

Philosophy as engineering

Adam Morton

University of Bristol

Philosophy used to compare itself to mathematics. It aimed for certainty and proof, and an aristocratic oversight above the rest of knowledge. More recently philosophy has compared itself to science, or more accurately to a science. Philosophy is one discipline among others, aiming to find truths about the relations between thought and its objects, in a way that requires evidence from fallible sources, including evidence pre-digested by other sciences. I shall suggest a different comparison. Philosophy is like engineering. We are concerned above all with topics where theory and evidence are not in perfect agreement, and where practical needs force us to consider theories which we know cannot be exactly right. We accept these imperfect theories because we need some beliefs to guide us in practical matters. So along with the theories we need rules of thumb and various kinds of models¹. We need a kind of first aid: what to do till the scientists arrives. In some cases it may never arrive.

¹ I mean models in the sense of structured objects, representing the facts, which mediate between theories and data. See the introduction and articles in Margaret Morrison and Mary S Morgan *Models as mediators* (Cambridge: Cambridge University Press, 1999.)

Two consequences of this comparison are likely to be controversial. The need to accept beliefs on imperfect evidence may not seem consistent with actual philosophical practice, since we argue about every tiny detail and we subject even plausible suggestions to intense scrutiny. Moreover at the heart of the scrutiny there are often counterexamples - or examples that might or might not turn out to be counter - which need only be logically possible, or consistent with the best current science, in order to provide objections. And, second, the direction to practical purposes may seem far from the motivation of practising philosophers, obsessed with truth, real and not approximate truth, for its own sake.

I think the appearances are misleading. "Will it work?" is as important to philosophy as "is it true?" A special kind of "will it work?", though. The first task is to overcome the appearances. This is best done with examples.

Example one: epistemology One kind of epistemology has existed only since the time of Galileo. Since the early days of the scientific revolution European philosophers wrestled with a particular interconnected set of issues, concerning the sources of knowledge, the relative roles of perception, reason, and traditional belief. These issues were animated by two realisations. First, that an understanding of the world had become possible, to be obtained by use

of some mixture of experience, experiment, mathematics, free hypothesis-making, and common sense, which had never been available before. Second, that this new understanding threatened conflict with deeply entrenched beliefs, including those of religion. The depth and width of possible new knowledge and the seriousness of the threat to traditional belief were not clear.

A new project, with new opportunities and new dangers, needs a new set of rules for its participants. And a new vocabulary for stating these rules, and more generally for the commentary, praise, and criticism necessary to keep the project running. Any intellectual activity has such a vocabulary, so that we can then criticize and fine-tune our performance. Often it is a specialized vocabulary. For example in chess playing we talk of strategies and traps and gambits. These vocabularies are not needed to take part in the activity; they are for talking about it, not doing it. But often we do it better if we talk about it. Someone could play brilliant chess without knowing about traps and gambits - they're different from the concepts of a king or of checkmate - but no one could discuss chess usefully without something like them. No one could discuss chess usefully with us, that is. There could be quite different vocabularies for talking about how to write winning computer chess programs, or even for talking about the game with people who approached it very differently from most of us. Whenever we can do things more or less successfully there is room for a vocabulary for discussing how things went and how to avoid their going badly.

In the case of early scientists the philosophical program of finding the right normative vocabulary for fine-tuning the project took the form of modifying an existing rhetoric. The rhetoric is the language for describing beliefs and the ways in which they are acquired - in modern language "rational", "reason for", "(un)substantiated", "evidence" and their kin - and the language for describing people and their belief-acquiring traits of mind - "hasty", "intelligent", "deluded", "gullible" and their kin. These play a role in everyday life in our efforts to assess the reliability of one another's beliefs. The project of modern epistemology, from Descartes to Popper, has been to adapt this regulatory vocabulary to provide a description of a set of practices which gives a good chance of achieving the promised knowledge and a vocabulary for making the meta-judgements needed for regulating these practices. The project has produced no end of doctrines and problems. And it has been largely successful, in that as science has developed it has developed a regulative vocabulary adequate to its purposes, the language of data, hypothesis, explanation, evidence, and established knowledge. Professional and amateur epistemology have both played a role in this. Progress here is largely invisible - as is most philosophical progress - as once a vocabulary is in place it seems the natural and inevitable way of describing its subject matter.

