

Experimental Philosophy, Edited by Joshua Knobe and Shaun Nichols. New York: Oxford University Press, 2008. Pp. xi + 244

Reviewed by Richard Holton, Department of Linguistics and Philosophy, MIT

Like many marriages, that between science and philosophy has had its ups and downs. Some great philosophers have been great scientists too: Aristotle, Descartes and Leibnitz spring to mind. Some, whilst not being scientists themselves, have been sympathetic, and have made use of scientific findings in their philosophical work. Some have been uninterested, thinking that science had little bearing on philosophy. And some have been outright hostile, defining their work in opposition to science.

Of course, such attitudes do not just affect the great. The relationship with science has waxed and waned across the whole philosophical community. The period after the Second World War saw much analytic philosophy isolated from science, as the focus turned to conceptual analysis. The concern was not directly with the world as science investigates it, but with our concepts of the world; not, for instance, with causation, but with what we mean by 'cause'. In this way philosophy could find its own niche, safe from anything that science might tell it. Concepts could be explored by reflecting on one's own intuitive judgments about possible cases, something that can be achieved, as it is often put, from the armchair.

Over the last thirty years or so, the pendulum has swung. Conceptual analysis has shown its limits. Reflection from the armchair, with its connotations of pipe tobacco and the SCR, is in retreat. What is demanded by your contemporary philosopher is a desk, and a computer, and a connection to the web. Few people now do work in philosophy of mind without knowing the relevant psychology or neuroscience; or do philosophy of language without knowing the linguistics; or do philosophy of physics without knowing the physics. Scientifically informed philosophy is, once again, the norm.

But having scientifically informed philosophers is one thing; having philosophers actually do science is quite another. Now we have a new movement, Experimental Philosophy, whose adherents do just that. Like any self-respecting modern movement it has its own conferences, its own blog (experimentalphilosophy.typepad.com), its own abbreviation ('x-phi'), and several articles about it in the mainstream press. It has been making quite a splash.

This volume provides an admirable introduction to experimental philosophy. Edited by two of its foremost practitioners, it collects seven influential articles that exemplify the approach—two on knowledge and semantics, three on moral responsibility and two on when actions are intentional—and adds four more pieces that discuss the methodology.

So why do these philosophers think that they should get into the business of doing science, rather than leaving it to those who are better trained and better equipped? Obviously there is nothing special about having philosophers perform experiments. The

point is rather that they want answers to empirical questions that scientists are not interested in asking. The advocates of experimental philosophy have asked the same sorts of questions that philosophers have addressed to themselves in the armchair—Can I know something if I have arrived at my belief from luck? Can I be responsible if my actions are determined?—but have asked them of ordinary people. So what was once answered by a priori reflection is now answered by experiment: we discover, for instance, what percentage of ordinary subjects will judge that agents whose actions are determined are nonetheless responsible for them. The science in which empirical philosophers are thus engaged is psychology, and the standard method is the questionnaire.

In a curious way then, empirical philosophy represents something of a return to old-style conceptual analysis. The focus is still on the nature of our concepts, as this is revealed by the judgments that we are prepared to make about possible cases; it is just that the ‘we’ extends beyond the circle of professional philosophers. As a result the approach has drawn much of the same criticism that has been directed at conceptual analysis. The critics contend that philosophy should be concerned with the nature of the world, not with the nature of our concepts of it. Just as one is not going to find out much about physics or mathematics by asking ordinary people for their intuitive judgments about quarks or complex numbers, so one is not going to find out much about knowledge, or the nature of moral responsibility, by asking what ordinary people think about that.

A simple response to this is to insist that insofar as philosophers make claims about our ordinary concepts or our ordinary judgments, it is good that those claims are well substantiated: we need to be sure that these are the judgments of ordinary people and not just of philosophers whose views are likely to be skewed by their own theories. That is fair enough, but provides nothing to answer the critics who think such claims are without interest.

A more substantial rationale is a deflationary one outlined by the editors in the introduction to this volume. In many areas we have come to realize that we are prone to systematic errors and biases: this is so in our logical reasoning, for instance, and in our assessments of probabilities. Once we realize what they are, we can understand why our intuitive judgments are often contradictory, and can often discount some of them as the product of these errors. One possibility that experimental philosophy can investigate is that we are prone to similar errors in some of our philosophical judgments. Certainly the papers in this volume show that judgments are extremely sensitive to factors that one might not expect to be relevant: judgements of what constitutes knowledge differ between Western and Asian subjects; judgments of whether agents are morally responsible in deterministic worlds are more likely to be positive if the case is described concretely than if it is described abstractly; actions are more likely to be judged intentional when the consequences are bad than when they are good. None of this shows that there is error in the judgments; but it does suggest that there are confounding factors.

This deflationary role is clearly philosophically important; and diagnosing sources of moral error can have important practical implications. However, I suspect that at a purely philosophical level there is potentially an even more fundamental role that experimental philosophy can play. A central philosophical debate, going back at least to

Plato, concerns the objective status of our thought. Science serves as a paradigm for something objective. But what of our moral thinking: do we read ethics off the world, or project it onto it? Such questions are notoriously hard to make tractable, but one promising line has drawn the distinction between those areas where our concepts are dictated by the world, and those where they reflect our own responses. A useful example here is colour. A line of thought from Locke onwards claims that all that holds the physically diverse class of red things together is that we see them as red. This is not to deny that they are really red; it is rather to say that the concept of red is one that derives from our responses, and not from the nature of the world as it exists independently of us.

Could ethics be the same? A number of thinkers have argued, on a priori grounds, that it is, but it is unclear how a priori considerations alone could ever establish such a conclusion. Perhaps what we need is some empirical input. After all, it was a set of empirical findings about colour—that it did not figure among the central explanatory scientific categories; that objects under a microscope lack colour—that influenced Locke; and it has been a set of further empirical findings—that the physical bases of colour are hugely diverse, and that our judgments of colour are so context sensitive—that have made an account like his so plausible. Similarly, it is only when we come to know quite what our moral responses are that we will be able to start to assess whether a parallel account of morality has any real plausibility. We will need to know exactly when we judge people responsible, and what the effects of factors like acquaintance and distance are on these judgments. We will need to know, not just how we judge, but how we actually act. And we will need to know how we respond in ways other than judgments nor actions: what our emotional responses are, for instance, and whether any other responses are elicited by typical moral situations. It is only when we have such data that we will be in a position to assess how best to formulate an account of ethics that treats human responses as central, let alone to know whether it is true. And what is true of ethics is equally true of other domains that might be treated in similar ways: certain epistemic notions, for instance, or semantic.

Much of this is work that needs to be done in the lab. A lot of it though is just the kind of work that is represented in this collection. It will not solve philosophical problems on its own; none of the authors featured here claims that it will. But it could tell us a great deal about what a philosophical solution will be like. We have much to learn.