science I have been trying to put into practice the 'activity-based analysis' of scientific knowledge (Chang 2011a; 2014). For the sake of consistency and continuity, I will begin by quoting from my best-publicized previous attempt to characterize epistemic activities, from my book *Is Water H_2O?* (Chang 2012a, pp. 15–16):

> An *epistemic activity* is a more-or-less coherent set[23] 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). An important part of my proposal is to keep in mind the aims that scientists are trying to achieve in each situation. The presence of an identifiable aim (even if not articulated explicitly by the actors themselves) is what distinguishes activities from mere physical happenings involving human bodies . . .

There I noted a similarity with Soler's notion of an 'argumentative module', which is 'individuated and defined as a unit on the basis of its aim: on the basis of the question it is intended to answer, or the problem it tries to solve' (Soler 2012, p. 235).

I will now develop the concept of epistemic activity further, starting with some simple modifications. The notion of 'epistemic' should be broadened to include the evaluation and use of knowledge, as well as its production and improvement. And it should be stressed that most of the time mental and physical activities are combined together. In addition to mental and physical activities, Bridgman (1959, p. 3) also noted the importance of 'paper-and-pencil operations' in science; similarly, Ursula Klein (2003) has stressed the importance of the use of 'paper tools' in chemistry.

One major issue to be worked out is the different levels of aims. Specific activities are designed to achieve proximate aims that help agents achieve their ultimate desiderata. Suppose someone asks me, while I'm striking a match, what I am trying to achieve. My answer may be 'I'm trying to light a match', or 'I'm trying to get a combustion-analysis of an organic compound

[23] In a slightly earlier formulation (Chang and Fisher 2011, p. 361) I used the word 'system' instead of 'set'; that usage is avoided here, as it is not consistent with the explication of 'system of practice' that I have now made.

going.' Both are cogent answers, but they get at two different kinds of aim. The first answer addresses what I will call the **inherent aim** of the activity: getting the match to light is the whole point of the activity itself, regardless of why one is engaged in that activity – that may be in order to light a Bunsen burner with it, or to burn down a house, or just to watch and admire the marvellous phenomenon that combustion is. These latter reasons might be called the **external functions** of the activity, as they are consequences of the successful execution of that activity. An activity is *defined* partly in terms of its inherent aim, which exists regardless of any external functions that the activity may serve (match-lighting is not match-lighting if one does not at least intend to light a match).

Finally, I should stress that any description of an activity we can give is a *programme* of action, whether in terms of retrospective understanding or as prescriptive guidance. Such a programme is bound to be abstract, in the sense of not including all the features that are present in each instance of its execution. So, any activity that we can describe is not precisely instantiated in our actual doings, and there is no uniquely right way to identify and classify activities out of the stream of doings that we continually carry out in life. In this sense, epistemic activities (and also systems of practice) are 'ideal types' in Max Weber's sense of the term. An ideal type is a concept derived from observable reality but with conscious simplification and exaggeration. Weber says: 'an ideal type is formed by the one-sided accentuation of one or more points of view' according to which 'concrete individual phenomena . . . are arranged into a unified analytical construct' (quoted in Kim 2019). As Sung Ho Kim explains (2019, sec. 5.2): 'Keenly aware of its fictional nature, the ideal type never seeks to claim its validity in terms of a reproduction of or a correspondence with reality. Its validity can be ascertained only in terms of adequacy.' As I will discuss further in Section 1.4, there is a hermeneutic dimension to activities: the operational coherence of an activity is about pragmatic sense-making, on the part of either the agents themselves or others who analyse their actions.

Systems of Practice

Epistemic activities normally do not, and should not, occur in isolation. Often they form a network that is dense enough and large enough to deserve to be called a 'system of practice'. To refer back to my 2012 publication again:

> A *system of practice* is formed by a coherent set of epistemic activities performed with a view to achieve certain aims . . . Similarly as with the

> coherence of each activity, it is the overall aims of a system of practice that define what it means for the system to be coherent. (Chang 2012a, p. 16; also Chang 2014, p. 72).

I now want to consider more carefully how activities come together coherently to constitute a system of practice. (Any kinds of activity may come together to form a system of practice; I tend to speak of epistemic activities since I am mostly addressing issues relating to knowledge.) The various activities in a system of practice are not merely performed in the same general setting; rather, they come together in very particular ways to meet certain aims. The system coordinates various activities for the satisfaction of the system-level aims.

Unlike an activity, a system of practice does not have one single inherent aim. In fact, that may be considered the chief difference between an activity and a system. Recall the example introduced in Section 1.1: the chief overall aims of the Lavoisierian system of chemistry were the knowledge of the composition of various substances, a good classification of all chemical substances, and the explanation of chemical reactions. There is both dependence and independence among such aims. Lavoisier's was a 'compositionist' system of chemistry (see Chang 2011b; 2012a, sec. 1.2.3): knowing the compositions was essential for the other two aims of explanation and classification, because in this system there was a commitment to make explanations and classifications on the basis of compositions. On the other hand, classification and explanation were largely independent aims, in the sense that one could be pursued without the other at least to an extent.

To take a different kind of example: consider the game of soccer – not an individual match, but the 'game' as an institution. This may be considered a system of practice. The whole institution of soccer does not have a unitary inherent aim, so it is not a single activity, while particular activities *within* soccer have unitary inherent aims: for example, the inherent aim of goal-keeping is to prevent the other team from scoring, and the inherent aim of passing is to give the ball to another player on one's own team. But isn't *winning* the inherent aim of the game itself? That may be said to be the aim of each team's engagement in a match, but not of the whole sport, to which 'winning' does not apply. If we ask seriously about the aims of the whole system of soccer, there is no one clear answer, but multiple answers: to provide entertainment for the people, to make income and profits for some individuals and entities, to promote health and fitness in society, to contribute to community solidarity, etc.

The Coordination of Activities

Now, there is something similar between how different activities come together to serve system-wide aims, and how an activity is formed by the coordination of many different doings. More thought is needed on the coming-together of doings. When I previously characterized an activity as 'a set of operations', I did not define what an 'operation' was, thinking that it was innocuous enough as something similar to 'activity' but as a convenient term for sub-components of an activity (Chang 2014, pp. 74–5). More problematically, I implied an overly atomistic picture, as pointed out in Léna Soler and Régis Catinaud's (2014, p. 82) cogent critique of my ideas. I believe that I have now found a way of avoiding the atomistic perspective.

A central puzzle is that even the simplest activities seem complex and susceptible to further analysis in terms of constituent activities, with seemingly no end to such analysis. For example, take one activity from the Lavoisierian system of chemistry: combustion-analysis. This activity incorporates various other activities: setting something on fire, capturing the combustion-products (gases) by means of other chemicals, weighing the resulting compounds, and making percentage-calculations. But even those activities in themselves seem to consist of other activities. For example, the activity of weighing-with-a-balance includes the placing of samples and weights on balance-pans, reading the number off the scale, and certifying the weights used as correct standard weights. But the certification-activity, without which the whole activity of weighing lacks validity, is itself a very complex thing! It may consist in ordering the weights from a reliable supplier, or comparing them to a more trusted set of weights, or checking them against certain natural phenomena (e.g., the weight of a certain volume of water at a certain temperature, in the conception of the originators of the metric system). Whichever method we opt for, it seems clear that this 'component' activity of weight-certification is not going to be simpler than the whole activity of weighing-with-a-balance.

The analysis of activities is not an atomistic or reductionist enterprise. It isn't quite right to say that an activity is 'composed of' simpler activities, since 'composition' implies an atomistic ontology too strongly. I prefer to speak of *constituent* activities, rather than *component* or *elementary* or *basic* activities. There is no lowest level of description, no rock-bottom of atomic activities, and no clear end to the process of analysis. Yet in many situations we do gain useful insights from analysing an activity into its *apparent* constituents, and the analysis should be carried out wherever it is productive, and only and as far as it is productive.

So how is the incorporation of activities to form another activity done? A couple of social metaphors might be instructive here. The United Nations isn't simply all of its member nations put together, because the functioning of the UN requires shared institutions, routines and purposes, which do not exist within individual member nations. Nor is it the case that the member nations wholly belong to the United Nations. So the whole here is both more *and* less than the sum of the parts. Likewise, when individuals come together to form an association (Neighbourhood Watch, Alcoholics Anonymous, or what have you), the association is not just the individuals put into a set, and each individual member has a very complex life outside the realm of the association, and e also belongs to a number of other associations. It is similar when activities come together to form another activity. In order to bring the constituent activities together, something outside those activities must be imposed in order to connect them; the most important factor is the overall aim of the integrated activity, which the constituent activities are brought together to serve.

