40 See my "Timothy Williamson's *The Philosophy of Philosophy*," *Analysis* 69 (2009): 109–16; and Williamson's "Replies to Kornblith, Jackson and Moore," *Analysis* 69 (2009): 125–35.

41 Thus, Williamson remarks, "I expect armchair methods to play legitimately a more dominant role in future philosophy than [Kornblith] expects them to – of course, such difference in emphasis can result in widening differences in practice." "Replies to Kornblith, Jackson and Moore," 126.

42 Ibid., 126.

43 In "The Role of Intuition in Philosophical Inquiry," in M. DePaul and W. Ramsey (eds), *Rethinking Intuition* (Lanham, MD: Rowman & Littlefield, 1998), 129–41; in *Knowledge and Its Place in Nature*; in "Appeals to Intuition and the Ambitions of Epistemology," in Stephen Hetherington (ed.), *Epistemology Futures* (Oxford: Oxford University Press, 2006), 10–25; in "Naturalism and Intuitions"; and in "Timothy Williamson's *The Philosophy of Philosophy*."

44 "Replies to Kornblith, Jackson and Moore," 128.

45 "What Am I to Believe?," in Steven Wagner and Richard Warner (eds), *Naturalism: A Critical Appraisal* (Notre Dame, IN: University of Notre Dame Press, 1993), 148.

46 "Indispensability of Internalism," 54.

47 Ibid., 62.

48 This is the way that Conee and Feldman characterize internalism. See *Evidentialism: Essays in Epistemology* (Oxford: Oxford University Press, 2004), 56: "The justificatory status of a person's doxastic attitudes strongly supervenes on the person's occurrent and dispositional mental states, events, and conditions."

49 This is going to be so on any way that one tries to make sense of the direct/indirect distinction. I do not mean to be endorsing any particular way of making out such a distinction, or, indeed, even the claim that we can make good sense of it.

50 "Indispensability of Internalism," 54.

51 And various extramental items as well. BonJour rightly points out that, on standard internalist accounts, we have such unproblematic access by way of reflection to at least some a priori knowable truths. "Indispensability of Internalism," 55. I will focus in the text, however, on our access to relevant features of our mental lives. If internalists cannot secure this, then internalism is committed to an extremely broad skepticism, and, indeed, the coherence of the entire position is threatened. More than this, there is every reason to believe that the kind of argument I make against unproblematic access to mental items can easily be generalized to cover the cases in which BonJour believes we have unproblematic access to a priori knowable truths. I have discussed these problems about the a priori in "The Impurity of Reason," *Pacific Philosophical Quarterly* 81 (2000): 67–89. It would take us too far afield from the issues under discussion here to pursue these questions about the a priori.

52 "Indispensability of Internalism," 53.

53 "Replies," 292.

54 BonJour seems to come quite close to acknowledging this. At one point he remarks that,

> Certainly it would be a very unusual brand of scepticism which would challenge whether my belief that B is justified by raising the issue of whether I do in fact accept B, the normal sceptical claim being precisely that certain beliefs which are in fact held are nonetheless unjustified.
>
> (*The Structure of Empirical Knowledge*, 81)

But claims about which skeptical challenges are common in the history of philosophy, or which are unusual, tell us nothing about which kinds of claims are unproblematic *tout court*. And it is this sense of what is unproblematic, rather than the dialectical sense, which is epistemologically relevant.

55 I am indebted to Timothy Williamson for conversations on this topic on a number of occasions. In addition, I have presented versions of this paper at Fortaleza, Brazil; the University of St Andrews; and the University of Cincinnati. I am grateful to audiences on all of these occasions for helpful comments and criticisms. Thanks too to Matthew Haug for helpful suggestions.

13 Methods in analytic epistemology

KIRK LUDWIG

In this chapter, I defend the program of conceptual analysis, broadly construed, and the method of thought experiments in epistemology, as a first-person enterprise, that is, as one which draws on the investigator's own competence in the relevant concepts. I do not suggest that epistemology is limited to conceptual analysis, that it does not have important a posteriori elements, that it should not draw on empirical work wherever relevant (and non-question-begging), or that it is not a communal enterprise. Although discussion in the space available will necessarily be brief, and many points must be elided altogether, I aim to sketch salient features of the landscape, clarify issues, set aside some confusions, and outline responses to some recent challenges.

In the next section, "What Are Concepts?," I sketch a traditional account of concepts and conceptual truths. In the section following, "What Is Conceptual Analysis?," I review a broad conception of analysis as encompassing not just reduction but also articulation of conceptual connections. In "How Could Conceptual Analysis Tell Us Anything about the World?," I address the charge that in studying epistemic concepts we turn away from our proper target of study, the actual phenomena of knowledge, justification, and so on. In "What Is the Role of Thought Experiments?," I give a brief overview of the method of thought experiments. Then in "What Are the Lessons of Experimental Philosophy?," I address objections to thought experiments that have their source in "experimental philosophy." Finally, in my concluding section, "Is 'Knowledge' a Natural Kind Term?," I address the charge that pursuing conceptual analysis in epistemology is misplaced because "knowledge," "justification," "evidence" and so on, are natural kind terms, and hence that we must engage in empirical research to discover their real essences.

What are concepts?

Concepts, in the sense we are concerned with, are common (general) elements in different thought contents. For example, the thoughts that chess is a

strategic game and that chess is a popular game are distinct but share the concepts of chess and of games. The concepts in a thought content, and how they are combined in it, determine its truth conditions, which in turn individuate it.[1] Concepts in turn are individuated by the systematic contributions they make to the truth conditions of thought contents. Concepts fall in different categories. There are monadic, binary, triadic, etc., concepts expressed with one-place, two-place, three-place predicates, etc. But there are also logical concepts such as that of negation, conjunction, disjunction, universal and existential quantification, and so on. I will call the conditions for the correct deployment of concepts their application conditions. If a noun or adjective "*F*" expresses a concept, we specify it as the concept of *F*, and we say that something falls under the concept of *F* iff it has the property of being *F*.

