Thought Experiments and Experimental Philosophy

PREPRINT

in the *Routledge Companion to Thought Experiments*,
edited by Mike Stuart and James Brown (2016).

Kirk Ludwig
Philosophy Department
Indiana University

1. Introduction

Much of the recent movement organized under the heading "Experimental Philosophy" has been concerned with the empirical study of responses to thought experiments drawn from the literature on philosophical analysis.[1] I consider what bearing these studies have on the traditional projects in which thought experiments have been used in philosophy. This will help to answer the question what the relation is between Experimental Philosophy and philosophy, whether it is an "exciting new style of [philosophical] research", "a new interdisciplinary field that uses methods normally associated with psychology to investigate questions normally associated with philosophy" (Knobe et al. 2012), or whether its relation to philosophy consists, as some have suggested, in no more than the word 'philosophy' appearing in its title, or whether the truth lies somewhere in between these two views.

Section 2 distinguishes different strands in Experimental Philosophy. Section 3 reviews some ways in which Experimental Philosophy has been criticized. Section 4 considers what would have to be true for Experimental Philosophy to have one or another sort of relevance to philosophy, whether the assumptions required are true, how we could know it, and the ideal limits of the usefulness Experimental Philosophy to philosophy. Section 5 is a brief conclusion.

---

[1] Contemporary Experimental Philosophy is roughly 15 years old, with the first papers that practitioners look back to for inspiration appearing in 2001, e.g., (Weinberg, Nichols, and Stich 2001). An important precursor was *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry* (DePaul and Ramsey 1998). At the time of this writing, the experimental philosophy bibliography in Philpapers has 824 entries, across philosophy of action, language, mind, ethics, epistemology, metaphysics, the nature of experimental philosophy, cross-cultural research, and miscellaneous other topics. Not all of these papers, however, fit the paradigm of survey philosophy examined here. There are some earlier antecedents to the use of the survey method in philosophy in the work of Arne Naess (Naess 1953). Two useful collections are (Knobe and Nichols 2008, 2014).

2. Varieties of Experimental Philosophy

Broadly construed, Experimental Philosophy is philosophy informed by empirical work.[2] Experimental Philosophy in this sense stretches back to antiquity. Here we will be concerned with a narrower conception of Experimental Philosophy characterized by the adoption of the survey as a central methodological tool, with a yes-or-no question, or a graduated range of answers from "strongly disagree" to "strongly agree" or the like, typically about a scenario in a philosophical thought experiment. Given this, Experimental Philosophy, in the sense we are interested in it, might more aptly be called "Survey Philosophy."

The most important division among experimental philosophers is between those who conceive of its project negatively and those who conceive of it positively. The negative (x-phi⁻) and the positive projects (x-phi⁺) both share the assumption that a central philosophical activity involves eliciting (what are often called) intuitions about actual and hypothetical cases, the latter involving conducting what we call thought experiments, with the aim of conceptual articulation or analysis. The traditional attempts to provide a satisfactory analysis of the concept of knowledge surrounding the literature on the Gettier cases (Gettier 1963) is a paradigm of the sort of activity they have in mind. In this case, the judgment (or intuition) that a subject with a justified true belief that p that is based on a justified false belief does not *thereby* know that p is taken to show that justified true belief is not sufficient for knowledge.

The negative project seeks to show that the results of surveys of undergraduates or others without much philosophical sophistication cast doubt on the probative value of intuitions and the use of thought experiments by philosophers. The negative project argues that the standard use of thought experiments makes certain empirical assumptions which can easily be tested by the survey method, and that (surprisingly) the assumptions (never tested) turn out to be false (Nichols, Stich, and Weinberg 2003; Alexander and Weinberg 2007; Stacy Swain 2008; Weinberg, Nichols, and Stich 2001; Machery et al. 2004; Liao et al. 2012). Among these alleged assumptions are that philosophers' intuitions are shared by everyone, that they are not biased, based on irrelevant factors, or theory driven, and that they are not relative to cultural or socioeconomic background, and the like.

There are two main ways of taking the positive project. The first is the Continuity Account, and the second the Psychological Account.[3]

---

[2] An example is the use of experimental work from social psychology to argue that virtue ethics makes false empirical assumptions about character traits (Doris 1998; Doris 2002; Harman 1999, 2000). This is still a matter of controversy (Kamtekar 2004; Snow 2010). (Rose and Danks 2013) make a plea not to use the label just for the narrower conception.

[3] And there can be mixtures as well, and sometimes it is a little difficult to tell, and the lack of clarity is often aided by a failure to distinguish between different senses of the world 'concept' or confusion about what sense of 'concept' is at issue in philosophical analysis.

1. The Continuity Accounts. X-phi is an enterprise with the same goals that philosophers have when they use thought experiments (the investigation of the application conditions of words or concepts, or of entailment relations between propositions, or implications of sentences), except that we crowd source the answers (Nahmias et al. 2006; Knobe 2003a, 2003b, 2004, 2006; Malle and Knobe 2001; Pettit and Knobe 2009; Buckwalter and Schaffer 2013; Genone and Lombrozo 2012). There are two further subcategories.
   a. Replacement Accounts: We add that survey philosophy should replace traditional "armchair" methods. (Alexander and Weinberg 2007)
   b. Supplement Accounts: We urge only that survey philosophy can provide a useful supplement to traditional methods. (Papineau 2011; Talbot 2013)
2. Psychological Accounts. X-phi is not the pursuit of traditional philosophical goals (or at least analysis) by appeal to surveys of the philosophically unsophisticated but instead a psychological inquiry,[4] as opposed to a philosophical inquiry, into
   a. concepts that philosophers have been interested in (Nichols 2011; Sarkissian et al. 2010; Knobe and Burra 2006), or,
   b. into psychological mechanisms generating responses, whatever they might reveal, that people have (in various groupings) about philosophical thought experiments (Young et al. 2006; Nichols and Knobe 2007; Sarkissian et al. 2011; Knobe 2007, sec. 2).[5]

I will consider each of these programs within x-phi, but focus most attention on the Continuity Account as the most promising case of the *positive* relevance of x-phi to the traditional projects of philosophy.

3. Criticisms of Experimental Philosophy

Experimental Philosophy has been criticized on a number of different grounds (not all these will be completely independent of, or consistent with, one another).

1. Negative x-phi has been directed against positive x-phi, on the assumption that it is pursuing what philosophers have traditionally been trying to do (Alexander, Mallon, and Weinberg 2010). If intuitions are not probative, then they are not probative period, whether you are gathering them in the armchair or in Central Park.

---

[4] There has been some dispute over the use of the term 'philosophy'. For example, Knobe has urged that the question how people think when they respond to surveys (asked as an empirical question about their psychology and not in the mode of conceptual analysis) *is* a philosophical question. To adopt this terminology, though, would obscure the issue whether these psychological studies bear on the projects for which thought experiments have been deployed in philosophy in the tradition in the last century.
[5] Some experimental philosophers gloss this as answering the question what concepts ordinary people operate with. In these cases, it seems likely that they have in mind by 'concept' something like 'psychological-causal structure that guides behavior and speech'. It would be a mistake to think this is the same project that the tradition in philosophy was interested in when asking for an analysis of, for example, the concept of justice.

2. Critics have argued that x-phi has been taken in by an illusion widespread among philosophers themselves, namely, that contemporary analytic philosophers rely on intuitions as a source of evidence for philosophical theories. Herman Cappelen in his (2012) book *Philosophy without Intuitions* argues that this is simply an illusion, and experimental philosophers have been taken in by it. While philosophy goes the way it always has, experimental philosophers are engaged both in their critical and in their constructive projects in pursuing a will o' wisp. X-phi⁻ is tilting at windmills while x-phi⁺ is simulating in surveys a practice that philosophers have never actually been engaged in. (See also (Cappelen 2014); and see Cohnitz and Häggqvist in this volume for discussion, section III.)

3. An allied criticism of x-phi⁺ is that concepts philosophers are interested in are not amenable to conceptual analysis because they are of natural kinds and, hence, their real essence is to be discovered by empirical investigation of the world, not in empirical investigation of speakers' dispositions to classify things under the concept. (Cf. Hilary Kornblith's view that knowledge is a natural kind and therefore not amenable to investigation by conceptual analysis (2002).)

