

ISSUE 92 SEPTEMBER/OCTOBER 2012

UK £3.50 US & CANADA \$7.49

Philosophy Now

a magazine of ideas

Raymond Tallis
When Did Time Start?

Helen McCabe
John Rawls & Justice

Christopher Norris
Philosophy Inc.

Mark Phelan
Experimental Philosophy



Philosophers
on
Philosophy

ISSN 0961-5970

09>



9 770961 597048

Experimental Philosophy As An Elephant

Mark Phelan dismisses common misconceptions of Experimental Philosophy

Experimental philosophy is a recent movement that attempts to use experimental methods to cast light on philosophical problems. The difficulty in saying anything more detailed about the movement is that it is large and still developing. But the failure to offer a more complete specification has often left the movement's observers to characterize it as something less than it is. If experimental philosophy is a growing elephant, then observers of it have resembled the proverbial blind men, feeling aspects of the elephant, then describing what they've felt in a variety of incommensurate ways – as a rope when feeling the tail, a tree trunk when feeling a leg, etc – each false idea formed from an incomplete perception. Here I'll try to offer more insight into the nature of experimental philosophy by introducing four common characterizations of the movement, and describing ways in which practitioners of experimental philosophy have superseded each characterization.

1: "The Survey Says... Incompatibilism!"

According to one early characterization, the experimental philosophical method involves polling people's reactions to philosophical thought experiments, and adjudicating philosophical debates on the basis of the results, in the same way in which we pick senators or members of parliament: majority rules. This first-past-the-post approach to solving philosophical problems is not what experimental philosophy is about, nor has it been since its earliest instances.

The experimental philosophy movement was born mostly in New Jersey, with a couple of papers published in the early aughts. In 2001, Jonathan Weinberg, Shaun Nichols, and Stephen Stich (let's call them WNS) published 'Normativity and Epistemic Intuitions' in *Philosophical Topics* (2001), which they followed two years later with 'Metaskepticism: Meditations in Ethno-Epistemology'. Although these early papers do assess people's reactions to thought experiments, they do not engage in the crude method of polling as a means of adjudicating philosophical debate.

In both papers, WNS explore how different groups of people ascribe knowledge to a protagonist in a variety of philosophical thought experiments. For example, WNS ask people about the following story:

Bob has a friend, Jill, who has driven a Buick for many years. Bob therefore thinks that Jill drives an American car. He is not aware, however, that her Buick has recently been stolen; and he is also not aware that Jill has replaced it with a Pontiac, which is a different kind of American car. Does Bob really know that Jill drives an American car, or does he only believe it?

This is a version of what has come to be known as a 'Gettier case' – a case in which someone is justified in believing something that is true, but is nonetheless thought (by at least a large number of philosophers) to lack knowledge. Edmund Gettier's pithy 1963 paper described several such cases and touched off decades of debate, and was taken by many to overturn the theory of 'knowledge as justified true belief' that had held sway since Plato first formulated it. However, WNS found that around 60% of people from East Asia and the Indian Subcontinent think – unlike most professional philosophers in the West – that Bob does not merely believe but *really knows* that Jill drives an American car in the above case. On the other hand, three-quarters of Westerners share with the (Western) philosophers the intuition that Bob *does not know* but *only believes* that Jill drives an American car. WNS's findings suggest that Western intuitions about knowledge differ from Asian/Subcontinental intuitions about knowledge *vis-à-vis* Gettier cases – which raises the possibility that the philosophers' reaction to Gettier cases was nothing more than an adventure in ethno-epistemology. What's more, WNS found significant differences in the way people of diverse backgrounds reacted to a variety of thought experiments, thus further challenging the purported universality of analytic philosophy's methodology and findings. These early experimentalists seem to have embraced something like the philosophically-standard role for intuition in confirming truths, but broadened the class of those whose intuitions count. Instead of counting only the intuitive opinions of a few hundred genteel east-coast or British academic intellectuals, experimentalists adopting this perspective attempted a wider survey of philosophical intuitions.

