

## *Introduction*

### **What Is the Problem?**

The aim of this book is to develop a realistic philosophical vision of scientific knowledge, which can be truly relevant to scientific work and useful in social and political life. In a world filled with so much misinformation, ignorance, prejudice, deception and mistrust, where can we turn for reliable facts, insightful theories, and guidelines for action? I am old-fashioned enough to believe that science and the scientific attitude constitute our best hope. But we are often distracted by an impossible ideal of scientific knowledge as proven universal truth about some ultimate reality. Unrealistic ideals can have harmful consequences. This is quite evident in the current crisis in the public's trust in science. When the advocates of science imply that science is the giver of unquestionable truth, a false equivalence opens up when actual science fails to live up to its overblown image. People will think: if evolution is not completely proven, it is 'just a theory' to be treated equally with creationism. They will think: if it is not completely proven that human activity is causing climate change, and moreover some scientists disagree with the idea, then it is just a majority *opinion* among scientists, not inherently superior to *their* own opinion.

Faced with challenges to established scientific knowledge, many scientists, philosophers and concerned citizens succumb to an authoritarian or even autocratic impulse, which they often honour with the label of 'scientific realism': science gives us the Truth about Reality (or at least an approximation to it), and anyone going against the verdict of science is simply misguided (if not malicious). There is a widespread conceit that modern science has *basically* the right answers, or at least the right *methods* for getting the right answers. Many 'realists' maintain this idea like an article of faith, even though we have no direct access to the unobservable aspects of Reality and the historical track-record of science shows serious fluctuations in scientists' views about the most basic aspects of the

universe. In the philosophy of science, what usually goes by the name of (scientific) realism comes down to the notion that our best scientific theories must express some version of the final answers to the most fundamental questions about nature. I cannot help expressing the feeling that this widespread philosophy of scientific realism amounts to an appropriation of the term ‘realism’ to describe a most *unrealistic* doctrine. I want to propose a realistic kind of realism, close to what William Wimsatt has called ‘realism for limited beings in a rich, messy world’, which is ‘a philosophy of science that can be pursued by *real* people in *real* situations in *real* time with the kinds of tools we *actually* have’ (Wimsatt 2007, p. 5; emphasizes original).

One problem with the standard doctrine of ‘scientific realism’ is that it tells us very little about what we can *do* to improve scientific knowledge. We need an *operational* ideal of knowledge, one that we can actually work with, which reflects our best actual practices while also indicating how they can be improved further. Ideals are of course hardly ever achievable, but in order to serve a useful function an ideal needs to be something that has an effect of making us think and behave differently. Inventing a test for a virus with zero per cent rate of false negatives, or having no murders in a large city in a whole year, is an operational ideal, even if we never manage to meet it. Seeking absolute truth is not an operational ideal – there is nothing we can actually do in order to approach that ideal. According to the common picture of scientific knowledge, science should give us the true picture of the reality that exists well-formed ‘out there’ completely independently of our conceptions and our experiences. But such ‘reality’ is not accessible to us and there are no actual methods by which we could attain assured knowledge about it.

A realistic ideal of knowledge also recognizes that a mature science in a democratic society needs to allow debate, treating dissenters with respect in the spirit of what Philip Kitcher has called the ‘ideal conversation’ (2011a, ch. 2).<sup>1</sup> Even if effective practical action necessitates the overriding of certain views and the sidelining of certain voices, this should still be done with intellectual sympathy and compassion. Behaving like religious fundamentalists is not the way to beat religious fundamentalism. We must overcome the urge simply to shut down the opposition in science, just as the most mature societies have learned to do in politics. We should not look to science for the authority to justify repression. Science consists in

<sup>1</sup> Philip Kitcher and Evelyn Fox Keller (2017) have concretely imagined such ideal conversations concerning climate change.

*Pragmatist Notions of Knowledge, Truth and Reality* 3

continual inquiry and debate, not a collection of dogmas that are protected from criticism. As my childhood hero Carl Sagan said about science: ‘Its only sacred truth is that there are no sacred truths.’<sup>2</sup> Even though some dogmas are sometimes practically necessary, dogmatism should not be our way of life.