4. Another criticism of the continuity version of x-phi⁺ is that there are no (at least interesting) conceptual truths, since there are no (at least interesting) analytic truths, and, hence, the x-phi⁺ conceived of as pursuing conceptual analysis is aiming for a non-existent or uninteresting target (Quine 1953; Putnam 1965), and, in any case, in fact philosophy has all along been aiming at general synthetic a posteriori truths (Papineau 2013).

5. Yet another criticism that focuses on the relation of x-phi to philosophy is that it assumes both in its negative and its positive versions an overly simplified and narrow role for thought experiments in philosophical theorizing, focusing on an overly simplified conception of the "case method," whereas in fact thought experiments are used for illustration, to draw analogies, to raise puzzles (as in Thomson's contrasting the switch and fat man trolley cases (Thomson 1976)), to draw attention to the range of cases that a theory must deal with, to draw out the consequences of theories, to illustrate arguments and bring out the limitations of our use of language (e.g. thought experiments involving sorites series), and to draw out how we think about certain matters, assumptions we make or principles we reply on, without the suggestion that we are drawing attention to conceptual truths (Sosa 2007a, pp. 101-102). (See also Cohnitz and Häggqvist, this volume, section II, in this connection).

6. Even in the so-called case method, x-phi operates with a caricature of philosophical method. Philosophers aim to arrive at a reflective judgment about a case and then to review it in the light of other judgments (their own and others) and more general theoretical considerations. They do not simply record their spontaneous judgments and take the third person stance toward them as neutral observations to be explained. They do not present themselves with scenarios out of the blue and like a medium at a séance wait for the spirit to move them to say something. It is an intellectual exercise like figuring out how to construct a proof of something in logic, or figuring out a mathematical problem, turning something on all sides to get the right view of it, reviewing a range of cases, testing for things that might be misleading by "turning the knobs" as Douglas Hofstader puts it, and looking out for familiar pitfalls. We often enough (though not always) have a sense of not being clear, and we withhold judgment

until it becomes clearer (it is not a forced choice).   Where it seems relevant, many of us think it is important to be familiar with contiguous domains of scientific investigation, which, while not immune from conceptual confusion, often present us with important puzzle cases and problems and insights. And we do not do this like hermits in the woods: we try out ideas and thought experiments on others, give and publish papers, take criticism, make revisions, try out new ideas generated in this process, and so on. See (Ludwig 2007; Jackson 2011, sec. 5; Bealer 1998).

7. Experimental philosophy has been criticized more narrowly on methodological grounds (these are all intertwined—see (Ludwig 2010, 2007; Deutsch 2009; Kauppinen 2014, pp. 5-6; Williamson 2011; Sommers 2010) for general discussion; see (Cullen 2010) for empirical refutations of assumptions behind some celebrated survey results).

   a. Poor Design. In some surveys which have gotten wide attention, the scenarios or questions have been unclear, misleading, or ambiguous, and insufficiently informed by the relevant knowledge of the issues, so that the interpretation of the results depends upon further untested assumptions about how respondents understood the scenarios and questions.

   b. Misunderstanding Intuitions. Experimental philosophers (many) have simply misunderstood the sort of intuition (or judgment) sought in philosophical thought experiments. They are not "spontaneous judgments … for which the person making the judgment may be able to offer no plausible justification" (Nichols, Stich, and Weinberg 2003, p. 19), nor are they expressions of *what we would say* or *how things seems to us*. They are rather, for example, judgments that draw solely on the concepts contained in the question in relation to the description of the scenario (there are a variety of views about the form and basis of the judgment, but they *all* reject the spontaneous judgment account). (See (Booth 2014; Ludwig 2010; Goldman 2010, 2007; Sosa 2007b, 2007a, 2008) for further discussion.)

   c. Confusion about what data surveys supply. Surveys do not ipso facto elicit intuitions (6b encourages the conflation). Not every response to a question on a survey about a scenario expresses an intuition. The survey data then is not straightforwardly data about intuitions. To use survey's for traditional philosophical purposes, we to filter responses that are not intuitions, or show the noise level isn't so high it makes the data unusable.

   d. Controlling for Factors Relevant to Eliciting Relevant Responses. Surveys typically do not control for a variety of factors that are relevant to getting useful results, such as

      i. proper understanding of the point of the survey, namely, that it is to elicit responses that on the basis of the content of the scenario itself and the content of the concepts involved in the question about it;

      ii. the motivation of respondents in responding;

      iii. their level of effort;

      iv. the general intellectual capacities of respondents, including their capacity to make and keep track of relevant distinctions;

      v. responses to pragmatic implicatures;

      vi. extraneous factors in experimental design that introduce confounds like implicit biases and emotional responses that skew judgments.

e. Impracticality. Doing philosophy by the survey method is unwieldy, even if it can be done. Philosophers can run through a large number of scenarios relevant to assessing various aspects of an account in a short time. Doing the same thing using the survey method would take a lot more time, with nothing more to show for it in the end. Doing philosophy by surveys, while it would generate more journal articles, would slow progress to a glacial pace.

f. Failure to take into account the relevance of competence. Taking the responses of unsophisticated people to thought experiments to be on a par with those of philosophers rests on the false assumption that professional philosophers are no better than their undergraduates in sorting out subtle conceptual issues. Taking the responses of unsophisticated undergraduates and laypeople as being on a par with the responses of philosophers (even controlling for other factors) is to fail to recognize that people differ in how good they are at it and that one can develop a competence in conceptual analysis (inter alia, making distinctions, getting clear on the issues, understanding the point, framing questions and scenarios in the right way, being sensitive to things that might mislead, being ready with alternative cases to check for confusions, etc.) in the same way that one can develop a competence in mathematics or color matching and so on, and that training in philosophy develops such a competence.[6]

g. Over quick generalization from studies. Some of the most celebrated and provocative early studies of the x-phi movement (Nichols, Stich, and Weinberg 2003; Stacy Swain 2008; Weinberg, Nichols, and Stich 2001; Machery et al. 2004), as well as more recent studies (Buckwalter and Stich 2014) alleging a divergence between "intuitions" of women and men, have failed replication tests: on epistemic intuitions see (Nagel, San Juan, and Mar 2013; Adleberg, Thompson, and Nahmias; Kim and Yuan; Seyedsayamdost 2015; Nagel 2012); on theories of reference for names see (Lam 2010); on gender differences see (Seyedsayamdost 2014a; Wright 2010; Adleberg, Thompson, and Nahmias). See (Seyedsayamdost 2014b) for a more wide ranging critique. In this, x-phi shares in the larger replication crisis of psychology (see *Social Psychology*, 45(3), May 2014 and (Collaboration 2015)).[7]

---

[6] It is hard to deny that competence is relevant, but x-phi⁻ researchers have aimed to turn this into a challenge to philosophy by arguing we have no reason to think philosophers are more competent in setting up and responding to thought experiments than the completely untutored—alone apparently among all the academic disciplines, philosophical training induces no epistemic benefits. Or at least: we need a positive reason to think so before we can trust philosophers' judgments. All the more reason to do x-phi! Now: let us try the same form of argument on the mathematicians. In connection with this, see: (Ludwig 2007; Weinberg 2007; Weinberg et al. 2010; Grundmann 2010; Williamson 2011; Sosa 2011; Hitchcock 2012; Rini 2013; Sorensen 2014; Pritchard 2014; Cruz 2014; Hales 2012; Buckwalter 2014).

[7] Proponents of negative x-phi often repeat the conclusions of the early studies without concern for criticisms that point out methodological flaws in the studies, without concern for evidence that they are not replicable, and without concern for whether the conclusions are even coherent; see for example Stich and Tobias in this volume, section 3.1, who state flatly: "To sum up, philosophical intuitions have been found to vary with culture, academic discipline, gender, age and personality." But responses to surveys are not ipso facto philosophical intuitions, and philosophical intuitions, taken as the target response to a thought experiment in the method of cases, cannot vary with culture or anything else (given that the target is the expression of a

8. Failure to yield the right kind of knowledge of results (Ludwig 2013). In philosophy, like mathematics, we are interested in understanding whether something is so from the first person standpoint. It is one thing to be told that the Pythagorean theorem is true. It is another to see that (and why) it is true. If you survey high school math teachers, you will find that most think that the Pythagorean theorem is true. You may be justified in accepting that if they all say that it is true, then it is. But you don't thereby see that (and why) it is true. And if you are interested in mathematics, it is not enough to know that it is true, you want to see that it is true. You want to know on the basis of the exercise of your own reasoning abilities that it is true. And that is the same kind of knowledge that we seek in philosophy. The results of surveys cannot supply it.