Now I can attempt to present an overview of the ontology of epistemic activities and systems of practice. An activity is identified and individuated by its operational coherence in relation to a unitary inherent aim (and may have various external functions). Typically, an activity can be seen to incorporate other activities into it, to serve its own overall aim while the constituent activities each retain their own inherent aims, too. The constituent activities, in turn, incorporate other activities. We may pursue the analysis of activities into their constituents indefinitely, but this is not a situation of infinite regress, because constitution here is not reductive. Rather, the overall picture is more of a reticular one, a network of a great number of activities, each one incorporating various others, reminiscent of the 'rhizome' structure that Gilles Deleuze and Félix Guattari ([1980] 1987) speak about. As for a system of practice, it is formed by the coordination of activities without a single overarching aim – if there were a unitary inherent aim, then the coordinated whole would constitute a single integrated activity. A system of practice has multiple system-level aims, and each of those aims is not locatable within any one of the activities involved. The coherence of a system of practice consists in an effective coordination of the external functions of various activities for the achievement of system-level aims. How well the different aims within a system go together is also a matter of operational coherence.

1.4 Operational Coherence

Operational coherence is a key notion in my account of active knowledge, and a core part of the definitions of epistemic activity and system of practice, as seen in Section 1.3. Now I will try to give a more precise and detailed characterization of this concept, especially because it will play a crucial role in the reconceptualization of the notions of reality and truth in Chapters 3 and 4. In previous publications I defined operational coherence in terms of the harmonious relationship among the operations constituting an activity. By grappling with some puzzles arising from that definition, I now arrive at a more fundamental view. In short, operational coherence consists in **aim-oriented coordination**. A coherent activity is one that is well designed for the achievement of its aim, even though it cannot be expected to be successful in each and every instance. Operational coherence is based on pragmatic understanding; it consists in doing what *makes sense* to do in specific situations of purposive action.

What Is Operational Coherence? Intuitions and Illustrations

In the overview of this chapter (Section 1.1) I put forward operational coherence as a chief criterion for assessing the quality of active knowledge. In Section 1.3, I presented it as a key characteristic of epistemic activities. I will now give a more detailed and considered characterization of operational coherence. Let me start by motivating the concept again with some illustrative examples. Operational coherence is a pertinent concept in all sorts of activities, scientific and quotidian. In daily life we employ literally thousands of simple skills that require a good coordination of bodily movements, material conditions and mental concepts: drinking a glass of water, tying shoelaces, eating with chopsticks, riding a bicycle – or walking up the stairs, which is quite an achievement as contemporary robotics has learned. Take a very simple activity like match-lighting (which was, incidentally, so essential to the progress of chemistry, and even physics, for so long!). 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, in which I hold the match tightly, just so; I pull the head of the match across the rough strip on the box (and I break the matchstick – no, no, the correct move is to push it), at an appropriate angle and at the right speed with some abruptness; I stop the movement of that hand promptly once the flame comes on. I need to bring these operations together well enough for my match-lighting activity to be coherent. It is important to keep in mind that coherence doesn't pertain to a single act, but to a sustained and organized *activity* (and even a whole system of practice).

The same kind of coordination takes place in scientific and technological practice, only with more theorized, complicated and careful actions. An extreme case is the operation of the global positioning system (GPS), as discussed by Peter Galison (2003, pp. 285–9). GPS requires an intricate coordination of a range of material technologies (geostatic satellites, atomic clocks, electromagnetic signals, mobile phones, etc.) and various abstract theories (Newtonian mechanics to fly the satellites, quantum mechanics to run atomic clocks, and both special and general relativity to make corrections to the atomic-clock readings depending on the speed of the satellites and their locations in the earth's gravitational field). Each element of this enormously complicated set-up is carefully coordinated with the other elements, to enable a marvellous degree of operational coherence in the activities we undertake by means of it. And even in this theory-heavy set-up, it is not the case that the operational coherence follows from a single unified theory (most fundamentally, we have no unified theory of quantum gravity, not to mention a theory that encompasses both the theoretical physics and the engineering systems involved in GPS). Rather, the coherence is achieved in a highly ad hoc manner, applying selected aspects of various theories to different parts of the system in a judicious way designed to achieve the specific aims at hand.

In puzzling out what operational coherence is, it is also helpful to think about what happens when coherence is lacking. If I try to drink water 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 well we normally maintain the coherence of our bodily movements and adapt them to our external surroundings without even thinking about it. Another example, from the typical social life of a professional (in the days before the Covid-19 pandemic): you go to a conference, meet a colleague that you really like but don't know very well; you offer a warm handshake, your colleague offers a discrete hug or kiss; your greeting ends in an incoherent tangle. Incoherence may sometimes be traceable to false or mutually contradictory beliefs, but ineptitude of belief is certainly not the only reason for it. Incoherence can also arise due to the lack of basic capability (starting with weak eyesight or muscular strength), the use of inappropriate materials, poor timing between different operations, the application of mutually conflicting rules, and so on. And the examples just cited should make it quite plain that operational coherence is a matter of degrees, and not a precisely quantifiable one at that. It is necessarily a less precise concept than consistency, which is well defined through logical axioms.

Before going on further, I want to anticipate a common worry: the so-called coherence theories of truth and justification in epistemology may slide into relativism, idealism or constructivism. If coherence is simply a matter of a positive mutual relationship between our beliefs without anything else to ground any of them, then there is a legitimate concern that coherence does not provide any link between knowledge and reality. 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. James O. Young (2015) notes that more plausible versions take the coherence relation as 'some form of entailment' or 'mutual explanatory support between propositions'. Similarly Richard Foley (1998, p. 157) says, in relation to justification: '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.' Catherine Elgin (2017, pp. 71–3) also has a liberalized notion of coherence, laying out the minimum requirement that 'the components of a coherent account must be mutually consistent, cotenable, and supportive', and also considering the relations between multiple levels of commitment. However, even in the more sophisticated versions of coherentism, the problem of circularity remains.[24] *Operational* coherence is a wholly different matter. It cannot be achieved arbitrarily by decree, wishful thinking, or mere agreement among beliefs or people. On the contrary, in order to do things coherently we need to have an understanding and mastery of our surroundings. Operational coherence carries within it the constraint by nature. Through operational coherence the world outside the control of the mind is brought to bear on knowledge. In fact, in Chapter 2 I will argue that operational coherence is the only means by which reality can shape our practices, and in Chapters 3 and 4 I will show how operational coherence can ground the very notions of reality and truth.

Three Puzzles

In an earlier publication I defined 'operational coherence' as follows:

> 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

[24] Note, however, Thagard's point (2000, p. 77) that circularity is not necessarily regress: 'Coherence-based inference involves no regress because it proceeds not in steps but rather by simultaneous evaluation of multiple elements.'

> a coherent activity is successful, *ceteris paribus*; the latter condition serves as an indirect criterion for the judgement of coherence. (Chang 2017b, p. 111)

This definition of operational coherence is not particularly wrong, but it needs further development and reorientation. Since the old definition is out there in print, I feel obliged to show how my current ideas evolved from there, rather than just presenting the formulation that I have now reached. And aside from doing penance for having published a half-baked definition, there may actually be some benefit in showing how the recent changes that I have made were motivated.

My old definition of coherence left three puzzles unsolved. First, there is a problem with conceiving operational coherence as a relationship among the operations that constitute an activity. As Soler (2009, ch. 9; 2012) has convinced me, it is unhelpfully constraining not to allow ourselves to think of coherence in terms of the harmonious relationship between different *types* of aspects or elements of an activity (such as theoretical assumptions, bodily abilities, perceptions, social constraints, and properties of our tools). Still, it is difficult to think cogently about how such a heterogeneous set of things relate to one another, which is what had originally pushed me towards the ontological homogeneity of dealing only with operations. But that solution only masked the difficulty in any case, because making sense of the interaction between operations is not trivial after all, and understanding how *each operation* works requires making sense of the interactions between different types of elements within it.

Second, what exactly does it mean for a relationship to be 'harmonious'? Harmony is a musical metaphor.[25] I have also talked about how actions 'fit together', but that is a mechanical metaphor, and actions are not parts of a machine, any more than they are musical notes.[26] I left this question unresolved in my earlier publication. I confessed:

> 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. (Chang 2017b, p. 110)

[25] However, Liba Taub tells me that the musical notion of 'harmony' was probably itself a metaphor in the original Greek usage, drawing on the idea of things fitting together, as with the planks of a ship. There is another layer of meaning that Neurath's boat can take on, then!

[26] In Chapter 2 I will be criticizing the problematic uses of metaphors in 'correspondence realism', so I should be careful in wading into metaphors of my own!

But I was not quite satisfied with taking 'harmonious' as a primitive.

Third, what exactly is the relation between the success of an activity and its operational coherence? In my old paper I refrained from equating coherence and success, especially in order to allow for the possibility that a coherent activity could still fail due to some accidental extrinsic circumstance (and conversely, that an incoherent activity could succeed by accident). For example, I may do all my match-lighting operations sensibly, but be foiled by an unexpected gust of wind, a mischievous friend pouring a bucket of water all over me, or any number of other possible mishaps. If we can demarcate well enough the match-lighting activity itself from extrinsic accidents, then it would make sense to say that my match-lighting activity *is* coherent, but may occasionally be unsuccessful due to circumstances. But why exactly is it that operationally coherent activities should tend to be successful? What kind of mechanism or causal path might be involved in the production of success from coherence? I left this issue unresolved, too.