There is another source of the language and official problems of epistemology. That source is the debates among philosophers writing in Greek

and Latin in late antiquity about knowledge, which we anachronistically assimilate to the modern epistemological project. In fact, the aim of the disputes among Stoics, Epicureans and Peripatetics is essentially moral. The main question is how much we can presume know about the world, using the methods already available, not trying to describe any new methods. And this question is important not for its own sake but because it plays a part in disputes about the things one should rely on in ones life. Essentially, if you can know nothing about reality beyond your immediate experience there is no point to shaping your life around any beliefs concerning any such unknowables. It is an interesting question how much the vocabulary of these debates influenced the everyday regulative vocabulary of belief. (The answer could be: not at all.) In any case, by the time of the scientific revolution the terminology of ancient epistemology became one source for the language needed for framing a new and essentially different set of issues. Galileo's questions need a language, so that of Montaigne and Erasmus, deriving ultimately from Epicurus, Sextus, and Pyrrhus, is adapted². Centuries later, we have difficulty seeing that the language could have any other focus.

Epistemology as the search for the right metadiscourse for science is now a large part of the philosophy of science. The issues are hard, important, and

² See Michael Frede's introduction to Michael Frede and Gisela Striker, eds. *Rationality in Greek thought* (Oxford: Oxford University Press, 1996), and Julia Annas "Scepticism, old and new" in that volume.

obviously practical. The aim is to keep an actual enterprise on the road, with one eye on the ideal and another on the facts about the limited human beings involved. But much contemporary philosophy is not closely connected to the philosophy of science. It is often not clear what the *point* of the enterprise is, besides saying intelligent things about issues, such as scepticism that have become part of the philosopher's job description. The result often illustrates or puts pressure on positions in the philosophy of language or the philosophy of mind. (And often there is a more visible point to these positions.) Recently there has been an interesting development of 'virtue epistemology', which studies the traits of people which advance and hinder their epistemic projects. Virtue epistemology is concerned both with the traits that people actually exhibit, as described in terms of what we know about human psychology, and with our standard vocabulary for describing such traits. The practical point is pretty clear here. We need a regulatory discourse in our non-scientific belief-acquiring activities, and we so it helps to know what our actual meta-epistemic discourse is, and how well it works. My own view is that virtue epistemology ought to focus on fine-tuning our set of epistemic virtues and vices. Are there traits which we take to be virtues which actually are not? Are there virtue-terms which would be more effectively deployed if their emphasis were changed? My own suggestion about the most promising focus for these concerns is our handling of our own finiteness. Do we have a useful vocabulary for describing

how successful a person is in, for example, the rational response to her own particular kind and degree of irrationality³?

Example two: the philosophy of mind There are two rather different branches to the philosophy of mind. Both are shaped by the existence of what is now usually, perhaps unfortunately or misleadingly, called 'folk psychology': the body of beliefs, intuitions, explanatory routines, and belief-forming procedures that we use in everyday life to understand one another. (So I could have just called it 'the concepts of mind'.) One branch of the subject then takes a semantic, external, angle on folk psychology. This asks what real facts folk psychology might be about, and whether it represents them accurately. So questions about mind and brain are central, but also what Thomas Nagel has called the mind-body question, concerning the relation between folk psychology and the account of mind given by a not necessarily brain-centered scientific psychology. (To a very first approximation at the extremes are versions of dualism which say that the facts are just as folk psychology presents them, and eliminative materialism which says that the facts are so different that folk psychology has no real hold on them.)

³ For a more detailed account of epistemology along these lines see my "Recent work in epistemology" *British Journal for the Philosophy of Science* 51, 2000. For virtue epistemology see Christopher Hookway "Epistemic norms and theoretical deliberation", *Ratio* 12, 1999, pp. 380-398, and Linda Zagzebski, *Virtues of the mind* (Cambridge: Cambridge University Press, 1996).