I will come back to the last of these below in assessing the question of the relevance of x-phi surveys involving thought experiments and their use in philosophy.

4. What Are the Ways Experimental Philosophy Might be Relevant to Philosophy

What, put abstractly, are the ways in which x-phi *might* be relevant to philosophy?

1. It might give us reason to think that there is something deeply problematic about the use of thought experiments in philosophy (for some purposes at least).
2. It might give us reason to think that some response to a thought experiment is the correct one.
3. It might give us insight into ways in which people can fall into error when responding to thought experiments.

We take these up in turn.

4.1 Can Experimental Philosophy show that there is something deeply problematic about the use of thought experiments in philosophy?

The negative project in x-phi aims to show that intuitions are not probative. This can take either a moderate or aggressive form. The moderate form merely says that in some cases the results of x-phi show that we should be more cautious about the consensus philosophers have reached about some thought experiments (e.g., whether in Gettier cases subjects lack knowledge), and should be less sanguine perhaps generally than they (allegedly) have been. The aggressive form says that the results of surveys show that reactions to thought experiments never yield any knowledge.

---

competence to judge entailments), and in fact the data about even responses does not appear to be replicable in many cases. I do not advocate throwing out all the studies, which as I will say below, have their place in our thinking about philosophy, but I do advocate more methodological caution and greater epistemic humility than is often displayed by proponents of x-phi who see their surveys as relevant to philosophy.

It is difficult to take the aggressive form seriously. Thought experiments rely on our ability to tell when one proposition entails another. We describe scenarios as if about particular individuals, but in fact their content is general. In a Gettier case, we are asked whether if there is someone with a justified true belief that *p* in a circumstance in which he infers it from a justified false belief, he knows that *p* or doesn't know that p. We are supposed to answer on the basis of whether instances of the antecedent entail corresponding instances of the consequent.[8] A general skepticism about our ability to respond correctly to questions about thought experiments calls into question our ability to tell when one proposition entails another or does not. If we can't tell that with a reasonable degree of reliability in appropriate conditions and with adequate training, then all inquiry collapses. There is nothing special about the abilities that we call upon in philosophy. We call upon them in everyday life and in all areas of inquiry.[9]

The moderate form cannot be dismissed out of hand, but even there it is difficult to establish a conclusion of the form: in such and such an area or on such and such a question, we are simply not in a position to come to know whether a test proposition is true given the description of the scenario. We might be provided, however, with evidence for the presence of confounding factors, which can alert us to the need to guard against them.[10] We return to this below.

The mere fact (*if* it is a fact) of diversity of response to thought experiments, within or across cultural, ethnic, or socio-economic groups (Machery et al. 2004; Nichols, Stich, and Weinberg 2003; Weinberg, Nichols, and Stich 2001), does not show that the application conditions of concepts like the concept of knowledge are relative to different cultures or social milieus.[11] As many people have observed, if the participants are deploying the same concepts (and otherwise they just understand the words differently—and this may explain some diversity of responses (Sosa 2007a, p. 104)), all that could be shown by these results is that that not everyone is getting it right, and that there can be factors that correlate with different cultural, ethnic, or socio-economic backgrounds that can contribute to errors (Sosa 2011, p. 466; 2007a, p. 107; Ludwig 2007, 2010). It doesn't show that if there is a consensus among philosophers about what the right answers are, they are wrong, or that they don't know that they are right. Similarly, that you might get variations in responses to simple mathematics questions across different cultural, ethnic, and socio-economic groups

---

[8] On this way of understanding the question, we are not asking whether a necessitated conditional is true, or whether a counterfactual judgment is true. If the proposition that *p* entails the proposition that *q*, then necessarily, if *p*, then *q*, but not vice versa. Similarly for the counterfactual judgment that if *p* were the case *q* would be the case. Entailment, as understood here, is an internal relation between propositions. But if it is possible that *p* and not *q*, then *p* does not entail *q*, and so that suffices to answer the question negatively.

[9] See (Ludwig 2007; Jackson 2011, sec. 2); see (Talbot 2014) for discussion of the sizable empirical burden of showing that standard procedures in philosophy are not adequate.

[10] With respect to the worry that the presence of distracting factors or order effects in some cases shows a general problem because we can't tell when something like this is going on, see (Wright 2010).

[11] As noted above, the studies suggesting systematic variation in responses to thought experiments across cultures and socioeconomic groups have not fared well in replication tests.

does not entail that the consensus among mathematicians about the right answers are wrong or that they don't know that they are right.

4.2 The Positive Program of Contributing to Conceptual Analysis

What has to be true for x-phi (in the survey mode) to have a positive contribution to make to philosophy?

Responses to surveys are never uniform. Consider a yes-no survey on whether a concept applies in a scenario (for example, whether someone with a Gettierized belief has knowledge or not). There are three possible outcomes: a majority say 'yes', a majority say 'no', or the respondents are split evenly between 'yes' and 'no'. The positive program (on the Continuity Account) aims to use the results to say something about the application conditions of the concept of knowledge. If answers are split evenly, we can draw no conclusions. If a majority favors one answer, it is natural to suggest the majority is right (parallel considerations to those brought up below would apply if we held that the minority is right).

What has to be true for this to be so? An assumption that would justify it is that the respondents have a positive bias toward a correct response. Given this, and that their responses are independent, we can show that the larger the group surveyed, the higher the probability that the majority is right. In fact, as long as none is invariably right, we can show that after a certain number of voters (as we can call them) the majority always outperforms anyone in the group. More precisely, according to the Condorcet Jury Theorem, where agents' choices are independent of one another, they all have the same bias p ($1 < p < 0$), the decision rule is simple majority, and there are two alternatives one of which is correct, where P_N is the probability for N agents' of the majority being right, and prior odds as to which of the two alternatives is correct are even (Grofman, Owen, and Feld 1983, p. 264):

> **If $1 > p > \frac{1}{2}$, then P_N is monotonically increasing in N and $\lim_{n \to \infty} P_N \to 1$**; if $0 < p < \frac{1}{2}$, then P_N is monotonically decreasing in N and $\lim_{n \to \infty} P_N \to 0$; while if p – $\frac{1}{2}$ then $P_N = \frac{1}{2}$ for all N.

Even more impressively, the rate of convergence can be quite rapid. If p = .8 then $P_{13} > 0.99$. The following figure shows, in the top graph, the rise in the probability that the majority is right with the probability each individual is right for a group with 95 members.[12] The bottom graph shows the increasing probability that the majority is right, given that the probability that each individual is right of .6, with increasing numbers of members. Both of these show only data points for odd numbered groups—ties for even numbered groups reduce the reliability of the majority though it still tends to 1 in the long run.