Coherence as Understanding and Coordination

Now, as it turns out, all of those puzzles about the meaning of operational coherence have one common solution, the germ of which is contained in the following statement from my old paper: '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' (Chang 2017b, p. 111; emphasis added). That is to say, operational coherence is a hermeneutical notion, concerning a pragmatic kind of understanding. What is *operationally* coherent is what makes sense for us to *do*, and 'sense' here is framed by our aims. But what does *sense-making* have to do with success? Surely I can't be suggesting that if an activity makes sense to me it will tend to be successful? The success of an activity is not *caused* by its coherence; rather, the coherence of an activity *consists* in doing what is sensible to do if one wants to succeed. Coherence is *design* for success, and that design is based on empirical learning: it makes *sense* to do what we think will succeed, and it doesn't make sense to do what we think is unlikely to succeed. Coherent activities are carefully designed so that they *would* work. The coherence–success relation is not one of cause and effect, but a hermeneutic–pragmatic act of sense-making in the context of purposive behaviour.

But what exactly does 'making sense' mean? There is nowadays a sizeable literature in the philosophy of science on the nature of understanding, to which I want to make links in the end.[27] However, initially I must set out in a different direction. This is because most of the extant literature is about the understanding of natural phenomena, or about our understanding of scientific theories. For my purposes here, it is necessary to consider first of all the understanding of our own actions, because that is what lies at the heart of operational coherence; this is what Oscar Westerblad (forthcoming) calls 'operational understanding'. In that case sense-making has to be approached partly from a psychological angle, as in Paul Thagard's much-neglected work on coherence. Thagard focuses on what the mind can hold together happily without too much dissonance, taking coherence as a relation between various types of elements: concepts, propositions, parts of images, goals, actions, etc. These elements can fit together (or not) through a variety of 'coherence relations', which may 'include explanation, deduction, facilitation, association, and so on' (Thagard 2000, p. 17). Using these relations we can make sense of how the different elements of our activities work together.[28] I should also look to the field of hermeneutics for a deeper understanding of understanding, but the traditional focus there is on the understanding of linguistic texts, so the connection to the understanding of actions will be indirect.

My own direct approach starts by taking operational coherence as something to do with rational action. If we start with the standard notion of instrumental rationality (means–end relationship), what makes sense for us to do is whatever will facilitate the achievement of our aims. But the usual treatments of instrumental rationality tend to exclude the hermeneutic dimension, seeing the means–end relation as basically causal. In my own initial thinking I was also projecting operational coherence onto the material dimension of actions, and that is why I could only talk in terms of metaphors. Things 'fitting together' is meaningless, unless there is a purpose under which the fit is judged; common images like the planks of a ship fitting together are deceptive, because in those cases the purpose (making the ship watertight) is taken for granted and not mentioned. The coherence of an activity is not some mysterious harmony between things in themselves, but it is a matter of how

[27] See esp. De Regt, Leonelli and Eigner (2009); De Regt (2017); Grimm, Baumberger and Ammon (2017); Grimm (2018); Stuart (2018).

[28] Thagard views coherence primarily as a matter of constraint-satisfaction, but I think many of his insights can be reworked in the direction of understanding. Interestingly, what is most akin to my thinking on operational coherence is Thagard's view of 'deliberative coherence' (operative in the realms of ethics and politics): 'if an action facilitates a goal, then there is a positive constraint between them' (Thagard 2000, p. 127). As for the problems of 'truth' and 'epistemic justification', Thagard considers them only to involve propositions as elements of coherence (ibid., p. 25, table 1). So his epistemology is not designed to deal with active knowledge, but I believe that it can be adapted to do so.

we bring together things and actions in order to achieve our aims. That is the sense in which operational coherence is **aim-oriented coordination**.

Some brief examples will be helpful in illustrating the pragmatic sense-making involved in operational coherence. There is a cartoon from *The Far Side* by Gary Larson showing a young man sitting up in bed in the morning looking attentively at a sign on his wall: 'Pants BEFORE shoes.' The joke is about someone needing to write that down to remind himself, but in all seriousness we do all have such rules lodged in our heads and in our bodies. Putting your shoes on before your pants[29] doesn't make pragmatic sense; you know that if you've tried it. Trying to go in opposite directions simultaneously, or trying to be in two places at the same time, certainly doesn't make sense for us, though you may think a bit differently if you are dealing with electrons in the quantum mechanical double-slit experiment. Trapping ultra-cold atoms with a laser is a coherent activity, because the physicists have learned to understand the conditions and operations that enable this feat. I may practise archery coherently, based on a certain sense of my own strength, of the properties of the bow and the arrow and the surrounding air, of the location of the target, and of the basic laws of mechanics.

To understand coherence, again it also helps to think about cases of incoherence. Some incoherence can be purely mental. In colloquial usage we often say 'incoherent' when someone is talking gibberish – not understandable, 'not even wrong'. But it is more interesting and informative to consider activities whose incoherence consists in a mismatch between how we think, what we want, what we do, and the way things are. Return for a moment to the example of archery: if the arrow does not hit the target, then that is a failure. Now I may say: oops, my hand slipped as I was stretching the bow; or, damn, there was this sudden gust of wind that I hadn't expected; and so on, to give myself and others an *understanding* of the failure. So, I can maintain the coherence of my activity in the face of isolated failures because my understanding of the whole activity remains intact. But if I keep missing the target completely and my failures are inexplicable, then it is incoherent to keep doing what I am doing. Then I must investigate, in the way that Peirce and Dewey say a disturbed situation gives rise to inquiry. I need to start thinking and doing things differently. Maybe I have to pull the bowstring harder, or revise the way I assess the amount of force exerted, or get my eyes checked to see if I am seeing the target well enough, or even adopt different laws of mechanics. If

[29] Larson means 'pants' in the American sense, meaning trousers. But the guideline would make even more sense with the British meaning of 'pants' (underpants)!

https://doi.org/10.1017/9781108635738.002 Published online by Cambridge University Press

I keep doing what doesn't work instead of making such adjustments, then my activity is incoherent.

With the notion of operational coherence now framed in terms of pragmatic understanding, it becomes clear how the three puzzles that plagued my old definition of operational coherence can be solved. First, pragmatic sense-making can accommodate the coordination of heterogeneous types of elements, built around the central relationship between means and ends. Second, the elusive and metaphorical sense of 'harmony' is now reduced to the sense of pragmatic understanding. And finally, there is a positive relationship between the coherence of an activity and its success because we *write into* the design of a coherent activity our sense of what would succeed. We are safe from being 'untethered from reality' as long as our sense-making maintains a commitment to empiricism. Pragmatic understanding is not just 'in the head'. But what if people refuse empiricism itself, and insist that ignoring the lessons of experience *makes sense to them*? No amount of objectivist epistemology or self-righteous condemnation will stop such people; we can only win them over patiently by showing them the fruits of empiricism in the long run.[30]

Before I close the discussion of operational coherence, I should briefly address the difficulties involved in attributing coherence to other people's activities works. This is part of a general issue, which I have already touched on in Section 1.3 in relation to the identification of activities. I would argue that we are justified in interpreting the activities carried out by other agents as operationally coherent, even if they do not themselves articulate the aims and the coordination involved in their activities. Historians, including historians of science, confront this issue on a regular basis, as the past people who did not leave evidence of reasons for their actions cannot be further interrogated. It is a meaningful and instructive exercise to attempt an understanding of the activities of non-articulating agents as operationally coherent. I may be an 'idiot savant' of archery, who just somehow knows how to hit the target without much thinking or conscious planning; however, the astonished observers would attribute coherence to my activity if *they* could understand how my operations make sense in relation to the aim of hitting the target. This is not fundamentally different from how we deal with the general problem of other minds in practice – I make the decision to regard you as a conscious

[30] In not appealing to an absolute standard or authority for understanding, my view may be considered a relativist one. But relativism in the sense of rejecting absolutes is not a crude and bankrupt doctrine, as the imposing collection of recent works edited by Martin Kusch (2020) makes abundantly clear.

knowing agent with experiences that are similar to mine, even though I have no direct access to your experiences. This is just a reasonable way of living, based on a productive kind of respect.