A conceptual truth is true in virtue of its contained concepts and their mode of combination. It is a conceptual truth that the arithmetic mean of a range of numbers lies within it, as it is a conceptual truth that the arithmetic mean is always greater than or equal to the geometric mean. Similarly it is a conceptual truth that if something is completely transparent, then it is not colored, that for any rigid bodies *a*, *b*, *c*, if *a* is longer than *b* and *b* is longer than *c*, then *a* is longer than *c*, that no one knows that the moon is larger than the earth if it is not true, that no person is identical to two distinct people, and so on. These are conceptual truths in the sense that we can explain why they are true by adverting to facts about the application conditions of the contained concepts and their mode of combination.

The link between something's being a conceptual truth and our being in a position to come to know it can be expressed in three connected theses.

- *Possession*. To think that *p*, one must possess the concepts involved in it.
- *Competence*. To possess a concept *C*, one must be competent in its deployment, in the sense of being in a position to deploy it correctly in thought, on the basis of its application conditions, in response to conditions as one takes them actually to be or in response hypothetically to conditions so specified.
- *Recognition*. A judgment or thought that *p* which is an expression solely of one's competence in deploying the concepts involved in light of their mode of combination counts as knowing (or being justified in believing) that *p* on the basis of their application conditions.

To have a concept is to have a competence expressed in thinking rightly in response to conditions that are relevant to the truth conditions that individuate the thought. Having thoughts is linked to competences connected to their components and their structure and implies the possibility of knowledge on that basis.[2]

Possession has been challenged on the ground that when a speaker uses a sentence intending to use it as others in his community do, we are licensed (at least sometimes) in attributing to him an attitude whose content is given by the sentence, even if the speaker doesn't fully understand all the words in it (Burge 1979). This would allow (us to say) that someone had a thought though he failed to possess all the concepts in it. Whether this is right or not, clearly we could engage in such a practice. The possession condition then should be understood to exclude attitudes whose attribution rests on such a practice—which could not exist in the first place if there were no thoughts attributed independently of it.

Competence is a weak condition: it says only that one is in a position to deploy concepts one possesses correctly in thought.[3] It does not say that one invariably does, or that it is easy to see the right thing to think in response to conditions that are relevant.[4] Nor does it say that deploying a concept correctly in thought is always a matter of recognizing features of objects. One's possession of the concepts of disjunction and negation will be expressed in part, e.g., in how one responds to accepting that *p* when one believes that *q* or not-*p*. Even for non-logical concepts application might not be based on accepting a set of propositions, but involve rather certain experiences, e.g., thinking something red given how it looks, or application on the basis of similarity to a prototype, etc., given other beliefs.[5]

Suppose that it is a conceptual truth that anything that is red is colored. If one judges on the basis of the concepts involved that anything red is colored, one thereby knows that anything red is colored. A judgment one reaches on the basis solely of competence in the deployment of the concepts involved in it is not based on facts about experience, the deliverances of introspection, or memory. Hence, on one plausible way of understanding the claim, the knowledge one has is a priori.

What is conceptual analysis?

Conceptual analysis in philosophy subsumes two projects, one narrower and one broader, though critics often seem to focus only on the first of these. There is, on the one hand, the project of providing informative necessary and sufficient conditions for the application of a concept, i.e., reductive analysis. On the other hand, there is the project of tracing constitutive connections between concepts, propositions, and experience, and ordering

families of concepts, so far as that is possible, in terms of relative priority, i.e., conceptual elucidation (Strawson 1992, ch. 2; McGinn 2012, ch. 7). A concept or family of concepts is prior to another just in case one can have it without having the other but not vice versa.

Conceptual analysis is sometimes dismissed on the grounds that few concepts of philosophical interest admit of informative analyses. But the interest of getting a clear view of the conceptual structure of the world is hardly exhausted by an interest in reductive analysis. It was never in the cards that we would get informative necessary and sufficient conditions for every concept of philosophical interest. On pain of an infinite regress, we must find some concepts that we can't analyze in terms of other more basic concepts. And the concepts that structure most deeply how we think about the world are just those that we should not think capable of reductive analysis. Here we turn to elucidation.

Nor should conceptual analysis be rejected with conceptual atomism, the analog of foundationalism for concepts, which holds that all concepts can be reduced to a set of basic concepts. An example is a version of empiricism that holds that all simple ideas (blurring the distinction between sensation, experience, and concept) derive from sensory experience and that complex ideas are built up out of them. This is not a definition of conceptual analysis but a substantive hypothesis about the global structure of our concepts, and the way that our concepts hang together may not conform to the model. A sense of the possibilities is suggested by Davidson's program in the theory of meaning (see the discussion in Davidson (2001, 137)), on which the idea is that a family of interlocking concepts (of meaning, truth, belief, desire, intention, agency, rationality, etc.) may be illuminated by tracing constitutive connections with a distinct family of concepts used to describe canonical evidence for a theory involving members of the first family, without one-by-one reduction of concepts, or even a holistic reduction of a theory to a set of statements about evidence for it.[6]

How could conceptual analysis tell us anything about the world?

Why think that conceptual analysis tells us anything about the world? The view that it does not has been forcefully stated by Hilary Kornblith:

> On my view, the subject matter of ethics is the right and the good, not our concepts of them. The subject matter of philosophy of mind is the mind itself, not our concept of it. And the subject matter of epistemology is knowledge itself, not our concept of knowledge.
>
> (2002, 2)

> By bringing in talk of concepts ... in an epistemological investigation, we only succeed in changing the subject: instead of talking about knowledge, we end up talking about our concept of knowledge.
>
> (2002, 9–10)

Analysis of the *concept* of knowledge is fine as far as it goes, but if we are really interested in *knowledge*, shouldn't we put the concept aside and look at the phenomenon itself?

This is a false dilemma. Conceptual analyses (for "predicative" concepts) are standardly presented in a biconditional of the form [C] in which "$F(\ \ldots\ )$" expresses the concept of interest (I focus on reductive analysis, but the point extends to elucidations).

> [C] For any $x, y, z, \ldots$ $F(x, y, z, \ldots)$ iff ...