---

[12] These graphs were generated in Mathematica using source code due to Tetsuya Saito.

number of respondents $N = 95$.

probability that an individual makes the correct choice

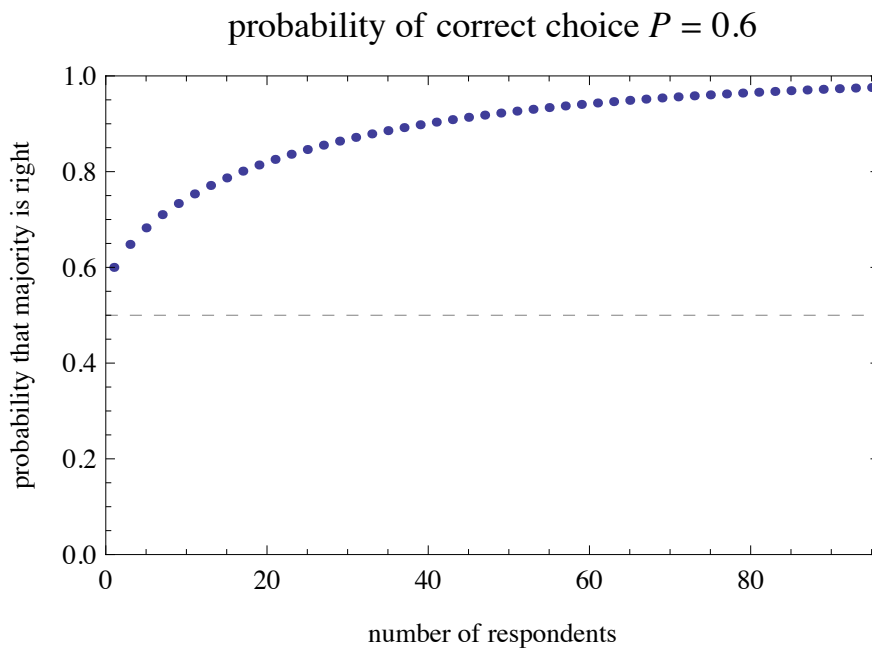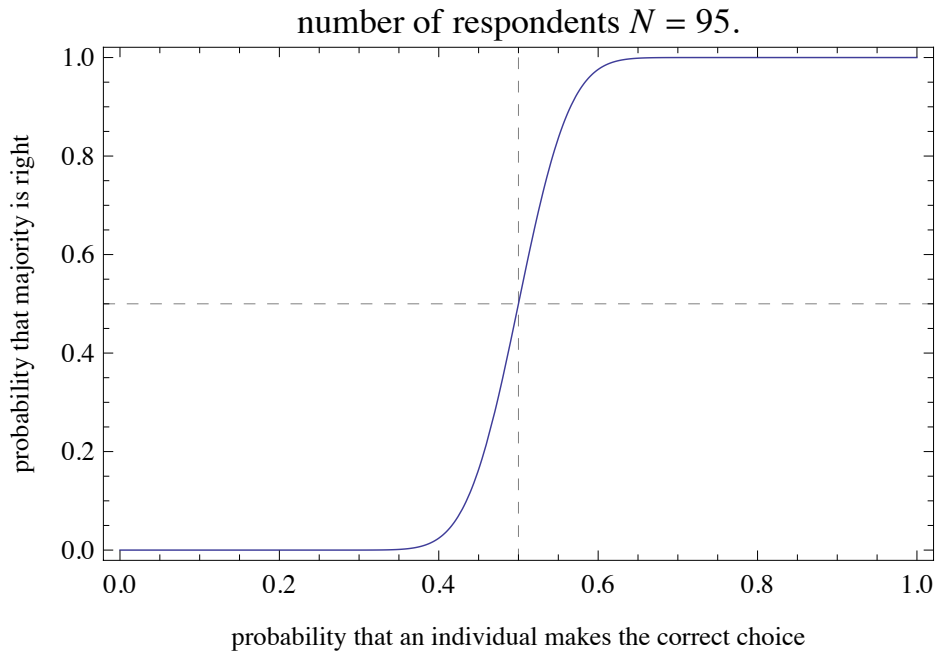## probability of correct choice $P = 0.6$

number of respondents

Figure 1

Even more striking are related theorems that relax some of these rather idealized assumptions, and, in particular, the assumption of uniform competence. Thus, with p_a

representing the average competence of members of the group, if the distribution of p_i is symmetric, then (op. cit. p. 268)

> **If 1 > p_a > ½, then P_N is monotonically increasing in N and $\lim_{n \to \infty} P_N \to 1$**; if 0 < p_a < ½, then P_N is monotonically decreasing in N and $\lim_{n \to \infty} P_N \to 0$; while if p_a = ½ then P_N = ½ for all N.

Furthermore (op. cit. p. 269),

> For heterogeneous groups, if p_i > 0.5 for all i, then the greater the size of the majority in favor of an alternative, the more likely is that alternative to be the correct choice.

Finally, there is a Generalized Condorcet Jury Theorem:

> **If p_a > 0.5 then $\lim_{N \to \infty} P_N \to 1$**; if p_a < 0.5 then $\lim_{N \to \infty} P_N \to 0$; while if p_a = 0.5, 1 − $e^{½}$ < $\lim_{N \to \infty} P_N$ < $e^{½}$, i.e., 0.39 < P_N < 0.61.

What this means is that, so far as the positive program goes, all we need to know is that the participants have a positive bias toward truth and are independent, and to have enough participants, in order to get results that are probative. Ladha (1992) has extended this to show that even when there are correlations among judgments, in a wide range of conditions for large groups the majority outperforms any individual. Abstracting a bit from details, let us call the assumption that enough members of groups surveyed on thought experiments have a sufficiently positive bias and make judgments sufficiently independently of one another for the majority results to have a significant probability of being right the Positive Bias Assumption.[13]

It is an empirical question whether for any given group, e.g., of philosophically unsophisticated undergraduates, or random people enjoying a Saturday in Central Park, the Positive Bias Assumption is correct, in the sampling context. In connection with this, I want to raise three questions.

1. Is it presupposed by the standard view of the method of thought experiments?
2. Do we have reason to believe it?
3. What would it take to establish the needed assumption?

I raise the first question because if it is a presupposition of the method of thought experiments in philosophy that for any random group of individuals (or for most random groups of individuals) most will have a positive bias toward the truth (let this stand in for the constellation of conditions required), then survey philosophers don't need to argue

---

[13] See (Talbot 2014) for an argument, based on these kinds of considerations, for the probative value of current philosophical practice with thought experiments, and the point at the end of section 3, #6.

against their traditionalist opponents that the assumption is correct. Their opponents will be committed to it already, and so they would have to regard the results as probative, and as the numbers surveyed get larger and larger, other things being equal, the results would have to be given more and more weight.

However, while the traditional philosopher is committed to it being possible to make advances in analysis by careful consideration of thought experiments, she is certainly not committed to thinking that most of those in groups being surveyed in the conditions of survey philosophy have a positive bias toward providing a correct response. Competence based accounts are committed to concept possession putting us in a position to articulate the application conditions of concepts. But they do not entail that it is easy, or that individuals might not for various reasons, in various circumstances, particularly for subtle questions, in an unfamiliar practice, be more likely than not to answer a question about a scenario incorrectly.

This makes the next two questions salient. Do we have reason to think that the Positive Bias Assumption is correct? There is no question that it is *always* correct. For example, the gambler's fallacy is quite common (similarly the hot hand fallacy, the assumption that one's probability of guessing correctly an outcome goes up after having guessed correctly a number of times in a row). Someone subject to the gambler's fallacy will think it more likely that the flip of a fair coin will yield a tails if it comes after a series of flips that turned up heads. The prevalence of the gambler's fallacy has been demonstrated both in laboratory and in real world settings (Croson and Sundali 2005).
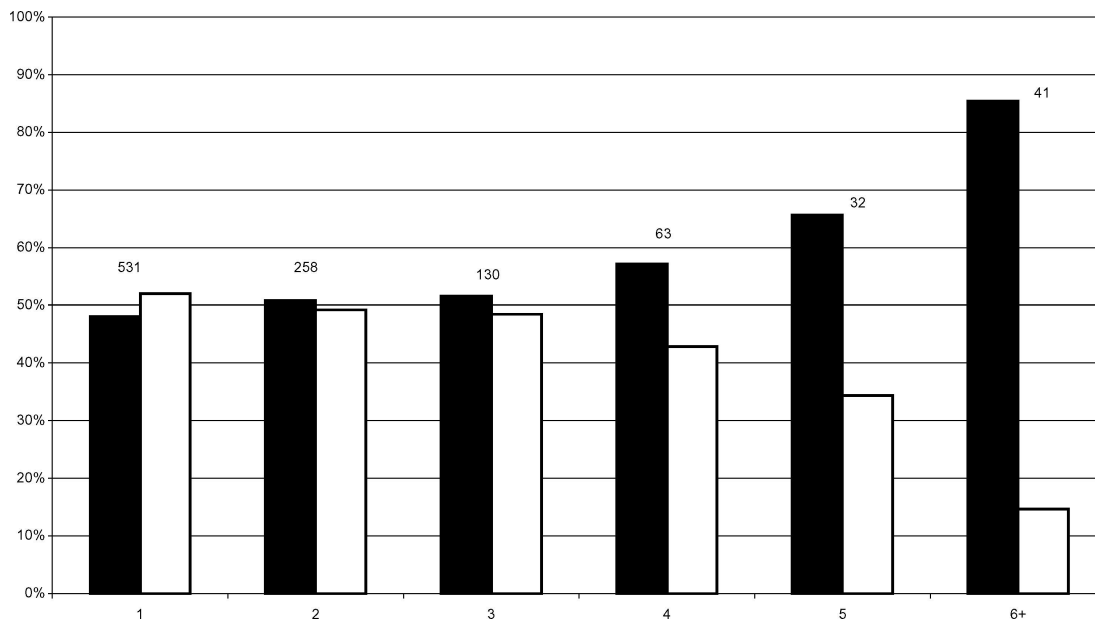
Proportion of gambler's fallacy outside bets after a streak of at least length N from data gathered at a large casino in Reno, Nevada, from 18 hours of security videotape of the roulette table. (Croson and Sundali 2005, 203)