The same may be done with animals. It is more natural to say that bees know how to collect nectar from the flowers and tell each other where the good flowers are, than to insist that they couldn't possibly possess knowledge because (we presume) they do not formulate beliefs in their tiny brains. I do think that spiders know how to build webs, squirrels know how to hide food (and even how to find it again, sometimes), dolphins know how to work together to corral fish, and migratory birds know how to navigate through thousands of kilometers of terrain with astonishing reliability. And AlphaGo surely knows how to play the game of go better than any human player. The attribution of knowledge in all these cases is based on our understanding of the agents' behaviours as operationally coherent activities. In daily life we rate such attributed knowledge highly and stand in awe of it. Why should epistemology ignore it? For my present purposes, however, it is enough that we make attributions or coherence, aims and knowledge in dealing with most humans and some very clear cases of other thinking beings. It is not necessary for me to engage in debates such as whether thermostats can be said to have consciousness and experience (Chalmers 1996, and responses to it).

1.5 Inquiry as Aim-Oriented Adjustment

In this section I take a dynamic view on active knowledge. That means looking into the nature of inquiry, the process through which we actively attempt to acquire and improve knowledge. I will elaborate on the idea that inquiry is the business of increasing operational coherence, a process of aim-oriented adjustment. In fully unrestricted inquiry, any aspect or element of the unsatisfactory initial situation facing the inquirers may be altered in order to bring about better operational coherence in their activities. Not only specific beliefs but capabilities and methods, and even aims themselves may be changed and improved in the process of aim-oriented adjustment. Various types of more restricted inquiry arise when certain elements are fixed. Starting from an examination of unrestricted inquiry helps us see the full range of processes of aim-oriented adjustment. Here I pay particular attention to inquiry processes that are usually neglected, including the crafting of concepts to handle novel experimental phenomena, the creation of new theoretical frameworks, and the adjustment of aims.

Aim-Oriented Adjustment

So far I have elaborated on the idea of active knowledge, whose quality consists in the operational coherence of epistemic activities. Now I want to take a dynamic view on the development of knowledge, and ask about the nature of inquiry. I start with a simple perspective: if good knowledge resides in operationally coherent activities, then the improvement of knowledge consists in the enhancement of operational coherence. Here I want to go clearly beyond the narrow image of inquiry as fact-finding, reorienting it and broadening it out into a picture of coherence-making. In Section 1.4 I stressed that operational coherence should be understood as a matter of pragmatic understanding; accordingly, I want to recognize the dimension of understanding in the process of inquiry, too. This view of inquiry is consonant with Peirce's and Dewey's mentioned in Section 1.1.

In Section 1.1, I stated that the enhancement of operational coherence was achieved through the process of **aim-oriented adjustment**, and I will now explain that notion in more detail. Let's start with a mundane example. Consider learning how to hammer a nail – how to make a tight grip on the handle, how to hit the nail on its head and how to recognize when you've done that, how to adjust the strength and number of the hammer-blows depending on the kind of wall, the kind of nail and other circumstances. Learning how to do this activity takes place through a process of trying out whatever it takes to improve the operational coherence of our activity in relation to the achievement of its aim. Generally, there is no pre-existing recipe for solving a real-life problem (rather than the kind of 'problem' that is an exercise laid out by the teacher who knows the answer already). Metaphorically, sometimes even literally, we have to twist and turn – try this and that, again and again, until something 'clicks' (or, until we settle into a comfortable and effective routine). This is quite similar to what Ludwig Fleck calls 'tuning'. As Andrew Pickering (1995, p. 121) puts it: 'The scientists tried varying the prototype recipe [for the Wassermann test for syphilis] in all sorts of ways and eventually arrived at a recipe that was medically useful.' We reach a bodily-and-conceptual understanding of the task as we go. Even learning to *see* is like this, and likewise for other modes of perception, too, as Alva Noë (2004) has vividly emphasized. Learning various second-person interactions, whose importance I emphasized in Section 1.2, also works largely by such aim-oriented adjustments: handshakes, explanation, co-authoring, telling jokes, boxing, ballroom dancing and moving furniture are just a handful of randomly selected examples.

It is instructive to revisit the classic example of learning to ride a bicycle, so memorably discussed by Polanyi. Initially the novice does not know how to keep himself from falling, and doesn't understand how that is done. The helpful older sister gives him tips like 'turn into the direction in which you are beginning to fall', but this advice makes no sense to the boy in the abstract, and when he tries to put it into practice, it doesn't work. However, trial and error eventually shifts something in the brain and the muscles, and he is able to ride, wobbly as he may be; at this stage the thing about turning into the direction of falling starts to make that conceptual–bodily sense. As his skill improves, he also begins to understand things like how a slight turn can be achieved by the shifting of body-weight without turning the handle. Through such learning-how-to improvements, his bicycle-riding continues to increase its operational coherence. If I ask myself whether what I am doing makes sense while riding my bicycle, I would say 'yes' in both ways: the way I move my muscles and shift my weight around makes inarticulate sense to me as I do it, and I can also understand what I'm doing by articulating it and putting it into practical rules like 'shift the body weight to make a slight turn'.

Unrestricted and Restricted Inquiry

In order to get a general sense of how inquiry works, I believe we must begin with a view of unrestricted inquiry, only afterwards asking how its character is modified when particular restrictions are placed on it. It is important to keep in mind that any elements of the problematic situation that gives rise to inquiry may in principle be revised or discarded, and new elements may also be introduced. The open-endedness of unrestricted inquiry is typical of what Kuhn ([1962] 1970) vividly described as 'extraordinary research', which often brings about revolutionary change. In such inquiry everything is subject to change, including presuppositions, methodology, aims, criteria of judgement, and the list of significant problems. In contrast, much of what is normally recognized as inquiry (or research) is of a much more restricted type (as in what Kuhn considers 'normal science'), and can only take place on the basis of some unrestricted inquiry that has previously been carried out. As an example of a most restricted type of inquiry, consider what we might call fact-gathering. Here the aims and questions and methods are all fixed, except for the actual contents of the blanks to be filled by well-regulated acts of observation. Fact-gathering is the simplest kind of inquiry that generates propositional knowledge as an outcome, but it would be a mistake to regard it as the prototype or paradigm of all inquiry. On the contrary, the successful execution of *unrestricted* inquiry is the foundation of all cognitive activity, because outcomes of unrestricted inquiry lie at the foundation of language-use, mathematics, experimental design, causal reasoning, theoretical explanations, and almost all other aspects of intelligent life. In more restricted types of inquiry so much is already settled, and the truly challenging and exciting unrestricted stages of the inquiry process tend to be hidden from view.

Even though Kuhn was wrong to imply a sharp dichotomy between 'normal' and 'extraordinary' research, his distinction still makes perfect sense if we think of it as pointing to the two ends of a spectrum. Where I really want to depart from Kuhn's perspective is in regarding 'normal' science as the *normal* state of science, or even definitive of science.[31] On the contrary, various kinds of restricted inquiry are only results of restrictions that are placed on inquiry as temporary and local expedients. Although in popular imagination Kuhnian extraordinary science is only associated with field-changing scientific revolutions, Kuhn himself did acknowledge that extraordinary research can happen at any scale. For example, the discovery of X-ray was a small-scale change but the process had the same character as the large-scale revolutions (Kuhn [1962] 1970, pp. 92–3). Kuhn's conception of normal science as 'research under a paradigm' expresses very well the restricted and well-prescribed character of much of mainstream inquiry in a scientific field. In contrast, extraordinary research is what scientists engage in when they feel that it is necessary to depart from the ruling paradigm in order to solve an urgent puzzle.

Recognizing unrestricted inquiry as the basic form of inquiry allows us to recognize that inquiry results in the development of all aspects of active knowledge, many of which have been neglected by philosophers. Within philosophy of science, attention has been limited mostly to several specific well-controlled types of inquiry: fact-gathering, hypothesis-testing, classification, theory-construction, the construction of theoretical explanation, and the development of observational methods. In the rest of this section I want to give some attention to some important aspects of inquiry that are not often discussed by philosophers of science, and bring out and highlight the pertinent processes of aim-oriented adjustment.

[31] See Kuhn's statement in his legendary debate with Karl Popper: 'It is normal science, in which Sir Karl's sort of testing does not occur . . . which most nearly distinguishes science from other enterprises. If a demarcation criterion exists . . . it may lie just in that part of science which Sir Karl ignores . . . In a sense, to turn Sir Karl's view on its head, it is precisely the abandonment of critical discourse that marks the transition to science' (Kuhn in Lakatos and Musgrave 1970, p. 6).