The other branch takes a structural, internal, angle on folk psychology. It asks about the relations between the different folk psychological concepts, and about their roles in explaining actions. Belief, desire, will, memory, sensation, pleasure, anger, fear, remorse: the list goes on. Each of these concepts is puzzling in well-known ways, the links between them are puzzling, and the way that they interact to form the complex of ideas that lies at the heart of human life is still largely mysterious.

Neither angle excludes the other, of course, though most works in the philosophy of mind fall clearly into one branch or the other. In both cases our focus is on our ordinary mental concepts. From the external angle we ask questions about the whole set, with respect to another privileged set, those of neurology or cognitive psychology. From the internal angle we ask about particular concepts, usually expressing our accounts in terms of other common sense psychological, descriptive, or normative concepts. Certainly the external angle suggests more drastic possibilities. Perhaps we ought to junk all talk of mind, perhaps we ought to operate with mental concepts as we would with useful but ultimately indefensible myths, perhaps we should try to formulate a hybrid mental/physical explanatory scheme. But if we take any of these possibilities seriously we are faced with detailed practical questions likely to snag on the fine grain of human life. How are we to instruct our children, amuse our friends, make our contracts, carry out all kinds of business, in the revised set of

concepts? In the face of these it is easy to see the appeal of a quietism such as is found in the work of Davidson and some followers of Wittgenstein. The line is that mental and physical concepts are so different in conceptual style that they cannot be mixed. Assuming there is no way in which we can ignore the psychological altogether, this would seem to leave us no alternative but to go on folk-psychologizing as before, with however bad a conscience⁴.

There is an assumption here, though, about the relation between the external and the internal angle. Quietism assumes that if we are stuck with folk psychology we are stuck with folk psychology as it is and always has been. But in thinking this it is accepting an assumption also made by eliminativism. For it is very far from obvious that there is any such thing as a single unchanging folk psychology. Our everyday ideas about mind change and vary, under social, scientific, and moral pressure. If we decide that the facts about brain and cognition do not fit the presuppositions of our ordinary thinking about mind, then we are faced with an extremely engineering-ish problem. How are we to do justice to the facts as we understand them while still preserving a structure which allows us to serve the varied range of tasks for which we now deploy mental concepts? And the facts surely do not fit the presuppositions - the only

⁴ No short list could do justice to the variety of work in the philosophy of mind, but the issues discussed in this section are discussed in David Rosenthal's introduction to his collection *The nature of mind* (Oxford: Oxford University Press, 1991). The essays by Ryle, Strawson, Anscombe, Davidson, and Feyerabend in that collection are classics representing the variety of relevant views. Issues about folk psychology are well represented in the essays and introduction

question is the size of the discrepancies - so we are surely faced with these practical questions. And in their mole-like ways philosophers are slowly feeling out the fine-tunings and alternatives that we will need.

Example three: the philosophy of language The most abstract part of philosophy, the most rarefied, according to some the most central, nevertheless the part about whose philosophical character students and non-philosophers are most often have doubts. Like the philosophy of mind there is an outside and an inside angle. Looking at language from the outside we try to understand truth, reference, communication, and various varieties of meaning, and try to see how these could have their origins in human conventions and rule-governed human practices. Looking at language from the inside we consider specific aspects of the languages we speak: quantifiers, names, demonstratives, vague predicates, natural kind terms. We try to understand how these relate to one another, and one by one we try to explain how it works, in terms of the others. The two angles are less separate than they are in the philosophy of mind. We try, for example, to understand whether there could be a language without vagueness or without demonstratives, and we wonder whether particular theories of truth apply more plausibly to languages exploiting some of these devices than those without them.