The standard explanation is that the gambler's fallacy is the result of a representativeness bias. People expect short sequences to reflect the underlying probabilities, and so think a string of heads should be balanced by a tail.

Similarly, data from surveys also show that we can't in general assume that the majority in a group is correct. This is shown by the fact that can be variations in the majority judgments between groups which differ with respect to culture and socio-economic background, at least on some occasions (Nichols, Stich, and Weinberg 2003; Machery et al. 2004), as well as relative amount of training in the philosophy and so on. Since different runs with different groups yield different results, we cannot say that whatever the majority says is correct. Moreover, the fact that ordering effects make a difference to the distribution of responses in some cases shows that judgments can be affected by factors that are not truth related (Liao et al. 2012).

We can seek larger groups, but it seems unlikely, as the case of the gambler's fallacy shows, that simply increasing the size itself guarantees that most members of the group will have a positive bias toward truth in the relevant conditions. There are other sorts of errors we would expect to be systematic as well, such as responding to standardized conversational implicatures, rather than to what is literally expressed in a scenario.

Thus, in some cases, we know that the majority result is not correct and this cannot be corrected simply by increasing the size of the group. Can we assume nonetheless that in most cases the majority result will be correct? Well, in most cases of what? We might say: in most cases in which the judgment is very simple, and there are no distorting or distracting features of the case or the conditions, and subjects understand what they are doing and are motivated to try to respond in the right way in response just to what the scenario says. Maybe so. But, first, how are we to test for whether or not there are distorting or distracting features of the case? We might find something that obviously is irrelevant to truth correlated with responses—like the choice of names of characters in a scenario or ordering effects or cultural background or personality type—but if there is nothing obvious, can we assume there are no distracting features of the case? Second, aren't many of the cases we are interested in not simple cases? Third, in the end, don't we want to know in *particular cases* whether or not the majority result is likely to be true?

This brings us to the last of the three questions I want to take up: what would it take to establish the Positive Bias Assumption? We have a method that purports to help us discern the truth about the domain of conceptual truths. The method is indirect in the sense that it appeals to data about the judgments of agents relative to the assumption that they have a positive bias toward the correct judgment. We third person observers of the majority response in a group of subjects to questions about the scenario in a thought experiment want to know whether this is a correct expression of their conceptual competence. If our only access to the domain is through their judgments, however, we have no independent method of checking on whether they are getting it right or not. It looks as if, in order to support the Positive Bias Assumption, we need to sample the domain independently of the response of the majority. Can we do this?

Yes, why not? Isn't that how we determine that the gambler's fallacy is a fallacy? It is a fallacy, we can see, because each flip of a fair coin is independent of the others, and, hence, for each flip, the probability that a heads or a tails comes up remains constant. On the assumption that the coin is fair, the probability that it comes up heads or tails is .5 each time it is flipped. That *follows* from the assumptions. We have independent access to the domain in question. We can see ourselves what the correct thing to say is. And seeing that, we can see that what most people may say or think about the probability of tails after a sequence heads is a mistake, and explain to them why it is a mistake, and why they made it.

So it is in the case of surveys of the philosophically unsophisticated. If we want to know whether the majority opinion is correct or probative, and so whether something like the Positive Bias Assumption is correct, we can check to see whether most of our subjects get it right. If they do, then that is evidence that the Positive Bias Assumption is correct. Then, having provided evidence for the Positive Bias Assumption, we can use the survey data to advance one or another particular thesis about a concept's application conditions on that basis.

But wait a minute! *If* to confirm the Positive Bias Assumption, in a particular case, we have to *independently* decide the truth of the matter about what we were interested in testing, then … we don't need it. Once we are in a position to use the Survey Data to contribute to a philosophical problem, we have settled the matter. Furthermore, it could give us no more justification than we already had, since any confidence we have in the majority being right, or even in the probability that the majority is right, rests on our confidence that the Positive Bias Assumption is right, and that rests on our confidence that we have independently made the right judgment. To try to raise our confidence by this means would appear to be a kind of illegitimate bootstrapping of our epistemic position.

Perhaps it was evident from the start that the solution offered was a poison pill. While it seems the most straightforward way for one to try to establish that for a given population the Positive Bias Assumption is correct, it would undermine the point of survey philosophy (on the continuity approach). It is obvious also, though, that one cannot simply assume that it is true. It needs defense, and it needs defense on a case by case basis. Is there an indirect method of arguing in particular cases that we should give some credence to the Positive Bias Assumption, especially when the results might seem to have the most potential for overturning philosophical orthodoxy, that is, when most philosophers don't agree with the majority in survey results? We could try to test for general competency and then try out thought experiments with subjects who have been confirmed to be good at making the right judgment. Perhaps that is what we do informally in philosophy already. But this seems unlikely to be much help in cases in which there is controversy.

4.3 Does x-phi deliver what we want?

Suppose that we can solve the confirmation problem for x-phi, so that we can provide reasons to think that in particular cases the Positive Bias Assumption is true without independently establishing the right result, so that it could actually contribute to advancing a thesis about the application conditions of a concept. Would this deliver what we wanted?

Why not? After all didn't we want to know what the application conditions for a certain concept were? Suppose that the concept were the concept of justice, for example, and suppose that we had someone construct a theory of justice using the results of surveys in an especially clever and sophisticated way. There we would have it, finally, we would know what justice is!

Suppose that a young Greek philosopher in antiquity is interested in the question what justice is, and she has learned of the Oracle at Delphi. She decides then to travel to Delphi to ask the Oracle. Suppose that the Oracle does not speak in riddles, and that it tells her that justice in the state is when individuals receive benefits according to their merits. Suppose that she justifiably believes that the Oracle has a high positive bias toward the truth. She writes down what the Oracle has told her, and she is justified in believing that it is true. Should she be satisfied?

Now suppose an undergraduate student is interested in the nature of justice. She decides to survey faculty who teach courses related to justice in political science, philosophy, cognitive science, sociology and so on. A majority of those she surveys agree that justice is that state of society represented by arrangements which people would choose were they to have been constrained to choose it behind a veil of ignorance, not knowing their natural abilities or position in society or their sex or race, culture, or individual tastes, but knowing that in any society there would be a distribution of abilities, differences in sex, race, culture, individual tastes, and so on. Suppose that she justifiably believes that the faculty she has surveyed have a positive bias toward the truth. She accepts what they say and is justified in believing it to be true on the basis of their authority. Should she be satisfied?

She should not be satisfied. Why not? The problem is not that what she believes is not true. Let us grant, in one or the other case, it is true, and in fact, we can grant, let us say, that our undergraduate student knows what justice is on the basis of authority. She knows what justice is in the sense that she can give a correct statement of it and she is justified (in a non-deviant way) in her belief. But what she doesn't have is any understanding of why *that* is what justice is. Or if she does, it does not come from the mere fact that she has ascertained that most of the people in positions of authority that she has asked about it agree on a characterization, and that the majority is likely to be right.