An instance of [C] counts as a reductive analysis provided that (a) expressions used on the right-hand side express concepts that are more basic than the concept being analyzed and (b) [C] expresses a conceptual truth. [C] is a material mode statement. It does not mention any concepts, or any words, and in particular it does not mention the concept of which it gives an analysis. For illustration, take an instance (Klein 1971):

> [K] For any x, for any p, x knows p iff (i) p is true, (ii) x believes p, (iii) p is evident to x, and (iv) there is no proposition q such that if q became evident to x, then p would no longer be evident to x.

This is not about the *concept* of knowledge: it is about *knowledge* itself. If it is true, then it says something about the conditions under which someone has knowledge.

Then is it not a conceptual analysis after all? Knobe and Burra raise exactly this objection (2006, 332): "The problem with such an account is that it seems to say nothing about people's *concepts*. (It would tell us, not about people's concepts, but about the actual properties in the world that these concepts pick out.)" This rests on a misunderstanding. [C] is a material mode statement. When it is put forward as a conceptual analysis, it is claimed to meet conditions (a) and (b). This is a claim about the statement, and if it is correct, the statement express a conceptual truth, and the right-hand side both *expresses* the application conditions of the concept of knowledge and *states* what it is for someone to have knowledge. If [K] meets the conditions, we can use it to give a formal mode statement about the application conditions of the concept of knowledge in [KC].

[KC] For any x, for any p, the concept of knowledge is true of the pair $<x, p>$ iff (i) p is true, (ii) x believes p, (iii) p is evident to x, and (iv) there is no proposition q such that if q became evident to x, then p would no longer be evident to x.

For this to be a conceptual analysis, it also has to meet conditions (a) and (b). From it one can infer [K]. As it is a mistake to think that [K], because it is about knowledge, cannot be used to express a conceptual analysis, so it is a mistake to think that [KC], because it is about the concept of knowledge, does not give us information about what knowledge is.[7] As Quine put it in the parallel case of giving the truth conditions of sentences: "The truth predicate is a reminder that, despite a technical ascent to talk of sentences, our eye is on the world" (1986, 9).

What are the sources of the curious view that conceptual analysis does not tell us anything about what our concepts are of? I think there are a number of things that have exerted an influence, sometimes perhaps in conjunction with one another.

One may simply be the failure to recognize that in specifying the application conditions of a concept of C one thereby gives information that suffices to produce a material mode statement about necessary and sufficient conditions for being C.

A second may be the conflation of the psychological study of concepts with a concern for understanding concepts in their role in fixing the contents, and, hence, truth conditions, of thoughts. The former psychological project – whether it concerns how concepts are realized, or contingent laws involving them – gives no insight into how concepts fix the truth conditions of thoughts, and so would not be thought to be relevant to the natures of the things which the concepts pick out. The thought that psychologists are studying the very concepts for which we seek analyses may also encourage the thought that analysis should be an empirical enterprise.[8] But this is confused twice over. First in thinking that psychologists are in the business of analysis and second in thinking their techniques are relevant to it.[9]

A third source is the conflation of the concept of F with a mini-theory about Fs, a folk theory of sorts, which may not accurately characterize Fs, and might in fact be radically mistaken.[10] If one thinks this, then one will not think that an analysis of the concept of F *ipso facto* gives you knowledge of Fs. But this is a category mistake. The theory would itself have to involve beliefs that include the concept because that is what fixes its subject. Thus, any beliefs one has about Fs presuppose possession of the concept. One's possession of the concept therefore could not be explained by one's having the theory. What it is to possess a concept is to have the ability to apply it correctly in accordance with a rule. The mistake of the mini-theory of concept possession is to confuse a competence in the deployment of a concept with having beliefs involving it.

A fourth source of confusion stretches back to Kant in one form, the idea that analytic truths do not provide ampliative knowledge. A purer form is found in the *Tractatus Logico-Philosophicus* (Wittgenstein 1961), namely, the idea that analytic truths are non-factual. If all conceptual truths are expressed with analytic sentences, we get the conclusion that conceptual truths are non-factual. The doctrine's origins in the *Tractatus* are tied to Wittgenstein's picture theory of representation and truth-functional theory of propositions. Wittgenstein held that atomic propositions represented atomic states of affairs. Atomic propositions were independent of one another. We represent how the world is by atomic propositions and truth functions out of them. A proposition has a sense to the extent to which it locates us in logical space. The logical apparatus of truth-functional logic aids in sketching positions in logical space, but the mechanism allows for limiting cases in which the truth values of the sentences constructed out of the connectives are insensitive to those of the contained atomic sentences, and so always true (tautologous) or false (contradictory). These do not locate one in logical space, have no sense, and hence are devoid of factual content. Identifying analytic truths with tautologies entails analytic truths are non-factual. This idea was adopted by the logical positivists, and associated with analytic truths by Quine in his criticism of the analytic/synthetic distinction (Quine 1953). But the framework in the *Tractatus* that made sense of it was not adopted along with the doctrine, and no substitute has been offered, though the view itself has continued to exert an influence. This is an example *par excellence* of the persistence of a theoretical dogma after its support has been removed.

A fifth source is the view that concepts correlate with semantic competence, which should be construed on the model of competence in the use of natural kind terms like "gold" or "oak," which point us to real essences rather than providing us with a way of directly apprehending them. Though not all terms are natural kind terms, I agree that competence with natural kind terms does not put one in a position to give their essence. Nor, however, does it suffice to grasp the concept of the kind (in the sense of "concept" we've discussed), because that is fixed by the kind property it attributes (if any). As competence with the term doesn't fix the kind property, it doesn't fix the concept of the kind either. So this doesn't bear on the present issue. I take up the question whether epistemic terms are natural kind terms in the concluding section ("Is 'Knowledge' a Natural Kind Term?").

What is the role of thought experiments?

Conceptual analysis requires us not just to make judgments on the basis of conceptual competences but to reflect on how those judgments express their structure. We exercise conceptual competence *inter alia* when we make judgments on the basis of our beliefs and perceptual representations of the world and when we form and act on intentions in pursuing goals in the light of what we want most to get. Most of the judgments we make are not conceptual truths, however, because they are the products of antecedent beliefs about contingent features of the world and present experience. We could make a start on identifying necessary or sufficient application conditions for concepts by considering what features are invariably present when we deploy them, or reflecting on patterns of inference. But since we are interested in the patterns induced by competencies, it is more expedient to draw on those competencies directly by asking when it is correct to judge one or another thing, in the sense of its following from the description of the conditions—e.g. whether someone who infers justifiably from a justified but false belief something which is true thereby knows it.[11] It may help to describe a scenario involving hypothetical individuals as placeholders for things with relevant features and then to ask whether it is correct to make certain judgments about them. In this case, we engage in a thought experiment.

