

THE COMPUTER REVOLUTION IN PHILOSOPHY (1978)

Aaron Sloman

[Book contents page](#)

This chapter is also available in [PDF format here](#).

CHAPTER 3

SCIENCE AND PHILOSOPHY

3.1. Introduction

Immanuel Kant's *Critique of Pure Reason* is widely acknowledged to have been a major contribution to philosophy. Yet much of the book can also be seen as an early contribution to theoretical psychology. For example, his claim that no experience is possible without some form of prior knowledge (a claim to which we shall return in the chapter on perception) is relevant to psychologists' attempts to understand the nature of perception and learning. His notion that perception and imagination require the use not of picture-like templates, but of rule-like schemata for analysing and synthesising images, has been re-invented by psychologists in this century and plays an important role in computer-based theories of perception.

So Kant's work illustrates the overlap between science and philosophy. There are many more examples. Einstein's approach to the analysis of concepts of space and time was influenced by his reading of empiricist philosophy. Frege's attempts to answer some of Kant's questions about the nature of arithmetical knowledge led him into logical and semantic theories and formalisms which have deeply influenced work in linguistics and computer science. Marx's sociological theories were partly based on Hegel's philosophy. More recently, work by philosophers of language, such as Austin and Grice, has been taken up and developed by linguists, and the psychologist Heider has acknowledged the influence of Ryle's *The Concept of Mind*.

Philosophers' analyses of some of our most general concepts, such as *cause, individual, action, purpose, event, process, good, and true*, are relevant to biology, to anthropology and developmental psychology, whether or not practitioners of these subjects are aware of this.

For instance, biologists studying the evolution of intelligence need to grasp what intelligence is, and how it includes the use of some or all of these concepts. A comprehensive anthropology would include cross-cultural studies of the most general and basic systems of concepts used by different peoples. And if developmental psychologists were to do their job properly they would spend a lot of time exploring such concepts in order to be able to ask deep questions about what children learn and how. (Piaget did this, to some extent. But I am not aware of university courses in developmental psychology which include training in conceptual analysis.)

Within artificial intelligence it is not possible to avoid philosophical analysis of such concepts, for the discipline of trying to design machines which actually behave intelligently and can communicate with us forces one into analysis of the preconditions of intelligent behaviour and our shared presuppositions. For otherwise the machines don't work!

These illustrations of the connections between philosophy and the scientific study of the world are not isolated exceptions. Rather, they are consequences of the fact that the aims and methods of philosophy overlap with those of science. In this chapter I shall try to analyse the extent of that overlap.

3.2. The aims of philosophy and science overlap

In particular, the greatest philosophers have shared with the greatest scientists the first three 'interpretative' aims listed in chapter 2, namely the aim of developing good concepts or thinking tools, the aim of finding out what sorts of things are possible, and the aim of trying to explain these possibilities. Their methods of pursuing these aims have much in common too, as will be shown below.

A fourth major aim that they appear to have in common is the aim of discovering limits to what is possible, and explaining such limits. However, in relation to this aim, the methods of scientists and philosophers tend to be rather different, insofar as philosophers often try to set up non-empirical demonstrations. And they usually fail.

By exploring the relationship between the aims and methods of science and philosophy we shall explain how it is possible for philosophy to be the mother of science, thereby perhaps making a philosophical contribution to the science of intellectual history.

Let us start with some reminders of the kinds of questions which have exercised philosophers. I shall ignore the many pseudo-questions posed by incompetent philosophers who cannot tell the difference between profundity and obscurity.

3.3. Philosophical problems of the form 'How is X possible?'

Many questions of the form 'How is X possible?' have been asked by philosophers. Some of them go back to the ancient Greek philosophers, or further, while others were first formulated much more recently. It will be seen from the long list of examples which follows that more and less specific versions of the same question can be asked. Detailed analysis in philosophy or science leads to the formulation of more specific questions, concerned with more of the fine-structure of X. Increasing specificity increases the scientific interest of the question. In particular, it should be clear that although the first question listed below is a philosophical one, more specific versions of it (e.g. the fourth one) look much more like psychological questions.