Some of the target concepts here clearly play a role in speakers' regulations of their linguistic activities. Meaning, notably. We ask whether someone really meant what they said, whether a word was used with one meaning or another, or what the meaning of an unfamiliar word is. It is pretty doubtful that the implicit assumptions behind these regulative practices can be explained in terms of one single concept of meaning, and so one vital philosophical task is to separate out the different meanings of meaning. The articulated collection of concepts of sense and significance that a good philosophical analysis of meaning would produce, should be capable of supporting a sophisticated linguistic meta-discourse. For example it should allow us to make a helpful separation between what words literally mean, what someone might reasonably be expected to have communicated with them on some occasion, and what the actual communicative intention was. It should allow us to do this better than our present vague talk of meaning. If not, it is not much of an analysis.

Other concepts in the philosophy of language play a central role in practices that relate language to other activities. Truth, for example, enters when we try to distinguish between norms, conventions or aims that are targeted at some other moral or social purpose from those that aim at believing that p when p . For example it might be wrong so assert that p , or even to think

that p , though the belief that p is in some sense justified. (p might be a fact about someone's private life, which you have no business discussing, indeed no honest reason even to be curious about, but, still, there is evidence for it.) If we try to explain this sense of justified, in which a belief can be justified even though it would be wrong to put oneself in the way of holding it, we will inevitably have to say something like: beliefs are in this sense justified when they result from processes which typically results in true beliefs.

Truth is so basic and simple an idea that it is hard to see how a philosophical theory could propose an alternative to it. (But perhaps that is a remark simply about the limits of our, or my, imagination.) But there is a lot of room for variation in the concepts that surround and connect to it. Tarski's account of truth links the truth of whole sentences to the satisfaction of their component parts. ("there are cats" is true if and only if something satisfies "cats".) And by doing this it suggests that behind the concept of truth there are other more fundamental concepts relating in more specific ways to the correspondance between words and world. One reason the concept of satisfaction is important is that it applies to the components of non-assertoric sentences, questions and commands. It thus suggests that analogs of truth apply even when we would not normally say "is true". The command "Let there be cats" is obeyed if something comes to satisfy "cat", and then the question

"Are there cats?" is answerable in the affirmative. So we can see a philosophical basis for a proposal that we call these true commands and questions.

Related to this, there are kinds and degrees of truth. There is truth in fiction, there may even be poetic truth; there is objective truth and truth that is the result of some convention about how we are to speak. It is not at all obvious what is the best and most coherent way of qualifying "true". One desideratum is keeping the relations between science and everyday belief clear (we may want to keep well apart claims that there is a real property of objects like a which a instantiates from claims that the rules of language allow a correct assertion of "a is P at t".) Another is allowing us to signal clearly the import of our communicative intentions (do we mean that p is literally true, or just that 'p' would be an adequate way of capturing some of the consequences of the facts?). There are clearly many others.⁵

A general pattern Similar remarks could have been made about most other areas of philosophy. I have not discussed ethics and political philosophy because with them the corresponding conclusions hardly need argument. Moral and political concepts clearly form networks whose structure is hard to

⁵ For a general discussion of recent accounts of truth see Richard Kirkham *Theories of truth: a critical introduction* (Cambridge, MA: MIT Press, 1992.). The issues of this section are closely related to the controversy over minimal theories of truth. See Paul Horwich *Truth* (Oxford: Blackwell, 1990), and my review of Horwich, *Philosophical Books* 32, 1991, 231-233.

understand, but which we need to negotiate in order to live our practical lives. So there are two vital tasks. First there is the task of giving us the information we need about the structure as it is, in order to find our way about when using it. (So students in beginning ethics courses not only learn the differences between standard positions about the nature of The Moral; they also learn for example the difference between something's being a right action and someone's having a right to do it, or the difference between an act's being wrong and its being right to prevent someone doing it.) The second task is that of tweaking, modifying, or fine-tuning the structure so that it can do its job better. (So we have debates about whether our intuitions concerning when it is permissible to cause one death to prevent others show something essential about moral judgement or are just quirks of the way we happen to think now⁶.) Both of these are very practical tasks, but neither can be carried out effectively without assumptions about the relevant facts, which determine whether and how a system of concepts can function. And neither can be carried out without assumptions about the structure of the system. We can in most cases only make the roughest of assumptions about either.