Though the scenarios involve hypothetical individuals, the import is general. An example will illustrate. In arguing that neighborhood reliabilism (NR) allows illegitimate bootstrapping, Jonathan Vogel describes the following thought experiment (2000, 353–56): Roxanne, who drives a car with a reliable gas gauge, believes what it reports, but she doesn't know that the gas gauge is reliable. However, she often notes how much gas is in the tank and what the gauge reads. The perceptual process by which she comes to believe what the gauge reads is reliable, as is the process by which she comes to believe how much gas is in the tank. She infers on such occasions that the gauge reads N and the tank is N. As deduction is a reliable process, her beliefs on these occasions are reliably produced. Roxanne concludes by induction (again a reliable process) that the gas gauge is always accurate, and then concludes that the gauge is reliable. Question: Does Roxanne know in virtue of the process she follows in coming to believe that the gauge is reliable that the gauge is reliable? Vogel says:

> I assume that bootstrapping is illegitimate. Roxanne cannot establish that the gas gauge is reliable by the peculiar reasoning I have just described. The challenge to NR is that it may go wrong here. On the face of things, it does improperly ratify bootstrapping as a way of gaining knowledge.
>
> (2000, 354)

Here the scenario functions as a schematic description of a type of case, and it is the type of case that we are to reason about. Hence, the question could be recast in terms of a universally quantified conditional, "Is it the case that for any x, if x ... , then x knows that the gas gauge is reliable?" where the intention is that one should answer on the basis of whether the antecedent states a condition conceptually sufficient for the consequent.[12] Thus, thought experiments draw on our ability to tell whether one proposition follows from another.

A creature can have concepts without the concept of a concept. But conceptual analysis requires the concept of a concept, and in particular the concept of conceptual entailment (what follows from what). Since having these concepts puts us in a position to deploy them correctly, this is both necessary and sufficient for the possibility of conceptual analysis.[13]

Thought experiments have played, and continue to play, a significant role in epistemological theorizing. Usually thought experiments do not establish outright an analysis, but rather provide starting points and test cases. We begin with observations that on the face of it express conceptual truths, such as that knowing p requires believing p and p being true. This being insufficient, we may propose that knowledge is e.g., justified true belief. Typically with concepts situated in a family of concepts related in diverse ways with others, it is not immediately obvious that the conditions proposed are correct. It need be no more obvious than it is obvious right off the bat that there is no greatest prime number. We can then test the proposal against judgments in cases where we have filled in details that may be thought to be relevant to whether the conditions are sufficient. The process has the familiar pattern of observation, hypothesis, prediction, and test, with the role of observation played by judgments with respect to scenarios in thought experiments, that is, judgments about entailment relations.

Thus, while the judgments we make, if properly based, are expressions of our competence, and are thus a priori, the analysis is often based on a form of inference to the best explanation. We may be able, once the proposal is formulated, to see directly that it is correct. In many cases, however, confidence that we have a correct analysis rests on the claim that we have surveyed representative cases and not overlooked anything important. In these cases, the justification we have is in part a posteriori (see Henderson and Horgan 2001). In this respect, though, philosophy seems no worse off than mathematics. Confidence that Peano's axioms

axiomatize the natural numbers rests in part on their entailing classical results.

Problems in conducting thought experiments arise from at least three sources:

Problems of design. A well-designed thought experiment has (a) a clearly characterized target proposition, (b) a clear, unambiguously described scenario, which constitutes a test case, and which is complete in the respects relevant to the test, and (c) a clear, unambiguous test question relevant to the target proposition (see Ludwig 2007, §2, for an example that fails these conditions).

Problems of execution. The subject should have a clear understanding of (a) the purpose of the thought experiment, (b) the scenario, (c) the questions, and (d) the possible responses, including the response that insufficient information is given or that a presupposition of the thought experiment is not met, and (e) the answer to the question understood literally should be based solely on conceptual competence.

Problems of presupposition. A thought experiment, perforce couched in language, presupposes that the words used express concepts, and that competence in their use involves grasp of the concepts they express. This presupposition fails for natural kind terms (see the concluding section, "Is 'Knowledge' a Natural Kind Term?") and for words that are semantically defective, as in the case of vagueness and the semantic paradoxes (Ludwig and Ray 2002).

We can check judgments by how they fit with other cases, with judgments by others (who are good "observers"), by fit with well-developed accounts in surrounding areas, and by various theoretical considerations. Withdrawing a judgment, we seek to explain the mistake and re-evaluate the case so as to see it in a different light. This is a form of the method of reflective equilibrium.

Reflective equilibrium is criticized sometimes for providing only internal justification (Stich 1988), and so being subject to the objection that there can be conflicting but equally coherent sets of judgments. But the suggestion is not that coherence makes for justification. The assumption underlying the method of reflective equilibrium in conceptual analysis is that most of the judgments we make under optimal conditions will be correct. We employ similar methods to correct mistakes about memory: on the assumption that we remember mostly accurately, we test cases by how well they cohere with the rest and with general knowledge; where our own resources give out we can appeal to others. Why accept the assumption in the case of thought experiments? First, to think about a subject matter requires being competent in the deployment of the relevant concepts. This guarantees

that we are in a position to make correct judgments about scenarios in thought experiments. Second, while it is consistent with this that in practice we typically fail to express our competences, there is (a) no special reason to think that this is so and (b) the supposition that we get it wrong generally throws all inquiry into doubt, since it undermines confidence that we can assess evidence for any hypothesis, including the skeptical one.

I turn to two objections to the traditional use of thought experiments in epistemology. The first objection challenges the reliability of relying on one's own conceptual competence. Here I distinguish two charges. The more radical is that "intuitions" are relative to such things as cultural or socio-economic background, and that in consequence "it is wrong for philosophers to assume a priori the universal validity of their own ... intuitions" (Machery *et al.* 2004, B8). The less radical is that given individual fallibility, a more reliable method of tracking correct responses is to take up the survey methods of the social sciences. The second objection is that most of the terms of interest in epistemology are natural kind terms, so that thought experiments at best reveal e.g., the stereotype of knowledge and not its essence.