Now, if WNS concluded from their results that, contrary to Gettier's indications, knowledge *is* justified true belief because there are fewer Westerners than Asians; or if they decided that there are two distinct concepts of knowledge, they would be



Buick and Pontiac: American cars. But which one? (see above)

acting consistently with the first conception of experimental philosophy. Their argumentative strategy would conform to the polled-solution-to-the-debate misconception. But this is not what they conclude. Instead they argue from their results to a challenge for analytic epistemology [epistemology is the study of the justification of knowledge or belief]: epistemologists either have to accept WNS's findings, and show how these are consonant with epistemology seeking intercultural truth via intuitive reactions to thought experiments; or accept the findings, and reject the reliance on intuition to maintain the intercultural reach of epistemology; or accept the findings, and reject the intercultural reach of epistemology; or explain away the results in a way that is consonant with the reliance on intuition and the universal reach of epistemology. Posing this challenge has been a fruitful move in the debate. Numerous responses in terms of these options have subsequently been articulated by professional epistemologists. However, since WNS ask experimental participants about knowledge and draw conclusions about *philosophical practice*, not about knowledge *per se*, it should be clear that experimental philosophers are not simply recommending philosophical theses on the basis of the majority reaction to experimental prompts, nor have they been doing so since the beginning.

2. "It's the Philosophy of Philosophy"

The early examples of experimental philosophy I've just discussed do not address the stereotypical questions of philosophy, such as: What is the nature of knowledge? Is determinism compatible with freedom and responsibility? Is what is right simply what is best? Instead they suggest a reevaluation of philosophical practice and a challenge to traditional conceptions of evidence in philosophical debates. So, if one read only the early papers, one might come to characterize experimental philosophy as a variety of *metaphilosophy* [a philosophical consideration of philosophy], as some critics have.

Here as elsewhere the problem lies in taking

hold of but a single part of the elephant. To be sure, a number of experimental philosophical projects are metaphilosophical in nature, but other projects are not. Consider another early paper, Joshua Knobe's 'Intentional Action and Side Effects in Ordinary Language' (*Analysis*, 2003). Here Knobe investigates what ordinary people would say about certain side-effect cases, in which someone foresees a certain result of his action, but doesn't care whether or not that result occurs. In each of Knobe's experiments, participants are asked about one of two cases. In one case, the side-effect is good, and in the other, bad.

In one experiment, participants read one or another story about an army lieutenant. In one story, the lieutenant's sergeant informs him that if the assault on Thompson Hill proceeds as planned, some men will be taken out of the line of fire and saved. The lieutenant responds, "Look, I know that we'll be taking them out of the line of fire, and I know that some of them would have been killed otherwise, but I don't care about what happens to our soldiers. All I care about is taking control of Thompson Hill." Here most experiment participants (and probably you too) say that the lieutenant did not intentionally save the men's lives by proceeding with the attack. Other participants read a subtly different story, to which they have a very different reaction. In this story, the sergeant informs the lieutenant that if the assault on Thompson Hill proceeds as planned, some men will be *put into* the line of fire, and killed. The lieutenant responds, "Look, I know that they'll be in the line of fire, and I know that some of them will be killed, but I don't care at all about what happens to our soldiers. All I care about is taking control of Thompson Hill." Here most participants said that by proceeding with the attack the lieutenant intentionally killed his men.

Now if Knobe claimed that his results called for a reevaluation of philosophical practice, his work would be about the philosophy of philosophy and conform to the second misconception. But this is not how Knobe interprets his findings. At his most explicit, Knobe identifies the question he is addressing as a question about how we should speak: shall we say that certain side-effects are brought about intentionally? But Knobe also suggests that knowing what people do say about such cases can help us adjudicate a philosophical debate about the true nature of action – specifically, the debate over whether one could bring about a side-effect intentionally. So, it seems most reasonable to understand Knobe as seeing his studies as offering strong evidence concerning what people say and think, and a small piece of evidence, which can be defeated by other considerations, concerning the nature of intentional action. Either way, Knobe is clearly up to something other than a metaphilosophical reevaluation of philosophical practice. And (as bears repeating) he is clearly up to something more sophisticated than a first-past-the-post approach to philosophical problems.