The Making of Experimental Meaning

In experimentation, the most unrestricted inquiry takes place when new kinds of phenomena are discovered and an intense discomfort develops in not knowing how to make sense of them. Then a nearly desperate sort of aim-oriented adjustment takes place, to create new meaning in the new phenomena. A good example is how scientists two centuries ago tried to make sense of the electromagnetic effect discovered in 1820 by Hans Christian Ørsted (see Figure 1.1). A metallic wire is laid above a magnetic needle in a direction parallel to the needle; the magnetic needle turns when the wire is connected to a battery and a current of electricity flows through it. Factually, there was no difficulty concerning Ørsted's result – the phenomenon was clearly observed, easily replicated, and never seriously disputed. The challenge was how to make sense of it: why did the flow of electricity have a *magnetic* effect, and why did the electric current push the magnetic needle in a direction *perpendicular* to its

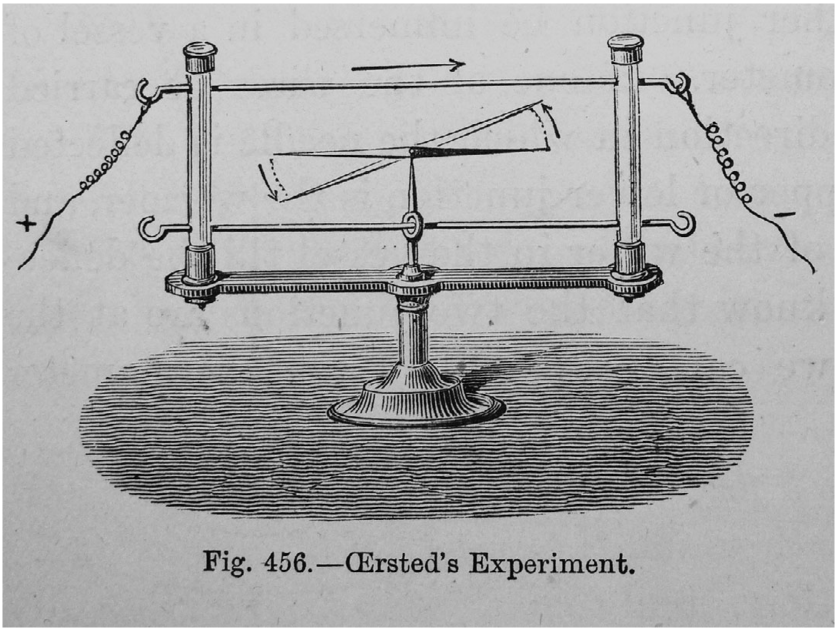

Figure 1.1 An experimental arrangement demonstrating the electromagnetic effect first discovered by Hans Christian Ørsted, from Privat-Deschanel (1876), p. 656, fig. 456.
Courtesy of the Whipple Library, Cambridge.

https://doi.org/10.1017/9781108635738.002 Published online by Cambridge University Press

own flow (rather than sweeping it along, as it were)? No mechanism for such an effect could be found within the dominant Newtonian scheme of forces between point-particles acting along a straight line connecting them.

In trying to understand how scientists tried to cope with this situation, I draw my lessons from Friedrich Steinle's ideas about 'exploratory experimentation', which were in large part developed through his study of how André-Marie Ampère and Michael Faraday developed their knowledge of electromagnetism (Steinle [2005] 2016). In exploratory experimentation, systematic high-level theory is absent and researchers start by focusing on discovering empirical regularities. The initial method employed is the systematic variation of all known experimental parameters that are presumed to be relevant. Sometimes this works out well, but when it does not, 'researchers consider the possibility that the failure might have something to do with deficiencies in the basic concepts of the field and thus feel encouraged to form and introduce new concepts that capture experimental findings' (ibid., p. 314). There is no algorithm to follow in that creative process, so the twisting-and-turning of aim-oriented adjustment enters the scene. Steinle points out that Ampère invented the concept of 'right and left of current' and that of a current circuit through his attempt to find the right concepts with which to express empirical laws about the electromagnetic effect.

Steinle builds on David Gooding's sadly underappreciated work on embodied agency and practical thinking in scientific work, which culminated in his book *Experiment and the Making of Meaning* (1990). Gooding follows how Faraday conceived concentric circles of magnetic influence around the current-carrying wire, in planes perpendicular to the direction of the current. Rather than adopting the approach dominant in France (led by Ampère and others) to reduce electromagnetic phenomena to Newtonian-style forces as much as possible, Faraday drew inspiration from his mentor Humphry Davy, who had placed magnetic needles on the plane perpendicular to the wire and saw them turn as to form a ring (Gooding 1990, p. 53). The same sort of ring pattern could also be shown with iron filings spread on a cardboard piece. Faraday began to see his 'lines of force' filling the space around electrically and magnetically active bodies; this idea later developed into the concept of fields in the hands of James Clerk Maxwell and others.

But through what process did Faraday come up with such concepts? Faraday's inquiry was a process of aim-oriented adjustment, employing various cognitive means to achieve the aim of explaining and extending electromagnetic phenomena. Based on his minute examination of Faraday's laboratory notebooks and other pertinent archival sources, Gooding concludes: 'Faraday is thinking *through* doing as well as *about* doing. Some of these thoughts are inherently ambiguous until articulated into configurations of real or imagined entities (images, models or concrete apparatus).' Faraday's creative thinking was a full blend of visual imagery, tactile sensation, reflections on laboratory observations, and the invention and use of new experimental apparatus. His early electromagnetic experiments were mostly *not* tests of well-conceived hypotheses: 'Faraday was experimenting to *realize* possibilities, not to *decide between* two distinct or incompatible interpretations' (Gooding 1990, p. 124; emphases original). Through continually evolving experimentation Faraday articulated previously unknown phenomena and created new meanings. At the same time he also created wondrous new apparatus like his famous rotation devices (see Figure 1.2), in which a vertically suspended current-carrying wire is made to rotate around the pole of a magnet underneath it (or a magnet is made to rotate around a wire), in a pool of mercury. Such devices served as embodiments of Faraday's new ideas. Through this line of inquiry Faraday not only opened the path to the establishment of modern technological civilization, but created new understanding embodied in a set of new operationally coherent activities.

Establishing New Theoretical Frameworks

The creative enhancement of operational coherence through aim-oriented adjustment happens in more theoretically focused inquiry, too. For example, consider the physicists of Einstein's generation grappling with the puzzle of the Michelson–Morley experiment (for full details see Holton 1969; Miller 1981; and Staley 2008). Assuming that light is a wave in the aether and the earth is moving around in the aether, the apparent speed of light should depend on the earth's motion, but there was no detectable variation that could be found. It was certainly not obvious what needed to be fixed in this situation in order to bring it into a state of harmony. Many ingenious accounts were given in order to explain why the motion of the earth through the aether would be undetectable, ranging from the idea that the earth dragged around the aether in its vicinity to the systematic changes in the observations of time and space coordinates proposed by Hendrik Antoon Lorentz (and similarly by George FitzGerald). Einstein's solution was more daring: he got rid of the problematic situation altogether by proposing his two postulates: the principle of relativity, and the principle of the constancy of the velocity of light (regardless of the motion of the source or the observer of light). With these postulates Einstein set about reconceiving the very concepts of space, time, mass and energy. Einstein showed that a very coherent system of practice could be built on this initially implausible basis, a system that included the activities of defining

https://doi.org/10.1017/9781108635738.002 Published online by Cambridge University Press

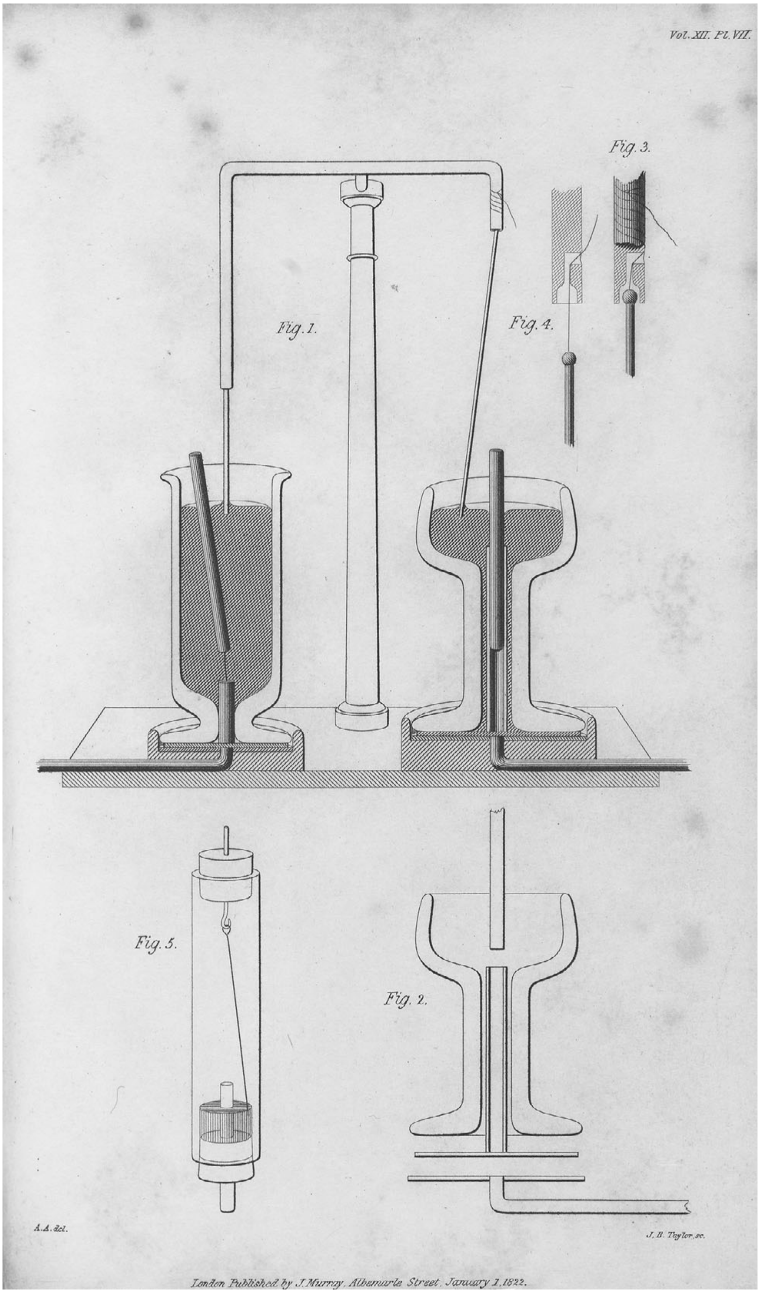

Figure 1.2 Faraday's compact rotation device, from Faraday (1822), plate 7. Courtesy of Cambridge University Library.