Here is a different way of putting it. Most philosophy is concerned with our box of conceptual tools. Sometimes it aims to sharpen tools already in the

⁶ See the symposium between Frances M. Kamm and John Harris "The doctrine of triple effect", *Proceedings of the Aristotelian Society*, supp. vol. 74, 2000. Also relevant is Tamara Horowitz "Philosophical intuition and psychological theory" in Michael R DePaul and William Ramsey eds. *Rethinking Intuition* (Rowman & Littlefield, 1998) pp. 142-160.

box, and sometimes it aims to add or subtract tools with an eye to the jobs to be done and the materials involved. This way of putting it has an advantage. It makes it no surprise that we are dealing with something at the same time very practical and very abstract. For the tools are *concepts*. They concern people's general patterns of thought, their use of implicit beliefs that cannot be mechanically expressed in available vocabulary, and their expectations that others will conform to general and often very subtle rules.

Patterns of thought, structures of concepts, cannot be read off the surface of what people say or how they react to particular cases. Suppose that there is a concept which if perfectly understood would lead to some simple conversational rule. (The concept might be 'knowledge' and the rule might be that you can never say 'she knows it but it is false.'). It does not follow that every person who possesses the concept will adhere to the rule, and expect everyone else to. For in some contexts people will not recognize the rule as applying, in others they will take it to be over-ruled by another, in others they will aim at following it but will typically fail. ('I know' also tends to indicate confidence in truth, so we can easily fall into taking confidence as necessary for knowledge, so that we can sometimes speak as if confidence were also sufficient for knowledge.) So a conjecture that concept related in a particular way lie behind the way we think and talk nearly always is indeed a conjecture, hard to verify and rarely exactly right. It is always an approximation to the complex

relation between some abstract never perfectly instantiated idea and the improvisations that link it to actual practice. One does not do justice to performance without seeing competence behind it; one cannot represent competence other than as something very different from performance.

Digression: concepts I might seem to have argued my case at the price of buying a controversial theory. For I have argued in terms of *concepts* and the *beliefs* they allow us to have. Moreover, I have spoken of better concepts that allow us to have beliefs that are in various ways better. And I have talked of the structures in which concepts are embedded. This sounds as if I am assuming that there are particular identifiable things called concepts such that there is a definite answer to the question whether the concept one person has, and uses a word to express, is the same as a concept used by another person. And while that assumption would not be at all obviously false, it would be far from obviously true. Certainly many philosophers have reasons for disbelieving it. It would be hard to reconcile with Quine's skepticism about meaning, for example. So it would be a risky assumption to use in supporting a claim about the nature of philosophy.

In fact, I think what I am arguing is consistent with a wide range of views about concepts and beliefs. I certainly need the following assumptions:

People assert and deny sentences of public languages, which can be true or false. The truth or falsity of a sentence often depends on properties of the words it contains, and the capacity to make true or false assertions depends on capacities to use these words in specific ways. When people assert sentences they cause one another to enter into states associated with these sentences, and these are closely related to the states between which people move when they think. Communication and thinking produce practical effects, and different ways of communicating and thinking, producing different states of mind, produce different practical effects. Some such states, and the use of sentences associated with them, can have better practical effects than others. Some states and sentences can be used to express truths that others can not.

I do need all these assumptions. But they are uncontroversial, within and beyond philosophy. They could be accepted by someone who believed in concepts and beliefs but individuated them extensionally, so that any two concepts true of all the same things are identical. And everything I have argued is consistent with this interpretation of concepts. The claims could even be accepted, with a little rewording, by someone who denied that thought is in any way conceptual. (The assumptions above do not use the words 'concept' or

'belief'.) For they are fundamentally claims about the activity of thinking, and its capacity to be influenced by the way we use our words⁷.