1. How is knowledge possible?
2. How is empirical knowledge possible?
3. How is it possible to acquire knowledge about the material world on the basis of sensory experience?
4. How is it possible to learn, from seen two-dimensional surfaces, that an object is three-dimensional and has unseen surfaces on the far side?
5. How is it possible to know anything about the past, the future, unobserved objects or processes, or other people's minds? (Cf. 16).
6. How is it possible to know that two events are causally connected?
7. How is it possible to know laws of nature or their explanations?
8. How is it possible to know truths of logic and mathematics?
9. How is it possible to know conditional truths, of the forms 'If P then Q' or 'If P had been the case then Q would have been'?
10. How is it possible for an infant knowing no language to learn one?
11. How is it possible to learn the language of a culture other than one's own?

12. How is it possible to use strings of symbols to describe our multi-dimensional world?
13. How is it possible to interpret flat patterns as pictures of solid three-dimensional objects?
(Compare question 4.)
14. How is it possible to use general concepts, such as *dog*, *triangle*, *game*, *taller*, or *between*, which apply to a very varied range of instances?
15. How is it possible to learn the names of, think about, or refer to, remote persons, places or events?
16. How is it possible to think about the past or future events? (Cf. 5.)
17. How is it possible to think or talk about non-existent things, such as Mr. Pickwick, Ruritania, the accident that nearly happened this morning, or the choice I considered but did not make?
18. How is it possible to think about minds other than one's own, or about another person's emotions or sensations?
19. How is it possible to have idealised concepts which go beyond the limits of what we can experience, such as *perfectly thin*, *perfectly straight*, *perfectly parallel*, *exactly the same shade of colour*, *exactly the same weight*, or *exactly the same shape*?"
20. How is it possible for a finite mind to think about such infinite totalities as the set of integers, the set of points on a line, or the set of all possible English sentences?
21. How is it possible to have a concept of a causal connection which is more than the concept of an instance of a well-confirmed regularity?
22. How is it possible to understand scientific theories referring to things which can never be perceived?
23. How is it possible to understand metaphors?
24. How is it possible to understand metaphysical questions?
25. How is it possible for a person, or a culture, to discover that its conceptual system is inadequate, and improve it?
26. How is it possible for there to be valid reasoning which is not purely logical, such as inductive reasoning or reasoning using diagrams?
27. How is it possible for an identity-statement, such as 'The Evening Star is the Morning Star', to be true, yet have a different significance from another identity statement referring to the same thing, such as 'The Morning Star is the Morning Star'?
28. How is it possible for two predicates, such as 'polygon with three sides' and 'polygon with three angles', to describe exactly the same set of objects yet have different meanings?
29. How is it possible for there to be formal, or syntactic, tests for valid (truth-preserving) reasoning?
30. How is it possible to have knowledge which one can use yet not formulate (e.g. knowledge of how one's native language works, or knowledge of the difference between Beethoven's and Schubert's styles of composition)?
31. How is it possible for there to be minds in a material universe?
32. How is it possible for physical and chemical processes to influence or even give rise to such things as sensations and feelings? (or vice versa?)
33. How is it possible for there to be such a thing as self-deceit, or unconscious beliefs, attitudes, desires, fears, or inferences?
34. How is it possible for a set of experiences, either at the same time or at different times, to be the experiences of one mind?
35. How is it possible for a set of experiences, beliefs, thoughts, decisions, intentions, and actions all to 'belong' to one mind?
36. How is it possible for deliberation, choice, or decision to exist in a deterministic universe?
37. How is it possible for a mind to continue to exist while unconscious?
38. How is it possible to think of oneself as being in a world whose existence is independent of one's

- own (or any mind's) existence?
39. How is it possible to distinguish moral or aesthetic evaluations from personal likes or dislikes, or to think rationally about moral problems?
 40. How is it possible to use moral language of a kind which does not reduce to descriptive or emotive language?
 41. How is it possible for a norm to exist in a community without being accepted by any individual in the community?
 42. How is it possible for democracy to exist in a state with millions of citizens with different and conflicting aims?
 43. How is it possible for a social system to be just?
 44. How is it possible rationally to weigh up short term and long term harm and benefit?
 45. How is it possible to search in a sensible direction for the solution to a problem without knowing what form the solution will take?
 46. How is it possible for an event to be temporally related to another distant event?
 47. How is it possible to identify and reidentify places?
 48. How is it possible for objects to change their properties and relationships (and remain the same objects)?
 49. How is it possible for there to be anything at all?
 50. How is it possible for people to invent philosophical problems?