https://doi.org/10.1017/9781108635738.002 Published online by Cambridge University Press

frames of reference, transforming physical quantities switching between reference frames (via the Lorentz transformation equations), deriving observable consequences of these transformations (including time-dilation, length-contraction, velocity-dependence of mass, and mass–energy equivalence), devising experiments to test these consequences, and explaining various well-known observations (including the Michelson–Morley experiment). After Einstein and his colleagues were done, the very fundamentals of physics had been transformed. This is a very large-scale example of how unrestricted inquiry creates new settled meaning in an initially incoherent situation. The desire for coherence and understanding provides both a powerful motivation and ongoing guidance for innovation.

Einstein's 1905 work on special relativity certainly created profound new understanding, but we need to ask more carefully: what exactly is the sense of understanding involved here, and how does it relate to active knowledge? Recall that pragmatic understanding is the business of making sense of our activities. Einstein's work is a useful case to consider especially because it was almost entirely theoretical; this serves as a reminder that epistemic activities do not necessarily involve physical operations with material objects, and that 'pragmatic' does not equal 'material' or 'practical' in a crude sense. It is also important to note that special relativity enhanced active knowledge not so much by increasing the coherence of existing activities, but by creating entirely new activities that turned out to be coherent. At the foundation of all activities in the relativistic system of practice is the activity of defining the frame of reference for any given observer, by putting in (by imagination) a lattice of rigid rods and clocks all moving at the same velocity. The activities of relativistic frame-setting and inter-frame transformations provided the basis for many new coherent theoretical activities, including those involved in relativistic quantum mechanics and the recasting of the old connection between electricity and magnetism. This new way of sense-making also led to plenty of new propositional knowledge, embedded in the new activities.

The Adjustment of Aims

In the search for better coherence, even the very aims of one's activities may be altered. That is to say, aim-oriented adjustment does not necessarily mean the adjustment of everything else in order to achieve a fixed aim; sometimes the best move is to adjust one's aim so that it becomes more realistic to achieve. I retain the notion of aim-orientedness even in that situation, because the operational coherence of an activity is still defined in relation to the aim of the activity at each moment, even though the situation is dynamic because our aims are not fixed in the long term. Aims can and should change in the process of inquiry, if we learn that what we were trying to do is not plausible and the best way to enhance coherence is to try for a different aim. For example, while engaged in the activity of hammering nails, we may find ourselves surrounded by steel walls; the reasonable thing to do in that situation would be to stop hammering, and find some other way of hanging the picture. It makes sense to stop trying to achieve an unfeasible aim. If we change the inherent aim of an activity, the activity ceases to exist. With a new aim we enter into a new activity, with a whole new judgement of coherence. And more positively, what we do well can often become something that we *want* to be doing, around which we also build other activities.

Talk of hammer and nail may sound idle in the context of the philosophy of science, but the modification of aims in search of coherence also happens in serious scientific practice. The case of Einstein and special relativity again provides an excellent illustration. Einstein did not solve the original problem of accounting for the result of the Michelson–Morley experiment (and other 'aether-drift experiments') in terms of the motion of the earth through the aether. Rather, he dismissed that problem by declaring the aether 'superfluous', and launched whole a new set of epistemic activities. His proposal was attractive to many other physicists who had become weary of their inability to solve the original problem. The course of inquiry can very well lead us to a reassessment of what it is that we ought to be aiming to know, and what questions are worth investigating.

1.6 Pragmatism and Active Knowledge

My thinking in this book has been strongly inspired by the pragmatist tradition of philosophy, so it is important that I explain what I take pragmatism to be, and why it is so relevant to thinking about active knowledge. Pragmatist philosophy is often concerned with clarifying the meanings of concepts in terms of their practical implications. More broadly, however, pragmatism is a philosophy that concerns itself with the nature of practices, which is therefore suitable as a framework for understanding the nature of active knowledge. Pragmatism as I see it is a thorough and relentless empiricism, which insists that experience is the only ultimate source of any kind of learning, and takes experience in its full sense as something that active knowers undergo. For pragmatists, empirical learning includes the learning of methods, all the way to the choice of appropriate logical axioms. Such a pragmatist perspective will helpfully

https://doi.org/10.1017/9781108635738.002 Published online by Cambridge University Press

inform all the other discussions in the remainder of this book. My interpretation of pragmatism is also offered as a small contribution to the recent revival of pragmatist philosophy in various quarters.[32]

Beyond Semantics: Pragmatism as a Philosophy of Practice

What is pragmatism, and what are its implications for the philosophy of science? A very good definition actually comes from an ordinary dictionary (*Webster's Ninth New Collegiate Dictionary*, 1986): 'an American movement in philosophy founded by C. S. Peirce and William James and marked by the doctrines that the meaning of conceptions is to be sought in their practical bearings, that the function of thought is to guide action, and that truth is preeminently to be tested by the practical consequences of belief'. The first part of this definition is nothing but a version of Peirce's 'pragmatist maxim', paraphrased by James here ([1907] 1975, p. 29): 'To attain perfect clearness in our thoughts of an object, then, we need only consider what conceivable effects of a practical kind the object may involve – what sensations we are to expect from it, and what reactions we must prepare.'[33] The Peirce–James pragmatist maxim naturally led to a semantic interpretation of pragmatism, which is perhaps the dominant interpretation today. According to Catherine Legg and Christopher Hookway (2021, sec. 2), the pragmatist maxim 'offers a distinctive method for becoming clear about the meaning of concepts and the hypotheses which contain them. We clarify a hypothesis by identifying the practical consequences we should expect if it is true.' In this way, pragmatism shares much with operationalism, and with the verificationism that was widely taken as a core doctrine of logical positivism.

James presented pragmatism as a 'method for settling metaphysical disputes that otherwise might be interminable' (James [1907] 1975, p. 28). A dispute is idle unless some 'practical difference' would follow from either side being correct. James opened his lecture on 'What pragmatism means' with an apparently trivial anecdote: on a visit to the mountains, his friends got into a 'ferocious metaphysical dispute' – about a squirrel! The animal was hanging on to one side of a tree trunk, with a human observer on the other side:

[32] The literature is too large and varied for me to survey adequately here. I will be discussing various individual authors' works in the remainder of the book. Some excellent collections of recent works include volumes edited by Cheryl Misak (2007a), Roberto Frega (2011) and Kenneth Westphal (2014).

[33] For the original formulation, see Peirce (1878).

> This human witness tries to get sight of the squirrel by moving rapidly round the tree, but no matter how fast he goes, the squirrel moves as fast in the opposite direction, and always keeps the tree between himself and the man, so that never a glimpse of him is caught. The resultant metaphysical problem now is this: *Does the man go round the squirrel or not?*

And here is James's solution of the problem:

> [The correct answer] depends on what you *practically mean* by 'going round' the squirrel. If you mean passing from the north of him to the east, then to the south, then to the west, and then to the north of him again, obviously the man does go round him, for he occupies these successive positions. But if on the contrary you mean being first in front of him, then on the right of him, then behind him, then on his left, and finally in front again, it is quite as obvious that the man fails to go round him, for by the compensating movements the squirrel makes, he keeps his belly turned towards the man all the time, and his back turned away. Make the distinction, and there is no occasion for any farther dispute. (Ibid., pp. 27–8; emphases original)

In this manner, the 'pragmatic method' promises to eliminate all irresolvable metaphysical disputes, and rather more important ones than the squirrel-puzzle, too.