The situation would be the same for the undergraduate who learns that the Pythagorean Theory is true (or the Condorcet Jury Theorem) from reading it in a textbook but not from being given a proof of it. She would know the Pythagorean theorem, but she would not know why it was true. She would not have any understanding why the area of the square on the hypotenuse of a right triangle is the sum of the areas of the squares constructed on

its sides. The textbook for her is an oracle, a known reliable source, but not a reliable source that conveys an understanding of the ground for the truth.   (Imagine that you told a freshman who is not too sophisticated mathematically but also not completely ignorant that 1 = .99999999.... She might accept it on authority but find it puzzling. Everything changes when you point out that 1 = 3 x 1/3 = 3 x .333333 ... = .9999999....)[14]

What is missing from the x-phi conception of philosophical method is only the most important thing about philosophy. What we want in philosophy is not to know things on the basis of their being the pronouncements of someone who is competence to say, but to know them out of our own competence, to see why they are true. Even if we could trust surveys, they would never give us the kind of knowledge we seek in philosophy. You might as well ask an oracle.

4.4 Can Experimental Philosophy help Philosophy?

Suppose you grant that x-phi is not a replacement or supplement for philosophy, even in the relatively narrow domain of philosophical activity that it concerns itself with. Is that the end of the line for x-phi? Is all the elaborate surveying of the unsophisticated to be relegated to some niche branch of psychology, for whoever might be interested in it? Or is there still some role for it to play as a kind of aid to core philosophical activities?

Experimental philosophy conceived of as a psychological enterprise can still be of some help. Psychology is not irrelevant to philosophical methodology. Our cognitive abilities are impressive, but we are not perfectly rational, we don't always get things right, and there are sometimes interesting explanations of this that only come to light with empirical investigation. The systematic study of the conditions under which we are liable to make mistakes helps us to guard against them. It can help us in teaching students as well, in choosing examples, and giving appropriate context.

That there are order effects in the judgments that people make about cases is an important insight for example (Schwitzgebel and Cushman 2012).[15] That people respond differently depending on the variance of factors that we are convinced are irrelevant to the correctness of a judgment provides a caution about our own judgments. This can help us to think about good thought experiment design, especially in the context of testing ideas with others. This doesn't show that we can never figure out what the right thing to say is, but it does show that figuring it out can be a complicated matter, and that we need to think

---

[14] Suppose that Survey Philosophy became firmly established as the dominant philosophical methodology. Philosophers cease to place any emphasis on their own reactions to thought experiments. They just survey their undergraduate students. And they tell their undergraduate students what the right answers are to traditional philosophical problems in accordance with what the majority say. Administrators find out about this, and decide they can cut out the middleman, and they find that student evaluations of philosophy courses go up because most students find that what they are taught in philosophy courses accords pretty well with what they already thought.

[15] Though as Cullen (Cullen 2010) shows, it is not always plausible that what generates order effects in survey responses has relevance to philosophical practice.

systematically about the subject matter and to take into account all of the possible factors that may incline us to error, and to try to correct for them so far as possible. All of this is familiar, of course, but details help.

How much does the identification of misleading factors in particular cases show about how far we can or cannot trust our judgments? In this volume, Stich and Tobias suggest that what it shows is that we can't get started using thought experiments in inquiry without first undertaking the experimental study of intuitions. This is an overreaction. They make this claim in responding to the point (made in an earlier paper of mine (Ludwig 2007)) that responses to surveys are not *ipso facto* philosophical intuitions.

> For Ludwig, a judgment or belief that is influenced by factors other than conceptual competence will not count as an intuition. ... A judgment that is influenced by order of presentation, or font size or fart spray is not *solely* the expression of conceptual competence. Thus when intuitions are characterized in this way, the studies pose no challenge to the use of intuitions as evidence in philosophy. But we are inclined to think that in making this move, Ludwig is hoist on his own petard. For the effects discussed in §3 are almost always *covert*; people have no conscious awareness that their beliefs or judgments are being affected. So [i] on Ludwig's restricted account, the only way to determine whether one's beliefs or judgments are intuitions is to do well designed and carefully controlled experiments. Rather than [ii] showing that the experimental studies are irrelevant to philosophical practice, [iii] Ludwig's restricted account of intuition leads directly to the conclusion that the sorts of studies that experimental philosophers have undertaken are required before we can begin to use intuitions as evidence, since without such studies we have no way of knowing which of our beliefs or judgments *are* intuitions. (p. nn; lower case roman numerals added)

There are a number of points to be made in response.

(1) The first is a point of interpretation. In the article that Stich and Tobias cite, I argued that extravagant claims made on the basis of surveys pose no threat to traditional philosophical methods, but I did not argue that surveys were irrelevant to philosophical practice, as here implied [ii]. Instead, I suggested, as I have above, that they can help us think about good thought experiment design and to identify pitfalls. In fact, except for the absence of the triumphalist tone, and my rejection of the (false and ungrounded) assumption that survey philosophy is essential to success in philosophy, my conclusion is not that different in its details from Stich and Tobias. Why all the fuss?

(2) The second is about the claim (in [i] and [iii]) that (a) the fact that in some cases there are factors that distort responses to thought experiments shows that (b) "we [generally] have no way of knowing which or our beliefs or judgments are intuitions" without first undertaking surveys about scenarios in thought experiments. In fact, the inference from (a) to (b) is a *non sequitur*. It doesn't follow from the fact that there are sometimes distorting influences at work in thought experiments that there always are. And it doesn't follow that because there are

sometimes distorting influences at work in responses to thought experiments that we can never know that our responses are intuitions—just as it doesn't follow from the fact that we sometime make mistakes in calculations that we never know that we have made a correct calculation! In response to the charge that there is no way we can ever tell whether we are in the good condition or the bad condition, and that if we can't, then we can never trust our intuitions, there are several things to say. The first is that it has not been demonstrated that we can never tell when we are in the good condition, and that it doesn't follow from the fact that if we are in the bad condition we can't tell (which has also not been demonstrated), that if we are in the good condition we cannot tell (cf. Bernard Williams's discussion of the cases of dreaming and anoxia (Williams 1978, appendix 3)). The second is that, waiving the last point, on the assumption that mistakes are not the norm, we can apply the method of reflective equilibrium to sort out what to say. The third is that there is empirical work that supports the common sense thesis that we are often implicitly aware of cases in which intuition instability (e.g., due to order effects) is a threat and that there are reliable methods for tracking it (introspective confidence and belief strength) (Wright 2010).

(3) The third is that the general claim Stich and Tobias make is self-defeating. Intuitions, on the view in question, are conceptually grounded judgments about entailments. Therefore, as a completely general thesis about judgments about entailments, it is clearly self-defeating to claim that we must first engage in surveys before we can ever tell whether our judgments about entailments are correct, since in evaluating what the surveys show us we must make judgments about entailments. For example, the judgment that order effects show that there is something problematic about responses to a thought experiment rests on the claim that order of presentation is irrelevant to the correct judgment. That order is completely irrelevant to the correct judgment and that judgments are influenced by order entails that the judgments are being distorted by truth-irrelevant features of the experimental set-up. I trust that readers feel confident that they know, in this case, that the one proposition entails the other. It a point of mild irony that the use of survey results to argue against the use of thought experiments in philosophy relies on an epistemology which could only be supported by the sources which it aims to undermine.

(4) The fourth is that it does not follow from the fact that, in some cases, there are factors that distort responses to scenarios, that to discover and correct this we need to do surveys. In actual philosophical practice, as noted earlier, sec. 3 #6, we aim to survey a range of cases, vary the cases studied, test them on others, compare results with previously established results, bring to bear general theoretical considerations, and when we notice a problem, try to get a clearer view of what is going on, in much the same way that we sort through a problem in geometry or mathematics. The journal literature in philosophy, among other things, carries out this task over time, just as the journal literature in the sciences carries out the task of self-correction and revision that is an on-going part of the advancement or science.

(5) The fifth is that to arrive at the right view, survey methods alone will not suffice, for in cases in which we suspect that there are distorting factors, the question of

whether the positive bias assumption is correct for a group surveyed becomes urgent. Intuitions are ultimately methodologically basic in this domain.