Many of the questions in the list have controversial presuppositions: it is often disputable whether the X in 'How is X possible?' is possible at all! Many attempts have been made to prove the impossibility of some X, for instance where X = meaningful talk about God or infinite sets, or rational discussion of moral issues, or even such obviously possible things as: change, a man over-taking a tortoise in a race, knowledge about the past, knowledge about material objects, or deliberation and choice.

Lunatic though it may at first appear, serious thinkers have put forward demonstrations that these are impossible. Equally serious thinkers have put great intellectual effort into attempts to refute such demonstrations. The process may appear a waste of time, but has in fact been very important. The discovery, analysis and, in some cases, refutation of such paradoxical proofs of impossibility has been a major, though haphazard, stimulus to philosophical progress and the growth of human self consciousness. It leads to a deeper understanding of the phenomenon whose possibility is in dispute. In some cases (e.g. Zeno's paradoxes) it has even led to advances in mathematics.

Often, a philosopher asks 'How is X possible?' only in the context of asking 'What is the flaw in so and so's alleged proof that X is impossible?' But there is also a more constructive philosophical tradition, first consciously acknowledged by Immanuel Kant, of granting that X is possible and attempting to explain how it is, in the light of careful analysis of what X is. This is the philosophical activity which merges into scientific theorising.

In what follows I'll try to analyse the similarities and differences in aims and methods: a step towards a scientific theory explaining the possibility of the growth of scientific and philosophical knowledge.

3.4. Some general similarities and differences between science and philosophy

One of the differences between science and philosophy concerns the range of possibilities attended to. Philosophers have mostly been concerned with possibilities which are known to everyone or at least to educated laymen in their community, whereas one of the characteristics of scientific research is that it uses sophisticated apparatus and techniques, and highly specialised explorations, in order to discover

new possibilities which are not discoverable simply by reflection on common experience.

I do not mean that all the possibilities discussed by philosophers are obvious: they may be known to all of us without our realising that we know them (like the possibility of truly unselfish action). Some of the things we know are not evident to us until we have engaged in the philosophical activity of digging up unacknowledged presuppositions. For instance, most people if simply asked how many different kinds of uses of language there are, are likely to come up with only three or four, such as the text-book favourites: exclamations, questions, commands and assertions (statements). But even though they do not think of more without prodding, they do in fact know of many possible uses of language not covered by this list, such as betting, congratulating, pleading, exhorting, warning, threatening, promising, consoling, reciting, calling someone, naming someone or something, welcoming, vowing, counting, challenging, apologising, teasing, declaring a meeting open or closed, and several more. (See J.L. Austin, *How to Do Things With Words*.)

Similarly, there are many psychological possibilities which we all know about, but do not find it easy to recall and characterise accurately when theorising about the mind. I shall draw attention to many examples in later chapters. So, both philosophy and science use specialised techniques to find out what sorts of things are possible, but their techniques and consequently the ranges of possibilities unearthed, are different. Philosophers dig up what we all know, whereas scientists mainly to extend what we know, about possibilities.

In both cases a preliminary characterisation of a kind of possibility is subject to correction, in the light of an explanatory theory.

One of the faults of philosophers is that they tend to ask questions which are not nearly specific enough. If one simply asks 'How is knowledge possible?' or 'How is knowledge of other minds possible?', these questions do not explicitly specify the requirements to be met by explanatory answers, since they do not describe in sufficient detail what is to be explained. They specify many requirements implicitly, because we all know a great deal about the possibilities referred to, but until they have been described explicitly, people can unwittingly select different subsets for consideration, and so philosophical debates often go on endlessly and fruitlessly.

The criteria listed in [Chapter 2](#) for assessing explanations of possibilities, presuppose that there are detailed specifications of the range of possibilities to be explained. Otherwise there is no agreed basis for assessing and comparing rival theories. This preliminary analysis of the range of possibilities to be explained is often shirked by philosophers.

Even when philosophers do a fairly deep analysis, it is not presented in a systematic and organised form but rather in the form used for literary essays. The result is that philosophers often simply talk past each other. (This also happens in psychology for similar reasons, as may be confirmed by looking at the cursory 'definitions' of mental concepts such as *emotion*, *memory*, *perception*, *learning*, etc., which precede lengthy chapters on empirical results and proposed theories.)