Even though I completely endorse the semantic tradition of pragmatism that James's squirrel example would seem to embody, my own emphasis is different. I follow Philip Kitcher's (2012, pp. xii–xiv) warning against the 'domestication' of pragmatism. Focusing on semantics can be a very effective method of domestication, making pragmatism look like a rather innocuous and interesting variation on normal analytic philosophy. I want pragmatism to be a philosophy that helps us think better about how to do things, not just about what our words mean in relations to actions. Recall the second part of the dictionary definition of pragmatism quoted above: 'the function of thought is to guide action'. Concerning the squirrel, one might wonder if what James advocates isn't just a matter of defining one's terms carefully. But I think that the sort of disambiguation offered by James is tied closely to potential practical ends. If my objective is to make a fence to enclose the squirrel, then I have gone around the squirrel *in the relevant sense*; if my objective is to check whether the wound on his back has healed, then I have failed to go around the squirrel in the relevant sense. It is the pragmatic purpose that tells us which sense of 'going round' we *ought* to mean. Semantics should be a tool for effective action. This is fully compatible with Huw Price's (1988) neo-pragmatist functionalism about truth and other notions.

Pragmatism as Relentless Empiricism

One very important reason why people often do not like to go beyond the semantic dimension of pragmatism is the fear of what happens if we go further and adopt the pragmatist theory of *truth*.[34] This issue needs to be tackled head-on. It is crucial that we reject the common misconception that pragmatism takes whatever is *convenient* as true. The 'pragmatic theory of truth' attributed especially to James is widely regarded as absurd, and this has contributed greatly to the disdain for pragmatism among tough-minded philosophers. Here is probably the most notorious statement by James:

> '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, as sounding too much like just 'convenient' or 'useful'. Or perhaps the word had quite a different connotation back then; that is for James scholars to debate. At any rate, the statement actually continues as follows:

> . . . 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. (James [1907] 1975, p. 106; emphases original)

In my view, what this passage really shows is James the staunch empiricist, declaring that the source of truth is experience, and that it is *futile to entertain any more grandiose notion of truth*. (I will have more to say about the pragmatist theory of truth in Section 4.6.) This provides an important clue to my interpretation of pragmatism, which is to understand it as a thoroughgoing, complete and relentless empiricism. Empiricism recognizes experience as the only ultimate source of learning, and refuses to acknowledge any higher epistemic authority. Here is James again ([1907] 1975, p. 31):

> Pragmatism represents a perfectly familiar attitude in philosophy, the empiricist attitude . . . both in a more radical and in a less objectionable form than it has ever yet assumed. A pragmatist turns his back resolutely and once for all upon a lot of inveterate habits dear to professional philosophers. He turns away from abstraction and insufficiency, from verbal solutions, from bad *a priori* reasons,

[34] My statement may be puzzling to those who treat semantics as a matter of truth-conditions. I prefer to take semantics as a study of meaning in a broader sense.

https://doi.org/10.1017/9781108635738.002 Published online by Cambridge University Press

> from fixed principles, closed systems, and pretended absolutes and origins. He turns towards concreteness and adequacy, towards facts, towards action, and towards power. That means the empiricist temper regnant, and the rationalist temper sincerely given up.

According to pragmatism as I see it, how philosophy serves life is by paying full and thorough attention to the experiences that constitute life.

In empiricism as it is generally presented in philosophy these days, the view of experience presented is extremely limited, seen as a matter of gaining information through sense-perception. Hans Radder (2006, ch. 2) rightly laments the 'absence of experience in empiricism'. The classical pragmatists had a feeling for the richness of experience, which I think all empiricists should recover. Cheryl Misak (2013, p. 12) argues that the early pragmatists were inspired by Ralph Waldo Emerson, who wanted empiricism, but not 'paltry empiricism'; for Emerson, experience of course included emotional and passional experience. James's 'radical empiricism' involved paying respectful attention even to religious, mystical and parapsychological experiences. It is an important part of pragmatism to take 'experience' as the full lived experience of human beings, recognizing its full range and all of its aspects. Pragmatism also understands experience in the context of action, which goes well with my conception of active knowledge. In taking such a well-rounded view of experience, pragmatism can look surprisingly different from what philosophers normally mean by 'empiricism'.

Let us consider further what a full understanding of experience involves. Empiricism is significantly perverted when it is taken to imply that we should assign absolute epistemic authority to results of sense-perception, ignoring other dimensions of experience. And sense-perception itself is much more complex than often imagined by many philosophers of science, involving much more than 'the five senses'. We have a great deal to learn from the phenomenologists in this regard, and there is much potential in a synergy between phenomenology and pragmatism.[35] Proprioception and muscular tension are inevitable underlying ingredients to all sensation, even when they are not subjects of conscious experience. Hacking (1983, p. 189) brings this consideration into scientific observation, too: 'you learn to see through a microscope by doing, not just looking'. He invokes George Berkeley's *An Essay Towards a New Theory of Vision*, 'according to which we have

[35] We could do worse than starting again with Maurice Merleau-Ponty ([1945] 1962), and then moving on to current authors such as Alva Noë (2004) and Mazviita Chirimuuta (2015). Long ago Herbert Spiegelberg (1956) noted an obvious affinity to phenomenology in the works of James, and made an in-depth comparative study of Peirce and Husserl.

three-dimensional vision only after learning what it is like to move around in the world and intervene in it'. Berkeley also points out that muscular sensations in the eye are an essential part of the experience of seeing (Berkeley [1709] 1910, pp. 16–20). Taking experience as it is really lived also means paying more attention to the active dimensions of experience, which is also to say that we focus on the *practices* of perception and observation. Even just looking at something is a concerted activity undertaken with a purpose. A focus on action also means a focus on the active agent, with a full awareness of the nature of epistemic agents as discussed in Section 1.2. We should also recognize experience as process. James's discussion of the 'stream of consciousness' referred to Henri Bergson's ideas on the passage of time, 'real duration', and memory (Bergson [1896] 1912). We need to take experience as a process of life, not as a set of statements describing contents of sense-perception. Common views on the nature of observation, experiment and empirical evidence need significant updating and revitalization.

The pragmatist view of knowledge is deeply rooted in humanism. Here I am using the term 'humanism' broadly, in the spirit of the now-forgotten German-British philosopher Ferdinand Canning Scott Schiller (1939, pp. 65–80). Humanism concerning science recommends that we understand and promote science as something that people do, not as a body of knowledge that exists completely apart from ourselves and our investigations. Science is a thoroughly human activity even when it is aimed at the production of the most abstract and impractical kind of knowledge. Empiricism itself can be seen as a form of humanism: ultimately, the only source of learning we have is human experience, shaped by human nature. As James famously put it ([1907] 1975, p. 37): 'The trail of the human serpent is thus over everything.' Perhaps this sort of humanism is not such a controversial stance (with its roots going back at least to Kant), but I think there is much value in spelling out its meaning and implications carefully. The spirit of humanism has been summarized in another way, and rather poetically, by Clarence Irving Lewis in his review of Dewey's 1929 masterpiece, *The Quest for Certainty*:

> Man may not reach the goal of his quest for security by any flight to another world – neither to that other world of the religious mystic, nor to that realm of transcendent ideas and eternal values which is its philosophical counterpart. Salvation is through work; through experimental effort, intelligently directed to an actual human future. (Lewis 1930, 14)

This passage is especially nice because it brings together the two pragmatist philosophers that I have found most inspiring, and ties together the humanist and the empiricist strands of pragmatism. Active knowledge fits very well into this empiricist-humanist spirit of pragmatism.

The most important thing about humanism as I see it is not the focus on the biological species *Homo sapiens*. Humanism is not incompatible with attention to artificial intelligence, animal cognition, or extraterrestrial intelligence. On the contrary, our understanding of non-human cognition would only be enhanced by a comparative view based on a full understanding of human cognition, and vice versa. The most fundamental point about pragmatism, as I take it, is that knowledge is created and used by some sort of epistemic agents, namely intelligent beings who engage in actions in order to live better in the material and social world. Recall Dewey's trenchant critique of the 'spectator theory of knowledge' (see Kulp 2009).

An Empiricist View on Methodology and Logic

The staunch empiricism at the core of pragmatism also encompasses methodology. In Dewey's view, method-learning is an empirical process as much as any other learning: 'through comparison-contrast, we ascertain *how* and *why* certain means and agencies have provided warrantably assertible conclusions, while others have not and *cannot* do so in the sense in which "cannot" expresses an intrinsic incompatibility between means used and consequences attained' (Dewey 1938, p. 104). Methodological rules are contingently generalizable, just like any general empirical statements: 'inquiry, in spite of the diverse subjects to which it applies . . . has a common structure or pattern . . . applied both in common sense and science' (ibid., p. 101). There is no special method for method-learning, which is just a kind of empirical learning from experience, our assumptions being conditioned by success. And the success of methodology is merely a matter of being 'operative in a manner that tends in the long run, or in the continuity of inquiry, to yield results that are either confirmed in further inquiry or that are corrected by use of the same procedures' (ibid., p. 13).[36] Dewey held strongly to the continuity of rules – of everyday inquiry, scientific method, and even logic (ibid., pp. 4–6), all of which arise from successful habits of thinking (ibid., p. 12). It was not a category-mistake that Dewey titled his last great work (published when he was nearly eighty) *Logic*, and subtitled it *The Theory of Inquiry*.[37]

[36] Dewey also says success is measured by 'coherency' but doesn't seem to say what exactly *that* means.