Finally, perhaps another way that x-phi is relevant to philosophy is that it can (potentially) throw up challenges to philosophical theories (though one might wonder whether it will raise challenges that wouldn't be raised anyway). For example, those who want to maintain the Simple View in the philosophy of action, that if someone intentionally Fs, then she intended to F, must confront the fact that there are circumstances in which many people say that someone intentionally F-ed but did not intend to F. Philosophers of course noticed this first, and indeed this was the impetus for the surveys (Harman 1976). Proponents of the Simple View need to offer an explanation for why people get it wrong. These explanations are of course empirical claims about distorting factors. It is not clear the extent to which these can actually be resolved using the survey method (or whether it is the best method (Levin 2009)), but it seems plausible enough that some testing of these theories could be conducted by the survey method, and that this can have some value.

5. Conclusion

What is the relation of x-phi to philosophy? A pessimist might say:

> The connection is that people who do x-phi are largely housed in philosophy departments, just as people who do religious studies are sometimes housed in philosophy departments. There is a social and an administrative connection. In addition, there is the appearance of the word 'philosophy' in 'Experimental Philosophy', and the fact that scenarios from philosophical thought experiments are featured in its surveys. They also sometimes publish in philosophy journals. In short: "experimental philosophy is a cuckoo-bird in the nest of philosophy" (Sorensen 2014).

Not so fast! X-phi is not promising as a positive program for contributing to the articulation of application conditions of concepts or conceptual connections. First, we have to have independent access to the domain to validate the assumptions that would enable us to use the data, which renders it otiose for our primary purpose, and, second, the kind of knowledge we want in philosophy is not supplied by asking an oracle. But even philosophers dismissive of x-phi⁻ and skeptical of the continuity account should be interested in the results of surveys, for it provides a map of the reactions of the unsophisticated (and in some cases of the sophisticated) and evidence bearing on influences on responses and mechanisms involved that are not related to truth. We can learn something from this, and it can help to make us more sophisticated about the construction and use thought experiments. In addition, some testing of empirical claims about errors that are made in response to thought experiments can be pursued fruitfully in this way. What then is the relation of x-phi to philosophy? X-phi is philosophy from the sociological point of view. From the standpoint of the tradition in philosophy, in so far as it is to be of aid to philosophy, it is more properly thought of as a non-philosophical activity or an adjunct to philosophical activity that has relevance to philosophy in the way in which the investigation of the foibles and frailties of human reasoning has relevance to the

practice of human reasoning. We should use all the tools we can assemble to advance our understanding.

References

Adleberg, Toni, Morgan Thompson, and Eddy Nahmias. 2014. Do Men and Women Have Different Philosophical Intuitions? Further Data. *Philosophical Psychology* 28 (5):615-641.

Alexander, Joshua, Ronald Mallon, and Jonathan M. Weinberg. 2010. Accentuate the Negative. *Review of Philosophy and Psychology* 1 (2):297-314.

Alexander, Joshua, and Jonathan Weinberg. 2007. Analytic Epistemology and Experimental Philosophy. *Philosophy Compass* 2 (1):56-80.

Bealer, George. 1998. Intuition and the Autonomy of Philosophy. In *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, edited by M. DePaul and W. Ramsey. New York: Rowman & Littlefield.

Booth, Anthony. 2014. *Intuitions*. Oxford: Oxford University Press.

Buckwalter, Wesley. 2014. Intuition Fail: Philosophical Activity and the Limits of Expertise. *Philosophy and Phenomenological Research* 89 (2).

Buckwalter, Wesley, and Jonathan Schaffer. 2013. Knowledge, Stakes, and Mistakes. *Noûs* 47 (1).

Buckwalter, Wesley, and Stephen Stich. 2014. Gender and Philosophical Intuition. In *Experimental Philosophy, Vol.2*, edited by J. Knobe and S. Nichols: Oxford University Press.

Cappelen, Herman. 2012. *Philosophy without Intuitions*. Oxford: Oxford University Press.

———. 2014. X-Phi without Intuitions? In *Intuitions*, edited by A. R. Booth and D. P. Rowbottom. Oxford: Oxford University Press.

Collaboration, Open Science. 2015. Estimating the reproducibility of psychological science. *Science* 349 (6251).

Croson, Rachel, and James Sundali. 2005. The Gambler's Fallacy and the Hot Hand: Empirical Data from Casinos. *The Journal of Risk and Uncertainty* 30 (3):195-209.

Cruz, Helen De. 2014. Where Philosophical Intuitions Come From. *Australasian Journal of Philosophy*.

Cullen, Simon. 2010. Survey-Driven Romanticism. *Review of Philosophy and Psychology* 1 (2):275-296.

DePaul, Michael, and William Ramsey, eds. 1998. *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*. New York: Rowman & Littlefield.

Deutsch, Max. 2009. Experimental philosophy and the theory of reference. *Mind and Language* 24 (4):445-466.

Doris, John. 2002. *Lack of Character: Personality and Moral Behavior*. Cambridge: Cambridge University Press.

Doris, John M. 1998. Persons, Situations, and Virtue Ethics. *Noûs* 32 (4):504-530.

Genone, James, and Tania Lombrozo. 2012. Concept Possession, Experimental Semantics, and Hybrid Theories of Reference. *Philosophical Psychology* 25 (5):1-26.

Gettier, Edmund. 1963. Is Justified True Belief Knowledge? *Analysis* 23:121-3.

Goldman, Alan. 2007. Philosophical Intuitions. *Grazer Philosophische Studien* 74:1-26.

———. 2010. Philosophical Naturalism and Intuitional Methodology. *Proceedings of the American Philosophical Association* 84 (2):115-150.

Grofman, Bernard, Guillermo Owen, and Scott L. Feld. 1983. Thirteen theorems in search of the truth. *Theory & Decision* 15:261-78.

Grundmann, Thomas. 2010. Some Hope for Intuitions: A Reply to Weinberg. *Philosophical Psychology* 23 (4):481-509.

Hales, Steven D. 2012. The faculty of intuition. *Analytic Philosophy* 53 (2):180-207.

Harman, Gilbert. 1976. Practical Reasoning. *Review of Metaphysics* 79:431-63.

———. 1999. Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error. *Proceedings of the Aristotelian Society* New Series 119:316-331.

———. 2000. The Nonexistence of Character Traits. *Proceedings of the Aristotelian Society* 100:223-226.

Hitchcock, Christopher. 2012. Thought experiments, real experiments, and the expertise objection. *European Journal for Philosophy of Science* 2 (2):205-218.

Jackson, Frank. 2011. On Gettier Holdouts. *Mind and Language* 26 (4):468-481.

Kamtekar, Rachana. 2004. Situationism and Virtue Ethics on the Content of Character. *Ethics* 114 (3):458-491.

Kauppinen, Antti. 2014. The rise and fall of experimental philosophy. In *Experimental Philosophy*, edited by J. Knobe and S. Nichols. Oxford: Oxford University Press.

Kim, Minsun, and Yuan Yuan. No Cross-Cultural Differences in Gettier Car Case Intuition: A Replication Study of Weinberg et al. 2001.

Knobe, Joshua. 2003a. Intentional Action and Side Effects in Ordinary Language. *Analysis* 63 (3):190-194.

———. 2003b. Intentional Action in Folk Psychology: An Experimental Investigation. *Philosophical Psychology* 16 (2):309-324.

———. 2004. Intention, Intentional Action and Moral Considerations. *Analysis* 64 (2):181-187.

———. 2006. The Concept of Intentional Action: A Case Study in the Uses of Folk Psychology. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 130 (2):203-231.

———. 2007. Experimental philosophy. *Philosophy Compass* 2 (1):81–92.

Knobe, Joshua, Wesley Buckwalter, Philip Robbins, Hagop Sarkissian, Tamler Sommers, and Shaun Nichols. 2012. Experimental Philosophy. *Annual Review of Psychology* 63 (50):72-73.

Knobe, Joshua, and Aruda Burra. 2006. Experimental Philosophy and Folk Concepts: Methodological Considerations. *Journal of Cognition and Culture* 6 (1-2):331-342.

Knobe, Joshua Michael, and Shaun Nichols. 2008. *Experimental philosophy*. Oxford ; New York: Oxford University Press.