In both philosophy and science, if progress is to be made, and seen to be made, the task of constructing an explanation of the possibility of X must be preceded by at least a preliminary characterisation of the range of possible kinds of X. This preliminary characterisation may be based on close examination of a wide variety of examples of X, taken from common experience, in the case of philosophy, or from specialised experiment and observation. The specification may include such things as the types of components, the types of organisation of those components, the types of behaviour, the types of functions, and the types of relations to other things, found in specimens of X, i.e. internal and external structures, functions and relations. In both philosophy and science, the construction of an explanatory theory will suggest ways of improving or correcting such 'observations'.

Having got a preliminary characterisation, that is, a preliminary answer to the question: What sort of things are X's? or What sort of X's are possible?, the scientist or philosopher can then begin to construct a theory describing or representing conditions sufficient to generate the possibility of instances of X (knowledge, perception, truth, scientific progress, change, falling objects, chemical processes, or whatever it is whose possibility is to be explained). Whether one is a philosopher or a scientist, the conditions for adequacy of an explanatory theory, and the criteria for comparing the merits of rival explanations of a range of possibilities are the same, namely the sorts of criteria listed in chapter 2.

Despite the overlap, there is an important difference. Often philosophers are content to find some theoretically adequate explanation of a set of possibilities without bothering too much whether it is the *correct* explanation. So they ask 'How *might* X be possible?' rather than 'How *is* X possible?', or 'What *could* explain the possibility of X?' rather than 'What *does* explain the possibility of X?' However, every answer to the latter necessarily includes an answer to the former, and in that way science subsumes philosophy, which is very like the relationship between A.I. and psychology (see chapter I). In practice, the difference between the two approaches becomes significant only when alternative answers to the first question have been formulated, so that something can be done to find out which is a better answer to the second.

3.5. Transcendental deductions

When one has such a theory T explaining the possibility of X's the truth of T *is a sufficient* condition for the possibility of X. However, it may not be the *correct* explanation, for instance if T itself is false. In general it is not possible, either in science or in philosophy, to establish conclusively that some theory is true: the most one can do is determine which, if any, of several theories is, for the time being, best. And even that is not always possible when a subject is in its infancy.

However, some philosophers have not been satisfied with this, and have tried to show that no other theory besides their own could possibly give the correct explanation. An argument purporting to show that T is not merely *sufficient* to explain the possibility of X, but also *necessary*, is called a 'transcendental argument'. (As far as I know, this notion was invented first by Kant.)

No attempts to construct valid transcendental arguments have so far been successful. For instance, Kant tried to show (in *Critique of Pure Reason*) that explaining the possibility of distinguishing the objective time order of events from the order in which they are experienced must necessarily involve assuming that every event has a cause; but quantum physics shows that one can get along without assuming that every event has a cause. Strawson tried to show (in *Individuals*) that our ability to identify and re-identify material objects and persons was a necessary part of any explanation of the possibility of identifying other things such as events, processes, states of affairs, pains, decisions, and other mental phenomena.

But he made no attempt to survey all the possible theories which might one day be formulated, including the varieties of ways in which computers or robots (and therefore people) might be programmed to use language, and his arguments seem to be irrelevant to the detailed problems of designing mechanisms with the ability to refer to and talk about things. (This criticism requires further elaboration.)

Such attempts at transcendental deductions are over-ambitious, for to prove that some theory T is a *necessary* part of any explanation of the possibility of X would require some kind of survey of all possible relevant theories, including those using concepts, notations and inference procedures not yet developed. It is hard to imagine how anyone could achieve this, in science or in philosophy. Scientists

rarely try: They are not as rash as philosophers.

One reason why philosophers feel they must bolster up their explanations with 'transcendental arguments' is that they dare not admit that philosophy can be concerned with empirically testable theories, so they try to show that their theories are immune from empirical criticism. However, I shall show below that this is inconsistent with the practice of philosophers.

We now look a little more closely at similarities and differences between methods of science and philosophy.

3.6. How methods of philosophy can merge into those of science

The procedures by which philosophy can make progress in the task of describing and explaining possibilities shade naturally into scientific procedures. So by describing such philosophical procedures and the processes by which they transform a problem, we begin to explain how it is possible for philosophy to contribute to science. The overlap with AI (when AI is done well) is specially significant.

The relevant philosophical procedures concern the following:

- a) Collection of information about what sorts of things are possible,
- b) Construction of new characterisations or representations of those possibilities (i.e. answers to the question 'What is X?'),
- c) Construction of explanations of those possibilities, and finally testing and refinement of explanatory theories. This last step can, as in all science, lead back to modifications of earlier steps.