[37] For a thorough treatment of Dewey's ideas on logic and methodology of science, see Matthew Brown (2012).

Dewey brings methodology down to earth: 'we know that some methods of inquiry are better than others in just the same way in which we know that some methods of surgery, farming, road-making, navigating or what-not are better than others' (1938, p. 104). Ways of reasoning have also developed through the course of the history of science, sometimes in quite fundamental ways, and often it is the development of very specific new methods that challenge the accepted general rules of thought. Looking back at the history of science, we can spot disputes about the best ways of reasoning in almost every period. Should we think of velocity and weight as quantities? The medieval Aristotelians had thought not, but they were defeated by the indisputable successes of quantification (Crombie 1961). Are thought-experiments legitimate ways of generating and supporting physical theories (Stuart, Fehige and Brown 2018)? Galileo, and later Einstein, thought so, and others disagreed. How about computer simulations (see, e.g., Galison 1997, ch. 8)? Such debates and changes are still ongoing: witness the dispute surrounding whether mathematical proofs done by computers should be considered valid (Burge 1998), or more practically, how best to put to use medical diagnoses made by artificial intelligence systems (Ahmad et al. 2021). Less outlandishly: which is the right sort of statistics to employ in empirical testing, frequentist or Bayesian (Mayo 1996)? What are the circumstances in which a randomized controlled trial is the appropriate method of hypothesis-testing (Cartwright 2011; Cartwright and Hardie 2012)?

It is an indisputable historical fact that science has learned new and better methods of inquiry in the course of its development. And the adoption of a new method is often accompanied by changes in fundamental epistemic standards as well. Take Alan Chalmers's discussion of Galileo's establishment of the telescope as a trusted means of making astronomical observations, superior to naked-eye vision (Chalmers 2013, pp. 151–5). This was a momentous event that had a deep implication reaching far beyond itself, as it was one of the key early instances in which the verdict of an instrument was deemed to be epistemologically superior to human sensation. I have told a similar story concerning the overriding of the sensation of hot and cold by thermometer-readings (Chang 2004, p. 47). These methodological changes were not sanctioned by some super-method or an overarching theory of physics, but only by some detailed case-by-case arguments about how the telescopic or thermometric observations should be given more credence than unaided perception. For example, the telescope showed some details that naked-eye observations did not show (such as the phases of Venus and the craters on the moon), and these details also made sense (if one adopted the Copernican theory, that is). Galileo also pointed out some inconsistencies in the human visual perception of the brightness and shape of planets and stars. Galileo was not able to argue from any first principles that the telescope should be granted epistemic authority (especially in the absence of a well-developed theory of optics at the time), but the use of telescopes in some concrete inquiries worked out very coherently.

It would have been the ultimate prize for a pragmatist to argue successfully that even logic was only pragmatically justified, rather than being a set of eternally valid 'laws of thought'. If so, the methods of science and any other rules we have in life are *of course* going to be revisable and evolvable. I think that is exactly what Dewey aimed to achieve. Logical rules are decreed to be true when you are in logic class, but when logic is employed in other settings it must be judged by its fruits, in terms of how well it supports operationally coherent activities. Logic for Dewey is still a normative discipline, about how we *should* think in order to think well. But the normative force of logic ultimately rests on its empirical success, not on any a priori requirements. For example, consider the cogency of the 'the principle of explosion' concerning material implication, which dictates that from a contradiction one can deduce any proposition. Is that really a productive way of thinking? I certainly have trouble thinking of any real-life situation in which we should go wild and believe any propositions at all, just because we have a contradiction on hand.

Dewey denies that logic exists at all apart from the subject matter of reasoning: 'all logical forms (with their characteristic properties) arise within the operation of inquiry and are concerned with the control of inquiry so that it may yield warranted assertions' (Dewey 1938, pp. 3–4). What people normally consider 'logic' is only the most general among the rules of thought that have been shown to be good, most abstractly expressed. In Dewey's view logic is a historical and empirical discipline in which we continually learn better ways of reasoning. Declaring that logic is a 'progressive discipline' (ibid., p. 14), he speaks of 'the needed reform of logic' (ibid., title of chapter 5) which should be based on a full historical awareness. For example, Dewey argues that Aristotelian logic was a system admirably suited for the science and philosophy of ancient Greece, but no longer suitable (ibid., pp. 82–93). As key elements of Aristotelian thinking that have been abandoned, he lists: essentialism, the emphasis on quality over quantity, static classification as the form of knowledge, the heterogeneous and hierarchical structure of the universe (again, one can see how logic and methodology blend into each other in Dewey's thinking). He faults his contemporary logicians for tending to retain the form of classical logic while abandoning the metaphysical and operational underpinnings of it.

One might think that Dewey was just a 'fuzzy' thinker who was ignorant about logic as practised by professional logicians. But no one could level the same accusation against C. I. Lewis, whose views on logic were very similar to Dewey's. Lewis had built a great and lasting reputation in the area of symbolic logic (for which he is still remembered) before he became known as a pragmatist. Lewis (1929, p. vii) himself stated that his pragmatist epistemology had in fact originated from his work in symbolic logic. Early on Lewis felt, like Dewey, a dissatisfaction with certain features of the fundamental notion of material implication in classical logic. Founded on the principle of excluded middle, classical logic dictates that *P* (materially) implies *Q* whenever it isn't the case that *P* is true and *Q* is false. Lewis thought that this notion did not correctly express what we mean by 'implication' in natural language, and it should be replaced by what he called 'strict implication', which did not have undesirable consequences like the principle of explosion that Dewey deplored. But critics pointed out that Lewis's strict implication had its own awkward features, and he was forced to admit that each system of logic had various virtues and vices. This pushed Lewis towards pluralism about logic, and about conceptual frameworks in general.[38]

There *are* different systems of logic, and anyone who wants to reason logically must start by adopting a particular system of logic. But the only plausible and non-arbitrary way of justifying the choice of a logical system would be on pragmatic grounds, because appealing to the rules of logic in arguing for this choice would clearly be question-begging. So it may actually turn out that the treatment of logic is the most convincing part of pragmatism! Lewis summed up his view as follows: 'the choice of conceptual systems for the interpretation of experience ... is a matter of pragmatic choice, whether that choice be made deliberately, or unconsciously and without recognition of its real grounds' (Lewis 1929, p. 300). Looking at the current proliferation of non-classical (or alternative) logics and the successful application of some of them in the design of intelligent systems, I think we must admit that Dewey and Lewis have been vindicated. The spirit of Dewey's work lives on in current work on 'anti-exceptionalism' about logic, as expressed here by Ole Thomassen Hjortland (2017, p. 631): 'Logic isn't special. Its theories are continuous with science; its method [is] continuous with scientific method. Logic isn't a priori, nor are its truths analytic truths.

[38] On Lewis's development see Schilpp (1968); Rosenthal (2007); Misak (2013), ch. 10; Stump (2015), ch. 5.

Logical theories are revisable, and if they are revised, they are revised on the same grounds as scientific theories.'[39]

To sum up: a truly empiricist philosophy of science, as I take pragmatism to be, should recognize clearly that scientific inquiry is itself a kind of human experience. According to the pragmatists inquiry is pervasive in life, as well as scientific practice; it could even be that experience is inherently inquisitive. And not only does inquiry engage with experience, but the process of inquiry itself is a type of experience. Learning from experience includes *learning from the experience of learning*. Philosophers need to pay attention to the processes of knowledge-production and knowledge-use, and ask how we can best organize and support the epistemic activities involved in those processes. If we conceive pragmatism as a philosophical commitment to engage with practices, pragmatist epistemology should concern itself with all practices relating to knowledge. And I believe that this is something that the classical pragmatists were seriously engaged in.

In the remaining chapters of this book I will carefully unfold the implications of pragmatism on epistemology and metaphysics, especially in the context of the philosophy of science. In this chapter I have focused on the pragmatic nature of knowledge and presented a pragmatist view of the process of inquiry. Central to my thinking throughout will be the notion of *operational coherence*, which I have defined in Sections 1.1 and 1.4. Operational coherence is the anchor of the kind of realism that pragmatists (and empiricists in general) can embrace, as I will explain fully in Chapter 5. Any sort of realism needs to take something real as the object of knowledge, so in Chapter 3 I also advance a pragmatist notion of reality, which is also based on the notion of operational coherence. Realism also involves the basic idea that through our inquiries we can learn truths about realities, and I will propose an updated pragmatist notion of truth in Chapter 4.

39 Hjortland traces this view to Quine, Maddy and Priest among others. For an introduction to non-classical logics see Priest (2008) and Gottwald (2020).