What are the lessons of experimental philosophy?

I cannot discuss in detail the now large literature surrounding experimental philosophy. But I wish to urge two points. First, surveys of undergraduate responses to thought experiments go very little way toward calling into question philosophical practice. Second, coming to a view from one's own perspective is ultimately necessary to assess such surveys and also essential to the aim of philosophy.

I begin with the charge that "cognitive diversity" across cultures or socio-economic groups undermines the probative value of thought experiments. Among the many papers in the genre I will focus for illustration on one that has attained the status of a classic in the field, "Epistemic Intuitions and Normativity" by Weinberg *et al.* (2001).[14]

Weinberg *et al.* characterize an epistemic intuition as "a spontaneous judgment about the epistemic properties of some specific case—a judgment for which the person making the judgment may be able to offer no plausible justification" (2001, 19). They then adduce evidence to show that "epistemic intuition" so characterized varies, for example, across cultural groups. The evidence consists in survey data of undergraduates responding to probes involving thought experiments that have appeared in the philosophical literature such as Gettier cases, Dretske's zebra case (1970, 1015–16), and Lehrer's Trutemp case (2000, 187). In the case of some probes, there were differences in majority responses across Westerners,

East Asians, and students from the Indian subcontinent. For example, in the case of a probe about a "Gettier case" involving Bob's thinking Jill owns a Buick and so an American car, though she recently replaced it with a Pontiac, students were asked whether Bob really knows or only believes Jill drives an American car. Seventy-four percent of Westerners said Bob only believes Jill drives an American car, whereas 57 percent of East Asians and 61 percent of the students from the Indian subcontinent said Bob really knows that she does. Thus, it looks as if "epistemic intuitions" and therefore concepts of knowledge must differ across cultures.

An initial mistake here confuses the issues. "Epistemic intuition" is given a stipulative definition as a spontaneous judgment about epistemic properties for which the person making it may be able to offer no plausible justification. But the method of thought experiments calls on us to respond on the basis solely of our understanding of the scenario and the question asked about it. Given this, and that concepts individuate the thoughts they are involved in, it is clear that there could be no relativity of the target response to cultural background or anything else, and no sense in which the concept of knowledge could differ across cultures. The most that relativity of response to cultural background or other factors could show (assuming shared concepts) is that there are errors traceable to something connected with those differences. This need not involve errors in the application of concepts, for it may involve errors in understanding the task, or differences in how unarticulated details are filled in, affected by different background assumptions, or ways of taking a question or statement. In the probe involving Bob and Jill, for example, if one assumes that most Americans who own an American car buy American cars generally, one might think Bob knows this also and so is justified in believing that Jill drives an American car independently of being justified in believing that she is now driving a Buick.[15]

But so what? So what if the varied responses of undergraduates to these questions can't all be taken to be judgments based on conceptual competence in response to the scenario, task, and questions, properly understood? Isn't the problem now that this just shows that *none of us* are very good at saying when one thing follows from another?

How could it show this? We know, after all, the correct answer in Gettier cases (properly described). The results show students, even a majority in some cases, can make mistakes, but we knew that. Does every mistake on a homework assignment in logic shake the foundations of the subject? No, not even if every student makes the same mistake. Analysis is a cognitive skill. It can be inculcated. It draws on basic shared competencies. But that doesn't entail everyone is equally good at it, or good at it right off the bat. Students are often not very good at recognizing

deductive validity, or in solving math problems, or elementary probabilistic reasoning. But many get better at it. They don't acquire new concepts, but get better at exercising them systematically, and beyond the usual range of cases they confront, and at a host of other related cognitive skills.[16]

What would show trouble? It can't be a general skepticism about our ability to recognize when one thing follows from another. That undermines all inquiry, including inquiry aimed at casting doubt on our abilities. Doubts about thought experiments in philosophy have to focus on something specific to them. But there is no special reason to think that we cannot identify entailment relations when we put a query about an entailment in the form of a thought experiment.[17] While as in any investigation there are methodological problems, we have also developed sophisticated tools for dealing with them (work on logical form and conversational pragmatics for example). Thought experiments should be approached with those tools in hand, and placed in the context of other thought experiments as well as theoretical considerations, both from within the field and from other fields. (For a case study, see the discussion of trying in Ludwig 2007, 145–46; see also Cullen 2010, for discussion of pitfalls in conducting surveys involving thought experiments.)

I turn now from these general skeptical concerns to the thought that, perhaps precisely because of the problem of identifying probative responses, we should move away from a first-person methodology to a third-person methodology.

It would be a mistake to dismiss how most people respond to surveys as completely irrelevant. One might even be encouraged by reflection on the Condorcet jury theorem to think that with enough participants, the probability that the majority is right becomes greater than that any individual is right. But as already noted, we can't take it for granted either. The Condorcet jury theorem assumes that everyone in the relevant class has a positive bias toward truth on the matter in hand and makes independent judgments (Dietrich and List 2004). The prevalence of the gambler's fallacy, as well as the mistakes students make on surveys about Gettier cases, shows this does not always obtain. We cannot assume most people are good at drawing relevant distinctions, or have the facility for the type of thinking involved. We cannot assume that most people are armed against the various pitfalls in conducting thought experiments. And for untutored or untrained subjects, we can't assume task understanding, even when we think we have explained the task clearly.

Ultimately, to assess whether the majority response on a survey is correct, we need to have an independent view of the matter. We need to have insight ourselves into the correct answer. That is how we detect the fallacy in the gambler's fallacy. Furthermore, this independent view of the right

response is exactly the kind of understanding that we seek in philosophy. Just as our interest in mathematics would not be served by knowing merely that a certain theorem was true because a majority of mathematicians endorsed it, so our interests in philosophy would not be served merely by knowing that a certain claim is true because most people endorse it. We want to see why it is true, to understand it ourselves. This is the most fundamental reason why surveys cannot replace the first-person approach to conceptual analysis.[18]

Is "knowledge" a natural kind term?