A first step is collecting information about the range of possibilities to be explained. For instance, before attempting to explain the possibility of knowledge one must ask 'What is knowledge?'. This involves collecting examples of familiar kinds of knowledge, and classifying them in some way. (Knowledge of particular facts, knowledge of generalisations, knowledge of individuals, knowing how to do things, etc.) Closely related possibilities should also be surveyed, e.g. believing, learning, inferring, proving, forgetting, remembering, understanding, doubting, wondering whether, guessing, etc. Functions of knowledge can then be listed and classified.

All this gives a preliminary specification of some of the *fine structure* of the range of possibilities to be explained, an answer to the question 'What is X?' (or, 'What are X's?'). One can go on indefinitely attempting to improve on the preliminary specification, by covering a wider range of cases, giving more detailed specifications of each, and revising the classification.

This process may at first rely only on what Wittgenstein (in *Philosophical Investigations*, Part I, section 127) called 'assembling reminders'. These are examples of possibilities which when stated are obvious to common-sense, since we have all experienced similar cases, though we may not find them easy to think of on demand, like the examples of possible uses of language noted above. Much analytical philosophy, and most of Wittgenstein's later philosophy, consists of this kind of common-sense natural history.

An obvious extension of this activity is the use of experiments, instruments, measurement, fieldwork, and other tools of science to find and describe new examples of X, or new facts about old examples. [Chapter 1](#) explained how artificial intelligence can contribute to this fact-gathering process in philosophy by providing examples of new forms of behaviour.

So the fact-collecting of philosophers merges into the fact-collecting of scientists. However, new empirical research may be premature if common sense knowledge about possible sorts of X's has not yet been made explicit and systematised. (Hence the futility of much psychological research, e.g. on decisions, learning and emotions.) So philosophical methods of analysis should come first in cases where relevant information is part of common sense for instance in the study of mind and society. (Some linguists have appreciated this, but few psychologists or social scientists. Fritz Heider was a notable exception: see his *Psychology of Interpersonal Relations*.)

In philosophy, as in science, fact collection is rarely useful unless guided by a problem or explanatory theory. The mere collection of possibilities is of little interest except insofar as a theory can be found to explain and organise them. And theories are important only if they help us solve problems or puzzles. How theories are generated is still largely an unsolved problem. No doubt chance plays a role, but individuals like Kant, Einstein and Newton would not have made so many theoretical advances if they had not employed (albeit unconsciously) rational procedures for making the best of chances which came their way.

Artificial Intelligence in its attempts to design intelligent (i.e. rational?) learning planning and problem-solving systems necessarily overlaps with philosophical attempts to explain the nature of theories and theory formation (as outlined in [Chapter 2](#))

3.7. Testing theories

Once a theory T has been found which meets some or all of the criteria listed in the [previous chapter](#) (see sections 2.5.4-6) for explaining the possibility of X's, the question arises whether it is the *correct* explanation. Whether in philosophy or in science, answering this question requires testing the theory on new examples of X, or new, more detailed, descriptions of old examples, in order to see whether it is sufficiently general and explains enough fine structure. The theory can also be related to other known facts to see whether it is inconsistent with them and therefore false: i.e. its plausibility can be tested.

Emotivism is a philosophical theory purporting to explain how it is possible to use moral language meaningfully. However, fact-collecting of the sort described above showed the theory to be insufficiently general, for it was unable to account for facts about moral language which were not at first obvious to proponents of the theory, but are part of common sense. For instance, the theory interpreted moral language as performing functions like expressing the speaker's emotions, evoking similar emotions in hearers and causing hearers to act in certain ways. This fails to account for the empirically established possibility of unemotional hypothetical discussion among rational people of what, morally, ought to be done in certain situations. So the theory must either be rejected or modified to deal with this use of moral language. (I have listed a range of facts which theories like emotivism cannot account for, and proposed an alternative theory, in my two papers on 'better': see bibliography.)

This example refutes the widespread assumption that philosophical theories are not empirically testable. The assumption is probably based on a misconstrual of what philosophers actually do when they use empirical facts to test or support their theories: they use widely known common sense possibilities, rather than facts based on specialised empirical investigation. So the work can be done in an armchair no laboratory is needed, nor fieldwork. (The situation is similar when a linguist investigates his own language.) Because the information is so readily available its *empirical* nature is not recognised. (R.M. Hare made related points in his 'Philosophical Discoveries').