———. 2014. *Experimental philosophy*. Vol. 2. Oxford: Oxford University Press.

Kornblith, Hilary. 2002. *Knowledge and Its Place in Nature*. Oxford: Clarendon Press.

Ladha, Krishna K. 1992. The Condorcet Jury Theorem, Free Speech, and Correlated Votes. *American Journal of Political Science* 36 (3):617-634.

Lam, Barry. 2010. Are cantonese speakers really descriptivists? Revisiting cross-cultural semantics. *Cognition* 115 (2):320–32.

Levin, Janet. 2009. Experimental philosophy. *Analysis* 69 (4):761-769.

Liao, S. Matthew, Alex Wiegmann, Joshua Alexander, and Gerard Vong. 2012. Putting the trolley in order: Experimental philosophy and the loop case. *Philosophical Psychology* 25 (5):661-671.

Ludwig, Kirk. 2007. The Epistemology of Thought Experiments: First Person versus Third Person Approaches. *Midwest Studies in Philosophy: Philosophy and the Empirical* 31:128-159.

———. 2010. Intuitions and Relativity. *Philosophical Psychology* 23 (4):427-445.

———. 2013. Methods in Analytic Philosophy. In *Philosophical Methodology: The Armchair or the Laboratory*, edited by M. Haug. New York: Routledge.

Machery, E., R. Mallon, S. Nichols, and S. Stich. 2004. Semantics, Cross-Cultural Style. *Cognition* 92:B1-B12.

Malle, Bertram F., and Joshua Knobe. 2001. The Distinction between Desire and Intention: A Folk-Conceptual Analysis. In *Intentions and Intentionality: Foundations of Social Cognition*, edited by B. F. Malle, L. J. Moses and D. A. Baldwin. Cambridge, Mass.: MIT Press.

Naess, Arne. 1953. *Interpretation and Preciseness: a contribution to the theory of communication*. Oslo: I Kommisjon Hos J. Dybwad.

Nagel, Jennifer. 2012. Intuitions and Experiments: A Defense of the Case Method in Epistemology. *Philosophy and Phenomenological Research* 85 (3):495-527.

Nagel, Jennifer, Valerie San Juan, and Raymond A. Mar. 2013. Lay Denial of Knowledge for Justified True Beliefs. *Cognition* 129:652-661.

Nahmias, Eddy, Stephen G. Morris, Thomas Nadelhoffer, and Jason Turner. 2006. Is Incompatibilism Intuitive? *Philosophy and Phenomenological Research* 73 (1):28 - 53.

Nichols, Shaun. 2011. Experimental Philosophy and the Problem of Free Will. *Science* 331 (6023):1401-1403.

Nichols, Shaun, and Joshua Knobe. 2007. Moral responsibility and determinism: The cognitive science of folk intuitions. *Noûs* 41 (4):663–685.

Nichols, Shaun, Stephen Stich, and Jonathan M. Weinberg. 2003. Meta-Skepticism: Meditations in Ethno-Epistemology. In *The Skeptics: Contemporary Essays*, edited by S. Luper. Aldershot: Ashgate Publishing.

Papineau, David. 2011. What is x-phi good for? *The Philosophers' Magazine* 52:83-88.

———. 2013. The Poverty of Conceptual Analysis. In *The Armchair or the Laboratory*, edited by M. Haug. New York: Routledge.

Pettit, Dean, and Joshua Knobe. 2009. The Pervasive Impact of Moral Judgment. *Mind and Language* 24 (5):586-604.

Pritchard, Duncan. 2014. Sceptical Intuitions. In *Intuitions*, edited by A. R. Booth and D. P. Rowbottom. Oxford: Oxford University Press.

Putnam, Hilary. 1965. It Ain't Necessarily So. *Journal of PHilosophy* LIX:658-671.

Quine, Willard Van Orman. 1953. Two Dogmas of Empiricism. In *From a Logical Point of View*. Cambridge: Harvard University Press.

Rini, Regina A. 2013. Analogies, Moral Intuitions, and the Expertise Defence. *Review of Philosophy and Psychology* (2):1-13.

Rose, David, and David Danks. 2013. In Defense of a Broad Conception of Experimental Philosophy. *Metaphilosophy* 44 (4):512-532.

Sarkissian, Hagop, Amita Chatterjee, Felipe De Brigard, Joshua Knobe, Shaun Nichols, and Smita Sirker. 2010. Is Belief in Free Will a Cultural Universal? *Mind and Language* 25 (3):346-358.

Sarkissian, Hagop, John Park, David Tien, Jennifer Wright, and Joshua Knobe. 2011. Folk
        Moral Relativism. *Mind and Language* 26 (4):482-505.
Schwitzgebel, Eric, and Fiery Cushman. 2012. Expertise in Moral Reasoning? Order Effects
        on Moral Judgment in Professional Philosophers and Non-Philosophers. *Mind and
        Language* 27 (2):135-153.
Seyedsayamdost, Hamid. 2014a. On gender and philosophical intuition: Failure of
        replication and other negative results. *Philosophical Psychology*:1-32.
———. 2014b. Reproducibility of empirical findings: experiments in philosophy and
        beyond, Philosophy, Logic and Scientific Method, London School of Economics,
        London.
———. 2015. On Normativity and Epistemic Intuitions: Failure of Replication. *Episteme* 12
        (1):95-116.
Snow, Nancy. 2010. *Virtue as Social Intelligence: An Empirically Grounded Theory*. New York:
        Routledge.
Sommers, Tamler. 2010. Experimental philosophy and free will. *Philosophy Compass* 5
        (2):199-212.
Sorensen, Roy. 2014. Novice Thought Experiments. In *Intuitions*, edited by R. Booth and D.
        P. Rowbottom. Oxford: Oxford University Press.
Sosa, Ernest. 2007a. Experimental philosophy and philosophical intuition. *Philosophical
        Studies* 132 (1):99-107.
———. 2007b. Intuitions: Their Nature and Epistemic Efficacy. *Grazer Philosophische
        Studien*.
———. 2008. A Defense of the Use of Intuitions in Philosophy. In *Stich and His Critics*,
        edited by M. Bishop and D. Murphy. Oxford: Blackwell.
———. 2011. Can There Be a Discipline of Philosophy? And Can It Be Founded on
        Intuitions? *Mind and Language* 26 (4):453-467.
Stacy Swain, Joshua Alexander, Jonathan Weinberg. 2008. The Instability of Philosophical
        Intutions: Running Hot adn Cold on Truetemp. *Philosophy and Phenomenological
        Research* 76 (1):138-155.
Talbot, Brian. 2013. Reforming intuition pumps: when are the old ways the best?
        Philosophical Studies 165 (2):315-334.
———. 2014. Why so negative? Evidence aggregation and armchair philosophy. *Synthese*
        191 (16):3865-3896.
Thomson, Judith Jarvis. 1976. Killing, Letting Die, and the Trolley Problem. *The Monist* 59
        (2):204-217.
Weinberg, Jonathan. 2007. How to Challenge Intuitions Empirically without Risking
        Skepticism. *Midwest Studies in Philosophy* 31 (1):318-343.
Weinberg, Jonathan, Chad Gonnerman, Cameron Buckner, and Joshua Alexander. 2010. Are
        Philosophers Expert Intuiters? *Philosophical Psychology* 23 (3):331-355.
Weinberg, Jonathan, Shaun Nichols, and Stephen Stich. 2001. Normativity and Epistemic
        Intuitions. *Philosophical Topics* 29:429-460.
Williams, Bernard. 1978. *Descartes: The Project of Pure Enquiry*. London: Pelican.
Williamson, Timothy. 2011. Philosophical Expertise and the Burden of Proof.
        Metaphilosophy 42 (3):215-229.
Wright, Jennifer. 2010. On intuitional stability: The clear, the strong, and the paradigmatic.
        Cognition 115 (3):491-503.

Young, Liane, Fiery Cushman, Ralph Adolphs, Daniel Tranel, and Marc Hauser. 2006. Does emotion mediate the effect of an action's moral status on its intentional status? Neuropsychological evidence. *Journal of Cognition and Culture* 6:291-304.