The general question facing us is whether the family of epistemic terms such as "knowledge," "truth," "evidence," "justification," "warrant," and the like, are all natural kind terms like "gold," "air," "water," "tiger," "gene," etc., discovery of whose real essences requires empirical investigation, or whether they are like "number," "circle," "cylinder," "necessity," and "logical consequence," whose essence is revealed in reflecting on our grasp of the application conditions of the concepts they express. It is beyond the scope of this chapter to take up the question with respect to all epistemic terms. I will restrict attention to "knowledge," which may be the best case for the hypothesis.

What are natural kind terms?

Let's take "gold" as our example. If (1) is true, then (1N) is true, but (1) is also, it seems, an empirical discovery. Thus, it seems to be both necessary and a posteriori.

(1) Gold is an element with atomic number 79.

(1N) Necessarily, gold is an element with atomic number 79.

How do the characteristic features of natural kind terms give rise to the view that (1), if true, expresses a necessary a posteriori truth?

We should distinguish natural kinds from natural kind terms. A natural kind we may take to be a stable explanatory kind relative to some range of phenomena. A natural kind term like "gold" is embedded in a practice which treats its purpose as that of "picking out" a natural kind, but not by way of our having been given the relevant property. The practice doesn't directly give us the kind but rather (a) involves a basis for the application of the term to objects or phenomena, (b) an explanatory relation the kind is to bear to the basis of application in at least most of the objects we apply it to, and (c) thereby a mode of identifying the kind

the term is to pick out, as, roughly, the kind, if any, which provides the best explanation of the relevant sort for the basis of its application in most actual cases (the kind property), and, possibly, (d) a default option for what the term picks out if there is no kind which explains the basis of its application in most actual cases (the default property). Competence in the use of the term, in the sense of counting socially as having mastered the practice, amounts to learning, along with its grammatical category, its role as picking out a natural kind, the basis of application, and the explanatory relation it is to bear to (most of) the instances picked out on its basis of application, and the default option. For "gold," the basis of application includes being (in typical circumstances of application) a malleable incorruptible yellow metal that dissolves in aqua regia. The intended explanatory relevance relation is something on the order of constitutive explanation, and for a kind to bear this relation to the basis of application of the term "gold" is for it to explain in virtue of the structure of the items to which the term is applied those features of it which constitute its basis of application. (Details won't matter so much as the form of the account.)

"Gold" is a mass noun like "snow" or "flesh" or "garbage." In logical form, mass nouns contribute predicates (Koslicki 1999). Thus, I represent (1) as having the form: For all x such that x is gold, x is an element with atomic number 79. The predicate "is gold" then is used to attribute a property to an object. But the property it is intended to attribute is not given by the basis of application, but rather is the relevant kind property, if any, and otherwise the default property. Suppose that the relevant kind property for the term "gold" is being an element with atomic number 79. Then: the truth conditions for "For all x such that x is gold, x is an element with atomic number 79" are given by "For all x such that x is an element with atomic number 79, x is an element with atomic number 79," and that is necessarily true. However, it is clear that empirical inquiry is required in order to identify the relevant kind. It therefore appears that "Gold is an element with atomic number 79" expresses a necessary, a posteriori truth.

What proposition is expressed by "Gold is an element with atomic number 79"? If we mean what determines the truth conditions for (1), "gold" contributes, not anything having to do with its basis, but instead the property it attributes. If that is being an element with atomic number 79, then the proposition is that anything that is an element with atomic number 79 is an element with atomic number 79. That proposition is a conceptual truth—indeed, it is true in virtue of its structure alone. A proposition is a priori if someone who grasps it is in a position to judge it correctly. By this standard, the proposition expressed by (1) is an a priori conceptual truth. And as the property attributed by "gold" is fixed by

what kind explains its basis of application, so is the concept expressed by it, in the sense of "concept" on which concepts are individuated by their contributions to the truth conditions of thought contents. (1) then turns out not to be a counterexample to only a priori truths being necessary. It expresses a logical and hence conceptual truth. The necessity involved in (1N) is old and familiar, not new and exotic.

Of course something here requires empirical investigation: what we didn't know is that the stuff to which we apply "gold" properly on its basis of application is for the most part something the fundamental constituent property of which is being an element with atomic number 79. That is, we didn't know prior to investigation what property "gold" attributes (or what concept, in the relevant sense, it expresses). And this is the same as saying that what required empirical work was discovery of what proposition "Gold is an element with atomic number 79" expresses.

This is no help, however, with the methodological challenge to epistemology. For it is clear that given the way "gold" is introduced, we cannot discover what (what we call) gold is without empirical investigation. Therefore, to the extent to which a domain of discourse that attracts philosophical interest traffics in natural kind terms, to that extent also traditional methods that presuppose grasp of the thoughts expressed by sentences in the domain of discourse are inapplicable. I turn now to the question whether "knowledge" in particular is a natural kind term.

Is "knowledge" a natural kind term?

One might be tempted to argue that "knowledge" is a natural kind term because knowledge is a natural kind: "There is a robust phenomenon of human knowledge and a presupposition of the field of epistemology is that cases of knowledge have a good deal of theoretical unity to them; they are not merely some gerrymandered kind, united by nothing more than our willingness to regard them as a kind" (Kornblith 2002, 10). From this we might pass to the thought that since "[u]nderstanding what that theoretical unity is is the object of our study ... it is to be found by careful examination of the phenomenon, that is, something outside of us, not our concept of the phenomenon, something inside us. In short, ... the investigation of knowledge, and philosophical investigation generally [should be pursued] on the model of investigations of natural kinds" (11). The suggestion that we should not examine our concept of knowledge but the phenomenon itself suggests that the claim that "the investigation of knowledge" should be pursued "on the model of investigations of natural kinds" is to be construed as the claim that we should treat it as on the model of the investigation into the natures of things we picked out with natural kind terms. Thus, it is natural to take the intent here to be

expressible in the following argument (even if not Kornblith's intent, this will clear the ground).

(1) For any kind *K* expressible using a term *T*, if kind *K* is a natural kind, then term *T* is a natural kind term.

(2) Knowledge is a natural kind.

(3) Knowledge is expressible using "knowledge."

(4) Therefore, "knowledge" is a natural kind term.