Rounding-off: perverse philosophical self-denial Some of these issues come to a head when the philosophical question turns on the relation between common sense and science. Then we can be crippled by finding ourselves caught between two wrong pictures. The more dominant is probably the one that comes from Quine. According to this view common sense is just a theory, like any other, which may be a usable approximation to the truth and in fact may be more suited for everyday practical activities than the theories that we have more evidence to believe are actually true. This picture cripples us when we consider possible changes in common sense. If a philosopher is to argue for a change in view then she should argue for what we have more reason to believe, and thus for bringing the common sense theory nearer to the scientific one. Moreover most of the evidence is already in the hands of the scientists (be they physicists, psychologists, or statisticians), so that the role of philosophers is reduced to that of popularizers and mediators. Almost as inhibiting is the opposed influential view, the one that comes from Strawson. According to this

⁷ For current work on concepts see Christopher Peacocke *A study of concepts* (Cambridge, MA: MIT Press, 1992), Andrew Woodfield, "Do your concepts develop", in Christopher Hookway and Donald Peterson, eds. *Philosophy and cognitive science*. (Cambridge: Cambridge University Press,

view common sense is the source of its own distinctive concepts, on which it imposes apriori constraints. There is no disparity between science and common sense, on this view, because they are incommensurable. One task of the philosopher is to explore the apriori conditions that thinking must respect if it is to remain within the common sense framework. But it cannot stretch or re-shape that framework⁸.

Applied to the most important interfaces between science and common sense, notably to ethics and to folk psychology, both these views are disastrous. The Quinian view suggests that we can only improve folk psychology by turning it into experimental psychology. And the Strawsonian view suggests that we cannot improve it without tampering with its essence, so that it is no longer commonsensical, no longer folk. On both views the idea of moral or folk-psychological progress seems to make no sense. My reasons for disagreeing with both views should be clear by now. Each focusses on a particular task, which is only a small part of what we can and should expect of philosophy. The Quinian ideal focusses on acquiring true beliefs, about nature, mind, and everyday life. The Strawsonian ideal focusses on understanding how we think now. Both focus on the concepts we actually have now. But we can also ask about concepts we could have, and routes we could take from here to there.

1993), and Jerry Fodor *Concepts: where cognitive science went wrong* (New York: Oxford University Press, 1998.)

⁸ See W. V. Quine "Epistemology naturalized", in *Ontological Relativity and other essays* (New York: Columbia University Press, 1969, and P.F. Strawson *Individuals* (London: Methuen, 1959.)

When we ask about possible concepts we realise how little we know about actual ones. We realise that giving an up-to-date account of the physical world is in a way easier (after the physicists have done the hard work) than describing the structure of the concepts that either physicists or non-physicists appeal to in forming their beliefs. The difficulty of the descriptive task emerges when we address the practical aim of finding the best ways of expressing and managing what we currently believe, or of removing obstacles to things we would gain from believing. Then we find ourselves grateful for any partial understanding that allows us even a little bit of progress⁹.

There are thus important consequences to saying that these central parts of philosophy are more like Engineering than like Science. While respecting considerations about what is true or supported by evidence their main focus is on the question "will it work?" In the case of ethics this suggests that we have two subjects, and it might sometimes help to keep them apart. The project of one of them, call it moral science, is to find the intellectually best supported positions about what we should do and what it means to say that that is what we should do. Taken like this, it is a very new project. As Derek Parfit points out, though humans have been formulating moral systems for millennia the idea

⁹ Nelson Goodman appreciated this point clearly in *Fact, fiction, and forecast* (Indianapolis: Bobbs-Merrill, 1955), when he urged us not to invent subtle justifications of induction in our ignorance of the patterns of inductive reasoning that we actually follow.

of this project has been clear to only a few people during a few scattered moments of history¹⁰.

In contrast, the project of ethical engineering is to find systems of moral ideas that we can use to think through moral problems in the context of everyday life. We want to find theories, norms, ideals, and strategies, which will in practice lead us to satisfactory outcomes if numbers of us decide to try them. This project can be carried out without a profound or exact definition of what would be a satisfactory outcome. We simply need to be able to identify things that are less than perfect about the way we now act and think. And finding things that we now do badly is easy, though deep and helpful diagnoses of the sources of our troubles are not so easily had.

¹⁰ See section 154 of the Concluding Chapter of Derek Parfit *Reasons and persons* (Oxford: Oxford University Press, 1984.)