You might wish that scientists themselves would set up and maintain a framework for democratic debate about scientific knowledge. I believe that some scientists did make such efforts back in the Enlightenment era, after the initial excitement of rebelling against traditional doctrines passed and science needed to look more openly and humbly at its own growing epistemic authority (see Chang 2012b). But a relentless process of professionalization has left many modern scientists oblivious to such needs. In that situation, it is appropriate that philosophers of science should take up the challenge. But the sad reality is that much traditional philosophical thinking has been caught up in the contemplation of the impossible ideal of knowledge mentioned above, and failed to provide insights that are relevant to actual practices of science. Even though much good work has been done in various corners of the ‘philosophy of the special sciences’ in communication with scientific practitioners, the spirit of such work has failed to permeate the discourse in general philosophy of science, not to mention epistemology and metaphysics done without particular attention to science. Among those who care about the place of science in society and culture, there is an urgently felt need for a general philosophy of science fit for engaging with real scientific practices. That is the need that this book seeks to meet.

### **Pragmatist Notions of Knowledge, Truth and Reality**

Here is the most fundamental step to take, in order to bring our philosophical views of scientific knowledge more in line with real practices: we should be thinking about what it is that we *do* in science. I find inspiration in the operationalist philosophy of Percy Bridgman (1956, p. 76): ‘it is better, because it takes us further, to analyze into doings or happenings rather than into objects or entities.’<sup>3</sup> This proposal will take us clearly beyond the default instinct of analytic philosophers to regard knowledge as

<sup>2</sup> This statement occurs in the PBS Television series *Cosmos*, episode 13. The equivalent passage in the book can be found in Sagan (1980, p. 333): ‘It [science] has two rules. First: there are no sacred truths.’

<sup>3</sup> Bridgman’s thoughts were often misunderstood, and he did not like the label of ‘operationalism’. See Chang (2009b; 2017c) for my interpretation.

consisting of justified belief in a set of true statements or propositions (or something of that sort). As I will explain in Chapter 1, the alternative perspective I promote takes science (and inquiry in general) as something that people do, consisting of epistemic activities with various aims whose achievement we can actually assess (unlike absolute truth). From this perspective, knowledge will be seen primarily as an ability, not confined to the possession of information. Having reliable beliefs is one aspect of the ability to conduct successful activities, but it is by no means the whole picture.

In this broadly operationalist spirit I propose that we rethink the philosophical notions of ‘truth’ and ‘reality’, so that they can be used without encouraging unrealistic and indefensible dogmatism. Truth and reality are perfectly meaningful concepts in the actual business of ‘representing and intervening’ (Hacking 1983), and they should stay in that realm of practice. We need to be able to debate the truth and falsity of statements, without claiming our truths to be absolute. We need to be able to say things like ‘The President is out of touch with reality on this point’, without in the same breath invoking some realm of metaphysical reality that we cannot ourselves claim to be in touch with. Articulating such useful notions of truth and reality, and a practical ideal of knowledge based on them, is something that I regard as one of the most important tasks of philosophy.

My proposed reframing of knowledge, truth and reality is built on the notion of ‘operational coherence’, which will be spelled out in Chapter 1. Very roughly, operational coherence is a matter of making elements of our activities fit together harmoniously so that our aims may be achieved. It may consist in something as mundane as the skilful coordination of bodily movements and material conditions in riding a bicycle, or it may be something as esoteric as the successful integration of material technologies and abstract theories in the operation of the global positioning system (GPS). Building on that notion of coherence, I advance a notion of truth in Chapter 4 according to which the empirical truth of a statement consists in the positive role it can play in facilitating operationally coherent activities. Similarly, in Chapter 3 I propose that what we should mean by something being ‘real’ is that it can be employed in coherent activities that rely on its existence and its basic properties. Truth and reality conceived in such a way are attributes grounded in our activities.