However, when the stock of relevant possibilities available to common sense is exhausted and has to be extended by more specialised empirical investigations, then philosophy merges into science. For instance philosophical investigations of the function of moral language and attempts to explain its possibility should, if properly conducted, overlap with linguistics and the psychology and sociology of morals. (Equally, the psychology and sociology, if done properly, would start with philosophical analysis of known possibilities.) For another example of philosophical use of empirical facts, this time from cognitive anthropology, see Bernard Harrison, *Form and Content*.

3.8. The regress of explanations

When a philosopher constructs his theory T, to explain a certain range of possibilities, it will not be long before someone asks for an explanation of the possibilities assumed in T. This may also lead towards scientific theorising and testing. For instance, Emotivism assumes (correctly) that it is possible for people to influence one another's actions and emotions by talking, and uses this to explain (wrongly) how moral language is possible. But the assumed possibilities also need explaining: and this leads directly into scientific studies of language and mind, e.g. studies of how utterances can influence attitudes.

Similarly, philosophers have often tried to explain the possibility of knowledge on the assumption that it is possible for things to be learnt from experience, and in particular that it is possible for ideas to become 'associated' with one another. But these assumed possibilities also need explaining, and this leads directly into scientific studies (in artificial intelligence and psychology) of ways in which information can be acquired and stored so as to be available for future use, and so as to enable one piece of information to 'evoke' another (which involves tricky problems of indexing and retrieval).

3.9. The role of formalisation

As specifications of phenomena to be explained become more detailed and cover a wider range of possibilities, so that more complex constraints have to be satisfied by the explanatory theory, it becomes necessary to invent special symbolisms and technical concepts in order to formulate theories which are sufficiently rich, powerful and precise.

In this way philosophy sometimes becomes more mathematical, as can be seen especially in the case of logic but also in philosophical studies of probability, in philosophy of science, and even in some branches of moral philosophy. Increasingly the formalisms of Artificial Intelligence will be used, as philosophical theories become more complex and precise, and too intricate to be evaluated without the aid of a computer. This parallels the ways in which scientific theories become more and more mathematical.

For instance, if, instead of the usual vague and general philosophical discussions of how perception can yield knowledge, an explanation is required *of how specific sorts* of perceptual experiences can yield knowledge *of specific sorts* of spatial structures, for instance an explanation of how certain views of a cube enable one to see that it is a cube with an interior and with faces on the far side, etc., then a mathematical formulation is inevitable. (N.B. 'Mathematical' does not mean 'numerical' or 'quantitative'.)

University courses in philosophy will need substantial revision if the appropriate theory-building and theory-testing skills are to be taught.

3.10. Conceptual developments in philosophy

In philosophy as in science, attempting to explain things can lead to new ways of looking at or thinking about the old facts, and this requires new sets of concepts. For example, the development of philosophical theories explaining the possibility of various uses of language can lead to criticism of old metalinguistic concepts or invention of new ones. Examples are: Kant's distinction between 'a priori' and 'analytic'; Frege's rejection of the subject/predicate distinction in favour of the function/argument distinction for describing sentence structures; the rejection by J.L. Austin and others of a four-fold classification of sentences into statements, questions, exclamations and imperatives; the discovery (explained, for instance, by J. Kovesi in his *Moral Notions*) that 'evaluative' is not a suitable label for the kinds of uses of language which have attracted attention in moral philosophy and aesthetics; modern criticisms of Kant's distinction between analytic and synthetic statements; and Kuhn's attempt to replace the concept 'scientific theory' with 'paradigm'.

My own attempt (in chapter 7) to replace crude distinctions between verbal and nonverbal symbolisms and reasoning processes with more precise distinctions is another example. My use of the concept of 'explaining how X is possible' is another. Further examples will be found in the chapter on numbers (chapter 8).

New concepts can change our view of what it is that we are trying to explain, so that a new specification is given of the old possibilities. Similar processes in the history of science have been described by Kuhn (1962, pp. 129-134), such as the change in the boundary between the concepts 'chemical compound' and 'physical mixture' resulting from the atomic theory of chemical composition.