Other early experimental work blends metaphilosophical thoughts with other reflections. In 'The Folk Psychology of Free Will' (*Mind and Language*, 2004), Shaun Nichols argues on the basis of his experimental research that children accept a model of free will according to which "agents have the capacity to cause actions, and for a given action, an agent could have

done otherwise." Then, later, in a more metaphilosophical bent, he explores a psychological explanation for "the historical recalcitrance of the philosophical problem of free will." The paper 'The Phenomenology of Free Will' by Eddy Nahmias, Stephen Morris, Thomas Nadelhoffer, and Jason Turner in *Journal of Consciousness Studies* (2004) has metaphilosophical overtones concerning the role of assumptions in debates about free will. But in it they also conclude that their own tentative (at the time) research into people's experience of free will provides some slight support for a compatibilist conception of free will (according to compatibilism, free will is possible even in a deterministic world operating through natural laws).

Unsurprisingly, as experimental philosophy grew its aims became more complicated. But experimental philosophers never proposed anything so crass as a majority-rules resolution of philosophical debates, and very early on, they included more than a metaphilosophical project.

3. "Intuition, Intuition, Intuition!"

Each of the articles so far cited involve a reassessment of intuitive evidence. The negative project described in the first section urges a wholesale revision of the use of intuition in philosophy on the basis of experimental research by broadening the sample class of those whose intuitions matter. On the other hand, the experimental projects discussed in the second section embrace the standard evidential role of intuition in philosophical practice, but attempt to reassess specific intuitions. Thus, despite their divergent aims, the papers cited so far give the strong impression that experimental philosophers are intuition-obsessed, inviting the retort that experimentalism does not reach beyond one limited source of philosophical evidence. Once again, this characterization is due to a failure to survey the full field of experimental philosophical projects.

Unfortunately, this 'intuitive' characterization of experimental philosophy has been aided and abetted by some prominent early descriptions of the movement by experimentalists themselves. For instance, in 'The Past and Future of Experimental Philosophy' (*Philosophical Explorations*, 2007), Thomas Nadelhoffer and Eddy Nahmias claim that all work in experimental philosophy "shares a commitment to using controlled and systematic experiments to explore people's intuitions and conceptual usage." Self-reflective characterizations of a new movement are necessary (I self-reflectively write); but if the movement is still developing, as experimental philosophy has been, they run the risk of being outpaced by new developments and becoming prominent *mischaracterizations* of it.

It is no longer correct (if ever it was) to say that all experimental philosophy projects involve the empirical study of intuitive assessments of particular cases – the 'thought experiments' (or 'intuition pumps') of traditional philosophy. Just as philosophers have traditionally appealed to a disparate assortment of evidence, so experimentalists have begun to systematically assess not only intuitive judgments, but other sources of evidence as well. For example, a team led by Eric Schwitzgebel has attempted to gain insight into the ancient and oft-asserted claim that philosophical reflection on ethical issues can improve one's behavior, by examining how professional ethicists behave in a

variety of circumstances. At least at face value, their results suggest that moral reasoning has little effect on moral behavior. Schwitzgebel and colleagues found that professional ethicists are no more likely to vote, or respond to student emails, than are non-ethicist philosophers and professors. Audiences in ethics sessions at philosophical conferences are generally just as likely to behave discourteously as audiences in non-ethics sessions. And ethics books (even obscure ones) are *more* likely to be missing from library collections than are books from other philosophical sub-disciplines. If, as Schwitzgebel *et al* argue, this is important evidence regarding the behavioral efficacy of ethical thinking, it is not evidence from intuition.