In terms of philosophical traditions, what I want to do is to reframe the basic discourse in the philosophy of science through a revitalization of

*Pragmatist Notions of Knowledge, Truth and Reality* 5

pragmatism.<sup>4</sup> Among the established philosophical frameworks, pragmatism offers us the best hope for facilitating proper attention to what it is that we really do when we gain and use knowledge. Pragmatism regards knowledge as an outcome of humble on-the-ground inquiry, and locates it in actual intelligent activities we carry out in life. I hasten to add that pragmatism has often been misrepresented and vilified within philosophy, so much so that many people who hold what I would consider pragmatist views have avoided the label of pragmatism. Even Marjorie Grene, who made a pioneering emphasis on how knowledge exists in the context of action, shrank sharply from the label of pragmatism, which she took as an insult even worse than relativism: ‘But pragmatism! If that is the kind of abandonment of any interest in cognitive claims we are supposed to be advocating, we are presenting our case very badly’ (Grene 1987, p. 69). There may be times when a word is so soiled by negative connotations that we should just give it up instead of engaging in futile attempts to correct the misuses and abuses of it. But I think the term ‘pragmatism’ can and should be defended and improved, not given up – as with ‘truth’ and ‘reality’.

In later parts of the book I will give a more considered and systematic view of what I take pragmatism to mean, but it would be helpful to start with a few brief notes of clarification. The common misrepresentation of pragmatism comes from two different directions. First, there is the charge that pragmatists disregard truth in any meaningful sense of the word, because pragmatism allows one to regard as truth whatever is *convenient* to oneself. Nothing could be a worse distortion of pragmatism. Whether our conceptions work out in practice is most definitely not up to our whims and wishes, nor is it a matter of what is ‘convenient’ to us in a shallow and limited sense. In order to do things coherently, we need to have a mastery of our surroundings. It is actually through operational coherence, not by the mirage of correspondence to the inaccessible ‘Real World’, that something objective is brought to bear on our knowledge.

In the second common line of misrepresentation, it is alleged that pragmatism concerns itself only with ‘practical’ things, and people have even blamed it for the currently prevalent devaluing of the humanities and other ‘impractical’ fields of study. True pragmatism is about how our conceptions work out in all practices, not just in the money-making or enemy-killing sorts of ‘practical’ practices. There is huge variety within human practices, in pursuit of diverse types of aims, and pragmatism is

<sup>4</sup> In many ways I follow the directions indicated by James Woodward (forthcoming).

concerned with all of them. For example, most of the experiments going on at CERN and other places of research in pure science are decidedly not ‘practical’, at least not immediately. Yet what goes on at such installations is the epitome of the empirical testing of theories, which pragmatists admire and cherish. The pragmatist vindication of the concept of imaginary number is in all the fruitful mathematical practices it enables, including practically useless ones as well as engineering applications. The pragmatist vindication of the fugue is in the experience of elation that countless listeners have felt in hearing Bach’s masterpieces; whether such musical pieces can be put to ‘practical’ uses such as helping soldiers march in step is irrelevant to their pragmatist appreciation.

What the spirit of pragmatism recommends is that philosophy should not be removed from the various practices in life, including science. What philosophers need to do is think different and unusual thoughts, but also about things that matter to people here and now. We do need the proverbial armchair in which to do philosophy, but only as a place of necessary and occasional retreat to allow us to think quietly and carefully about issues of actual life. It may be surprising to some to learn that such practical engagement was an inclination strongly expressed in the manifesto of logical positivism by the Vienna Circle (which I quote with apologies for the gendered language of the time):

Neatness and clarity are striven for, and dark distances and unfathomable depths rejected. In science there are no ‘depths’; there is surface everywhere . . . Everything is accessible to man; and man is the measure of all things. Here is an affinity with the Sophists, not with the Platonists; with the Epicureans, not with the Pythagoreans; with all those who stand for earthly being and the here and now. (Neurath et al. [1929] 1973, p. 306)

### Scientific Realism as Realistic Activism

The philosophical battleground most relevant to my concerns expressed above is the fierce and long-standing debate on scientific realism. Ilkka Niiniluoto opens his erudite and insightful book on scientific realism thus: ‘The philosophy of science in the twentieth century has been a battlefield between “realist” and “anti-realist” approaches’ (Niiniluoto 1999, p. v). I wish to take philosophers’ attention away from standard scientific realist attempts to show that the impossible is *somehow* possible, namely that empirical science can attain assured truths about what truly goes beyond experience. As I will argue in Chapter 2, it is time to accept the fact that we cannot know whether we have got the Truth about the World (and that

such thoughts are perhaps not even meaningful). Scientific realists go astray by persisting in trying to find a way around this fact, while anti-realists make the mistake of engaging unproductively with that realist persistence.