In philosophy and in science, conceptual changes generate new specifications of what needs to be explained, and so can lead to new theories. The process of growth of human knowledge seems to be full of 'feed back' loops.

3.11. The limits of possibilities

I have said a lot about overlaps between aims of philosophy and the first three aims of science, namely the discovery, description, and explanation of possibilities. But science attempts also to find limits of possibilities: laws of nature. Is there a counterpart in philosophy?

Certainly some philosophers have tried to show not merely how things are but also how they must be or cannot be. Empiricists try to show that all significant knowledge *must* be based on sensory experience. Rationalists try to show that certain important kinds of knowledge *cannot* be empirical. Dualists try to show that there *must* be more than a material world if consciousness as we know it is possible. Logicians try to argue that mathematical concepts *must* be definable in terms of logic, if they are to have their normal uses. Moral or political philosophers often try to argue that their own moral or political principles must be accepted if morality or society is to be possible at all. Such theses are often based on attempts at 'transcendental arguments', which I have already criticised as over-ambitious, in the discussion of Kant, above.

Kant claimed to have unearthed various laws and principles which were part of the fundamental constitution of the human mind, so that all human thought and experience necessarily had to conform to them. However, such claims are very rash, in view of the fact that both biological and cultural evolution are known to be possible. We have already seen that thoughts that were impossible for ancient scientists are possible for modern scientists. The same contrast can be made between children and adults. This suggests that insofar as human minds have a 'form' limiting the nature of the world

they experience, this form can vary from culture to culture and from time to time in one culture or even in one person, or robot.

The same is probably true of forms of language, society, morals, religion and science. If there are limits to this variation, they will have to be found by scientific investigations, not introspection or philosophical argument. The limits can hardly be studied before the mechanisms of individual and social development are understood, however. We must not try to fly before we can walk, even if we are philosophers.

However, there are many more mundane kinds of limits of possibility which philosophers characteristically attend to in their attempts to analyse familiar concepts. For instance, it is impossible for someone to be a spinster and married; it is impossible to admire someone for his honesty and breadth of knowledge yet never believe a word he says; it is impossible to be interested in botany yet never wish to look at or learn anything about plants; it is impossible to be intensely angry with someone yet not believe that person has done anything you dislike or disapprove of; it is impossible to drive a car with care and recklessly at the same time (though it is possible carefully to drive over a cliff, to commit suicide). These are not laws' limiting what is possible in the world. Rather, they express defining conditions, or logical consequences of defining conditions, for the use of our concepts. Kant called such propositions 'analytic'.

Making such 'definitional' necessities and impossibilities explicit is part of the task of analysing how our concepts work. This in turn is a useful means of drawing attention to the presuppositions we all make about what sorts of things are possible in the world, and about useful ways of sub-dividing these possibilities. Looking at such subtle differences as the difference between 'with care' and 'carefully' (which are different since they have different boundaries) we learn to articulate our implicit common-sense knowledge about possible configurations of human beliefs, motives, decisions and actions. This is a contribution of philosophy to psychology and AI. (See chapter 4 for more on this.)

The role of necessities and impossibilities in philosophy is a large topic, and I have by no means exhausted it. All I wanted to show here is that the scientific aim of discovering limits of what is possible in the world is not an aim philosophers can or should share unless they are prepared to go beyond philosophical argument.

However, it is important for philosophers to expose present limits of our conceptual and representational apparatus often as a first step towards overcoming those limits. I am trying to expose, and remove, limits of our normal ways of thinking about philosophy and science.

3.12. Philosophy and technology

A theory which explains old possibilities may have surprising new implications. Technology includes the use of ingenuity to invent previously unthought of possibilities which can be explained by available theories. But this is also a major part of pure science, as when the kinetic theory of heat explained the possibility of a lowest temperature and the theory of relativity was used to demonstrate and explain the previously unsuspected possibility of conversion of mass into energy, and of the bending of light by gravitation. The realisation of such new kinds of possibilities in suitable experimental situations can provide dramatic new support for the theories which explain them. So can new ways of realising old possibilities. Philosophical theorising can also lead to the invention of possibilities previously unthought of and possible new means of realising previously thought of situations. So philosophy, like science, has its technological application.