But the argument is unsound because the first premise is false. "Gold" is a natural kind term, and gold is a natural kind, namely, the element with atomic number 79. However, "element with atomic number 79," which picks out the same natural kind, is not a natural kind term. Thus, it doesn't follow from knowledge being a natural kind that "knowledge" is a natural kind term. And it doesn't follow from something's being a natural kind that we don't possess the concept of it in the sense that puts us in a position to specify what its essential nature is (to the extent possible) by analysis.[19]

We might at this point entertain other indirect arguments. But the claim that "knowledge" is a natural kind term is a claim about its use in the language. So let us instead ask directly how to test whether a term is a natural kind term.

A good test is provided by the kind of thought experiment Putnam used to bring out what's special about our practice with respect to terms like "water" and "gold" in the first place (1975). A hallmark of a natural kind term is that what property it attributes is fixed by what explanatory kind if any actually explains the features that constitute its basis of application in the samples to which we apply it. Fixing the practice, if the underlying kind that explains the features had been different, then the property attributed, and the propositions expressed by sentences containing it, would have been different. And if there were no underlying kind that explained the basis of its application, then either it would track a default property if provision is made for it, or not attribute any property at all, and so no sentence containing the term in a use position would express a proposition. We can then test the hypothesis by considering a circumstance in which a community of individuals associates with "knowledge" the same linguistic practices as we do, but where there are salient differences in the states they pick out which suffice for us to judge them not to constitute knowledge, though those states explain the basis for their application of "knowledge" to them. We then ask whether when they call those states "knowledge" they are speaking truly, as the hypothesis predicts.

Consider a possible circumstance in which we have doppelgängers who are bodies in vats (BIVs) in the style, say, of the 1999 film *The Matrix* by the Wachowski brothers. A supercomputer tracks outputs and regulates inputs to their brains (or bodies if you like), and generates coordinated experiences so as to take into account outputs from brains in determining, relative to the plan of the fictitious world they are presented with, inputs to other BIVs. For example, when it seems to my doppelgänger that another is speaking or moving, that is because the other's motor cortex is firing in a way that would be appropriate for that, and that in turn, through the mechanism of the supercomputer, generates inputs of the sort that generate in my doppelgänger experiences of the requisite sort for a body moving in the appropriate ways. The BIVs are therefore partly causally responsible for the co-evolution of their experiences. Their dispositions with respect to language use are to be exactly the same as ours, as is, modulo references to the self and time, the course of their experiences. They can, in a fairly straightforward sense, carry on "conversations" about "what goes on around them" (we can even imagine they move their lips and utter sounds, though otherwise immobile in their vats). I stipulate that we know that the actual world is not like this. The point is not to raise any skeptical worry about our own knowledge, but to test a linguistic hypothesis. *We* will judge (correctly) that *they* would not know very much at all, since most of their beliefs are false.[20] The question to focus on, however, is whether, for example, when my BIV doppelgänger says or thinks "Jones knows a lot about Volkswagens" he is (a) speaking the truth, (b) speaking falsely, or (c) not expressing a proposition at all. The hypothesis that "knowledge" is a natural kind term would seem to predict either (a) or (c).[21]

The states (in themselves and others) that most of their uses of "knowledge" would track would be states that are reliably connected, not with what they are about, but with features of the supercomputer that realize its model of the illusory world which their experiences represent to them (and perhaps certain other features corresponding to internalist constraints on knowledge—throw in whatever else seems relevant). For convenience, let us sum up these features pertaining to the proposition that p as the illusion that p. Now we can ask: when my BIV doppelgänger says "Jones knows a lot about Volkswagens" does he speak truly in his language? Is the proposition expressed by his sentence expressed in our language by "Jones is in a belief state that is caused by a process that reliably produces belief that p in circumstances in which Jones is presented with the illusion that p, and ... "? Alternatively, does he express no proposition at all?

On the face of it, neither of these suggestions has any plausibility. Were it us in that situation, and we were to wake up, as it were, and find we could leave the vats, and learned of what had happened, and were instructed in what our use of "knowledge" had actually tracked, what would we think or say? Would we say (in the language that we would in those circumstances speak): "We no longer have any knowledge but we used to"? Would we say: we were speaking neither truly nor falsely when we claimed such things as "Many people know where they live"? Or would we say: "We did not previously have any knowledge but thankfully we now do"? If we would say the latter, even given that we believe that "knowledge" formerly tracked states that were reliably produced by illusions, then we would be expressing allegiance to a practice that is incompatible with treating "knowledge" as a natural kind term. It is very clear that we would in fact say the latter, and it is very clear that this is exactly what our BIV doppelgängers (who have our linguistic dispositions) would say. But this is not what the hypothesis that "knowledge" is a natural kind term predicts.

I will leave it as an exercise to the reader to construct further test cases, but the preliminary result is that whether or not knowledge is a natural kind, "knowledge" is not a natural kind term. I dare say this is not much of a surprise. It was not to be expected that epistemic concepts, which must form the framework for our thinking about the rational investigation of any subject matter, should themselves be outsourced to the world that we are investigating.

Notes

1 I have in mind truth conditions in the sense in which we say that "p" in "s is true iff p" gives the truth conditions of s iff "p" translates s.

2 This is intended to capture the core of a traditional view that in one form goes back at least to Frege. See Peacocke 1992 for one presentation of the general form of the idea, and Chalmers and Jackson 2001 for another.

3 Most of Williamson's arguments against epistemic conceptions of analyticity in 2007, ch. 4, aim to show that grasp of conceptual truths is insufficient for assent or disposition to assent. I don't know that anyone ever maintained otherwise. He considers only one proposal for an epistemic approach in the last pages of the chapter, and gives a schematic and unconvincing argument against it, but in any case it is not the proposal advanced here.

4 Goldman suggests competence views may be alright in theory but little help in practice (2010, 135), but the truth is that we have made a lot of progress in conceptual clarification in non-ideal conditions.

5 Contrary to what is sometimes suggested (Stich 1988; Ramsey 1998), there is nothing in so-called prototype concepts per se incompatible with the tradition in analysis. Goldman makes this point (2007, 23).