I would like to reshape realism into a stance<sup>5</sup> that is useful for scientists and others who are actually engaged in the production, assessment and improvement of empirical knowledge: a *realism for realistic people*, which will be fully articulated in Chapter 5. I am trying to carry out a task that I already advocated a decade ago: ‘I think *realistic* people (including most empiricists and pragmatists) should re-claim the label of “realism”!’ (Chang 2012a, p. 217; emphasis original). This is a call that had already been made by many others. For example, there is some affinity between my line of thought and what Peter Kosso (1998, ch. 8, pp. 177–8) called ‘realistic realism’. I cannot remember whether I absorbed the ‘realistic’ trope from Kosso, but I would be happy if that were the case; in Section 5.1 I will also mention various other authors who have voiced a realistic spirit.

I hope that my articulation of an operational ideal of knowledge in this book will persuade both traditional ‘realists’ and ‘anti-realists’ that there is a more realistic version of realism that they can all subscribe to. The realistic realism I offer here is focused on genuine empirical learning, and provides a notion of knowledge suitable for conceptualizing how rational decision-making works in real life. Knowledge is only meaningful within the world in which we live. It is a futile and pernicious philosophical dream to seek what Hilary Putnam (1980a, p. 100; 1981, p. 49) called the ‘God’s Eye point of view’ on nature, an ‘externalist’ perspective from which we can tell the ‘real’ shape of the world. Roberto Torretti (2000, p. 114) blasts the ‘scientific realists’ who believe ‘that reality is well-defined, once and for all, independently of human action and human thought’ – yet ‘in a way that can be adequately articulated in human discourse’. The self-designated ‘realists’ hold that science aims to develop ‘just the sort of discourse which adequately articulates reality – which, as Plato said, “cuts it at its joints” – and that modern science is visibly approaching the fulfilment of this aim’. Torretti confesses that he finds it difficult ‘to accept any of these statements or even to make sense of them’. Epistemology should not be focused on the non-existent ‘final’ state of knowledge. And it is futile to insist that

<sup>5</sup> I use the term ‘stance’ in a similar spirit as in van Fraassen (2002) on empiricism and Kellert, Longino and Waters (2006) on pluralism. Here is a slightly later formulation from van Fraassen (2004, p. 128): ‘A stance consists of a cluster of attitudes, including propositional attitudes (which may include some factual beliefs) as well as others, and especially certain intentions, commitments, and values.’

a vision of scientific knowledge expressed from the God's Eye view is useful as a regulative ideal; nothing can be considered 'regulative' if it has no actual bearing on how we conduct ourselves, if it doesn't in fact help *regulate* our practices.

The realism I advocate in this book is an *activist* stance, as I will explain further in Chapter 5. When it is applied to science, it is the philosophical position that says science should do everything it can in order to gain more and better knowledge, as opposed to the position of a spectator who observes with satisfaction that science seems very good at finding knowledge, and that this knowledge must be a fair reflection of the state of the real world. Most scientists would clearly endorse the activist stance if questioned about it, though their actual conduct sometimes goes in a different direction. It is worth noting that the activist nature of the realism I advocate makes it a very different kind of doctrine from metaphysical realism (with which 'scientific realism' usually aligns itself, as discussed in Chapter 2). It makes sense that the pronouncements of metaphysical realism ('The external world is real!') leaves most practising scientists cold ('Go away, I'm busy learning something about nature!'). Realism as I see it is a stance in active service of scientific progress. It applies in a similar spirit to other areas of life, too, wherever we regard knowledge as a good thing. It is an injunction against resting content with what we believe, or what we think we know. Such a realist spirit is intimately connected to social and political progress as well.

### What Kind of Book Is This?

Many philosophers have an ambition to write a book that becomes the focus of philosophical debates everywhere. My ambition is focused more on writing a useful book. What I am trying to offer is a set of ideas that people in various walks of life can pick up as tools for their own work, ideas whose worth will hopefully be proven through productive use by those who want to understand and promote good practices relating to knowledge. In my own research, the immediate use of ideas in the philosophy of science has always been in the framing of the history of science. The ideas in this book are being applied to (and shaped by) some new historical research I am conducting at the same time, whose results will be published in a separate book (*How Does a Battery Work?*, forthcoming from the University of Chicago Press). I hope that the philosophical ideas developed in the present book will provide useful framing devices for other historians, too, as well as sociologists, anthropologists and others who make empirical