For instance, philosophers have tried to use theories of language to show the possibility of logical languages which in one respect or another (e.g. precision, clarity, economy of rules) improve on natural language, or social theories to demonstrate the possibility of improving on existing social structures, or epistemological theories to demonstrate the possibility of improving on prevailing standards of rigour in science or mathematics. Similarly there is a technological theme to this book, insofar as it uses a theory of the relation between philosophy and science in an attempt to show the possibility of new types of collaboration between philosophers and scientists who study man, or engineers who try to design intelligent machines.

3.12. Laws in philosophy and the human sciences

I have tried to show that philosophy and science have overlapping interests, and partially similar methodologies, and that philosophy can generate science. The affinities between science and philosophy seem to be strongest in the case of the sciences which study man. For it is unlikely that these sciences will, in the foreseeable future, go beyond theories which describe and explain possibilities (the things people and social systems can do).

It seems very unlikely that they will discover new laws with predictive content and explanations of those laws, apart from such trivial laws as are based on common sense, such as the law 'that no normal person in our culture calmly invites a total stranger to chop his leg off! Some alleged laws are very likely to be culture-bound regularities, modifiable by training, propaganda, or economic pressures. Other apparent laws 'discovered' by empirical research are in fact just disguised tautologies, true by definition, for instance: 'Other things being equal, people tend to choose alternatives which they believe will bring about what they desire most'; or 'Persons are more likely to believe a statement if it is made by someone they respect, other things being equal'.

But the lack of substantial laws does not leave the human sciences without content, for there are many kinds of social and psychological phenomena whose *possibility* is well known and needs to be explained, even though the prediction and explanation of *particular instances* is out of the question, since it depends enormously on highly complex individual histories, decision-strategies, beliefs, interests, hopes, fears, ways of looking at things, and so on.

To turn to the search for *probabilistic* or *statistical* laws, when the hope of *universal* laws has been abandoned, as so often happens, is to reject the opportunity to study and interpret the rich structure of particular cases as a way of finding out what possibilities they exemplify.

Insofar as there are laws and regularities to be discerned among all the idiosyncracies of human behaviour, they can hardly be understood and explained before the possibilities they limit have been described and explained. Outside novels, there are so far few, if any, rich and systematic descriptions or explanations of human possibilities, so the human sciences will need to join forces with philosophy in the study of possibilities for some time yet.

3.14. The contribution of Artificial Intelligence

But not only with philosophy, for in the new discipline of artificial intelligence theories are emerging, in the form of specifications for computer programs, which, for the first time, begin to approach the complexity and generative power needed for the description and explanation of intelligent behaviour while also accounting for immense individual differences (as pointed out by Clowes, in 'Man the creative machine').

When such theories are embedded in computers and shown by the behaviour of the computer actually to work, then this establishes that they do not rest on presuppositions of the type they are trying to explain. (However, at present, A.I. models explain only a very tiny fragment of what needs to be explained.)

It may turn out that the combination of skills and knowledge required to construct non-circular and rigorous explanations of any significant range of human possibilities cannot exist in any one scientist nor in any team of scientists, philosophers, and engineers, small enough to co-operate. Human possibilities may be too complex to be understood and explained by humans. But the time is not yet ripe for drawing this pessimistic conclusion, and even if it is true, that is no reason for not trying.

3.15. Conclusion

The best way to make substantial new progress with old philosophical problems about mind and body, about perception, knowledge, language, logic, mathematics, science and aesthetics, is to reformulate them in the context of an attempt to explain the possibility of a mind. The best way to do this is to attempt to *design* a working mind, i.e. a mechanism which can perceive, think, remember, learn, solve problems, interpret symbols or representations, use language, act on the basis of multiple motives, and so on.

Computers cannot yet do these things in a way which compares with humans, and perhaps they never will. But computer programs provide the only currently available language for formulating rigorous and testable theories about such processes. And only with the aid of computers can we explore the power of really complex and intricate theories. (Part two of this book elaborates on the kind of complexity involved.) So I conclude that in order to make real advances in problem areas mentioned above, philosophers, like psychologists and linguists, will need to learn about developments in the design of computing systems, programming languages and artificial intelligence models, even if they do not write programs themselves.

The ('meta-level') concepts *used for describing* computing systems, programming languages, hardware and software architectures, etc. are as important as, or perhaps even more important than the concepts *used in* programming languages.

The attempt to design a mind is a very long term research enterprise. I expect that it will provide the best illustration of the overlap between science and philosophy.

[Book contents page](#)

[Next: Chapter four](#)

Last updated: 4 Jun 2007