In other work, Joshua Knobe and Jesse Prinz examined how frequently people use different kinds of mental state ascriptions on the internet, and argued from their findings that people are ordinarily willing to ascribe beliefs and desires, but not experiences or emotions, to disembodied or distributed entities like corporations. No intuitions are invoked in these findings. And in my own experimental work, by systematically comparing how people paraphrase metaphorical and literal utterances, I have found evidence against the claim that it is distinctively difficult to express what a metaphor means.

Each of the approaches I have just described does something other than survey intuitive reactions to thought experiments. In these studies, both thought experiments and intuitive assessments are absent, as in many other recent projects in experimental philosophy. So if one characterizes experimental philosophy as the systematic investigation of intuitions about particular thought experiments, one has got hold of only part of the elephant.

4. "It's Not Philosophy"

According to the final characterization, experimental philosophy is not only not an intuition-obsessed philosophy of philosophy, it is not philosophy at all.

Some who adopt this perspective suppose that experimental philosophers are philosophy-hating philosophers (or amateur social scientists) who lack skills in assessing arguments, drawing careful distinctions, reasoning abstractly and other aspects of traditional philosophical methodology. Thus, the proponents of this first form of the 'not philosophy' characterization suppose that experimental philosophy is not really

衆
瞽
摸
象
之
圖



philosophy because it does not employ philosophical methods.

I hope I've shown by now that this is not the case. Experimentalists largely embrace traditional philosophical methods, and simply seek to augment or revise them. A slightly more sophisticated version of the same mischaracterization would admit that experimentalists do employ traditional philosophical methods, just poorly. But it would be tough to maintain this position after a careful examination of the experimental philosophy published so far, including pieces in top-tier philosophy journals, and works by some primarily non-experimental philosophers, specifically known for their care in assessing arguments and reasoning.

According to another variation on this theme, philosophy is essentially *normative* – it concerns how the world should be – but (this version goes) if experimental philosophy employs the experimental methods of science, then it can only answer *descriptive* questions – questions concerning how the world *is*. Therefore, this form of the 'not philosophy' characterization claims that experimental philosophy cannot contribute to the distinctly philosophical project.

Is this true?

This critique of experimental philosophy has occasionally been couched in terms of the familiar philosophical slogan that you can't derive an 'ought' from an 'is', as Hume argued in his *Treatise of Human Nature* (1740). Many philosophers interpret Hume's principle to be incompatible with naturalism – the view that the truth of normative principles can depend on descriptive facts about the world. But, as Charles Pigden pointed out in an article for *Philosophy Now* in Issue 83, Hume's arguments for his principle do not support this anti-naturalistic interpretation, nor did Hume intend it. Instead, Hume's

arguments support the more modest conclusion that you cannot derive a normative principle from descriptive facts *apart from using analytic bridge principles* – that is, apart from using premises which define the normative in terms of descriptive facts. Indeed, Hume himself was a naturalist, who defined virtue to be "whatever mental action or quality gives to a spectator the pleasing sentiment of approbation," and suggested, as an experimental philosopher might, that to determine what is virtuous "we then proceed to examine... what actions have this influence." (*Enquiry*, p.289). If it is reasonable to define the normative in terms of non-normative facts, then it is reasonable for experimental philosophers to expect to derive normative conclusions using their descriptive tools: so they *would* be doing philosophy. Of course, one might maintain that the normative *cannot* be defined just by describing features of the world (a position that was embraced by G.E. Moore). But even if the descriptive derivation of norms were impossible, the second form of the not-philosophy characterization should still be rejected, since much of philosophy is not essentially normative. Philosophers have clearly long been concerned with numerous questions of straightforwardly descriptive fact, such as: Are the objects of perception external to the mind? Are all concepts derived from percepts, or are some concepts innate? Is all human action self-interested? Thus, even if experimental philosophy could not contribute to the normative questions of philosophy, many descriptive questions would still be promising game. Those who maintain otherwise have failed to get hold of some elephant or other.

© DR MARK PHELAN, 2012

Mark Phelan is Assistant Professor of Philosophy and Cognitive Science at Lawrence University, Appleton, Wisconsin.