*What Kind of Book Is This?*

9

studies of science. I also hope that my ideas can help practitioners of science who want to make more explicit sense of the aims and methods of the work that they do so well. When I say ‘practitioners of science’ I have a broad range of people in mind: not only research scientists, but engineers, doctors, mathematicians, and teachers and students of all those fields, too. Even more broadly, I hope that this book will be useful to all those who want to think carefully about the place of knowledge in our individual and social lives, including policy-makers and the general public. The operational ideal of good knowledge that I am proposing here will hopefully aid clear thinking in all the areas where science meets politics and ethics.

I do not focus on the adversarial type of argumentation, except where I feel that it is really necessary for creating the space in which my own views can grow, as in Chapter 2. I will tend to bypass many of the ongoing cutting-edge debates in relevant areas of professional philosophy. This is not from a lack of respect for that work, but simply a matter of my own priorities. It is not my main concern to argue that other people have been wrong, especially when I know that many of them are superior intellects to myself who have put a great deal of effort into developing and defending their views. In fact, avoiding excessive and minute disputation is part of my aim. I do not wish to write the kind of specialized academic philosophy that is practically unusable even for most professional philosophers not working in the immediate narrow sub-fields within which the debates take place.

I am most interested in calling for a change of perspective or stance, rather than proving specific points. Sometimes I will merely be introducing useful tautologies, for example when I say that ‘all we can ever think or speak about are conceptualized entities’ (see Chapter 3).<sup>6</sup> I aspire to present an overall vision that you, Dear Reader, can subscribe to, even if you disagree with some particular points I make. And, in any case, I do not think that philosophy is the kind of enterprise in which one can prove points absolutely or win arguments decisively. I am passionate about the ideas I am presenting, but I seek to master the art of *respectful* denunciation, *peaceful* incitement, and *productive* frustration. I would like to be reasonable and rational without going on about Reason and Rationality, to be realistic without grand claims about Reality, and to be honest and true without banging on the table about Truth.

<sup>6</sup> Everyday examples of useful tautologies include ‘A man’s gotta do what he’s gotta do’ and the recently infamous ‘Brexit means Brexit’ (Theresa May, UK prime minister).

This is the first time I have ever attempted to write an entire book of abstract philosophy. In departing from my normal mode of work I felt that I was answering a call of duty, though the work has certainly been pleasurable, too. The ideas I am presenting here needed to be articulated, and I could not find them already put together in a clear, systematic and accessible way anywhere. So I was compelled to try to articulate them in my own way, for myself and for others. I am not even good at sustained abstract thinking, especially compared to the formidable thinkers who are the leading lights of academic philosophy, whose works justly inspire awe and admiration. But I take comfort in the thought that geniuses are rarely good teachers or explainers. And I did not in fact succeed in writing an entirely abstract book. On the contrary, the discussions to follow will be peppered with many concrete examples from the history of science, as well as everyday life, though there are no sustained historical studies. One caveat: you will find that most of my scientific examples are from physics and chemistry, not drawn in a balanced way from across the natural, social and human sciences. To be frank, this is because physics and chemistry are the sciences I know best. But I think my limitation can also be an advantage, because the kind of traditional philosophical views about science that I am trying to move beyond have largely been inspired by the physical sciences, especially physics. By including many examples from physics as illustrations of my points, I am making the point that even physics isn't like 'physics' (i.e., the common image of physics).

It is my own honest view that there are very few truly original ideas contained in this book. While I don't agree with the common adage that all philosophy is footnotes to Plato and Aristotle, I do think that it is difficult to have a completely original thought on any important issue that people have worried about for centuries. My thoughts have been inspired by a number of great thinkers past and present, ranging from Immanuel Kant to Nancy Cartwright. What I do think is original in my book is the particular way in which I synthesize the old familiar thoughts, which I hope will also enhance the value and meaningfulness of each component going into the synthesis. Again, my main aim is to present a positive view that you can use for your own purposes. I will make my best attempt to connect and engage with other authors' works as needed, but in general my focus will not be on fine-grained disputation or exegesis. Rather, I honour other thinkers by using their ideas as inspirations in thinking for myself, and by building directly on their work whenever I can. And I hope that you will do the same with my work.