6 I elide discussion of what must be an integral part of the overall project when we confront the fact that our access to the structures of thoughts relies on the analysis of the structure of the sentences we use to express them, namely, that a first step is analysis of the logical form of the domains of discourse we are interested in.

7 See Ludwig 2007, 131; McGinn 2012, ch. 5. Semantic descent is not a recent discovery.

8 Kornblith writes, for example: "If concepts are psychologically real, and also ... there is a well established tradition in experimental psychology that studies them, then what room is left for the armchair methods of philosophers, methods designed to illuminate the very same target?" (2007, 30).

9 Except when it takes the form of statistical studies of patterns of application of concepts to things in the context of subjects seeking to say what is true. This takes us in the direction of experimental philosophy, which I discuss below ("What Are the Lessons of Experimental Philosophy?").

10 Both Ramsey (1998) and Cummins (1998) seem to make this mistake. I detect this thought in Kornblith (2007, 37) as well.

11 This answers the paradox of analysis. Grasp of concepts is *competence in correct deployment*. An analysis is a proposition we come to know *on the basis of exercising competence in the concepts deployed* in it. See McGinn 2012, ch. 4, and Strawson 1992, 5–13, for essentially the same account; a similar line can be found in Fumerton 1983, and earlier in Myers 1971.

12 I have avoided the word "intuition." While I have a position (Ludwig 2010), given the dust raised by extensive debate about the word, it seems best avoided in favor of an independent characterization of what we are interested in. See Nagel 2007 for a historical review. Recently, Herman Cappelen has argued through case studies that philosophers do not rely on intuitions, including in this judgments based on conceptual competence (Cappelen 2012). In my view, Cappelen is looking too hard for metaphilosophical remarks in philosophers' texts. Though Vogel in the passage quoted doesn't talk about intuitions or conceptual analysis, the fact is the article is part of an ongoing discussion understood to be concerned with the analysis of (the concept of) knowledge.

13 Perhaps this itself has an air of illegitimate bootstrapping, for this reasoning is itself presented as resting on our understanding of the contained concepts, and there seems to be no higher court of appeal (Cummins 1998). But this is true of every fundamental source of justified belief (Goldman 2007, 5).

14 In the same vein see Nichols *et al.* 2003; Alexander and Weinberg 2007; Swain *et al.* 2008. See Ludwig 2010 for further discussion of the 2001 paper.

15 I draw attention to this (rather obvious) point in my 2010; Sosa (2008) made the point earlier.

16 One response is that it is an empirical matter whether philosophers theorize better with the aid of thought experiments than undergraduates (Weinberg *et al.* 2010). Yes, but survey results give us no reason to think otherwise. See Williamson 2011. In any case, resolution of the challenge presupposes we can come to a correct view. But who are we going ask about this? In this connection, see the next note. For a response that draws on empirical work, see Nagel 2012.

17 Weinberg (2007) argues, not that there is special reason to doubt the reliability of the method of thought experiments, but rather that it is defective because not open to independent error correction. For a response, see Grundmann 2010.

18 On a different view, experimental philosophy aims merely "to provide an account of the factors that influence applications of a concept, and in particular, the internal psychological processes that underlie such applications" (Knobe and Nichols 2008, 5). On this view, experimental philosophy doesn't aim to engage in or undermine conceptual analysis, but to help identify pitfalls in conducting thought experiments. It is hard to see any objection to or in this.

19 Goldman has suggested that it is incompatible with philosophical practice that knowledge is a natural kind (2007, 8; 2005). But the practice doesn't rule it out. A natural kind we pick out by a natural kind term doesn't cease to be one when we discover the kind property—but when we do, we grasp its concept. The standard practice seems incompatible with "knowledge" being a natural kind term, which is evidence against the claim. But I will suggest we can test the hypothesis more directly below.

20 I don't think that a form of externalism about thought content that undermines the description of the scenario is correct. But, in any case, the hypothesis should predict a result even relative to the hypothesis that externalism about thought content is false.

21 Perhaps one could appeal to a default property, but it is a mystery what it could be.

References

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

Burge, Tyler (1979) "Individualism and the Mental," *Midwest Studies in Philosophy* 4: 73–121.

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

Chalmers, David and Frank Jackson (2001) "Conceptual Analysis and Reductive Explanation," *Philosophical Review* 110(3): 315–60.

Cullen, Simon (2010) "Survey-Driven Romanticism," *Review of Philosophy and Psychology* 1(2): 275–96.

Cummins, Robert (1998) "Reflection on Reflective Equilibrium," in M. DePaul and W. Ramsey (eds), *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, New York: Rowman & Littlefield.

Davidson, Donald (2001) "Radical Interpretation," in *Inquiries into Truth and Interpretation*, New York: Clarendon Press. Original edition, 1973.

Dietrich, Franz and Christian List (2004) "A Model of Jury Decisions Where All Jurors Have the Same Evidence," *Synthese* 142: 175–202.

Dretske, Fred (1970) "Epistemic Operators," *Journal of Philosophy* 67(24): 1007–23.

Fumerton, Richard A. (1983) "The Paradox of Analysis," *Philosophy and Phenomenological Research* 43(4): 477–97.

Goldman, Alan (2005) "Kornblith's Naturalistic Epistemology," *Philosophy and Phenomenological Research* 71: 403–10.

——(2007) "Philosophical Intuitions," *Grazer Philosophische Studien* 74: 1–26.

——(2010) "Philosophical Naturalism and Intuitional Methodology," *Proceedings of the American Philosophical Association* 84(2): 115–50.

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

Henderson, David and Terry Horgan (2001) "The A Priori Isn't All It Is Cracked Up to Be, But It Is Something," *Philosophical Topics* 29: 219–50.

Klein, Peter (1971) "A Proposed Definition of Propositional Knowledge," *Journal of Philosophy* 66(16): 471–82.

Knobe, Joshua and Aruda Burra (2006) "Experimental Philosophy and Folk Concepts: Methodological Considerations," *Journal of Cognition and Culture* 6 (1–2): 331–42.

Knobe, Joshua Michael and Shaun Nichols (2008) *Experimental Philosophy*, Oxford: Oxford University Press.

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

——(2007) "Naturalism and Intuitions," *Grazer Philosophische Studien* 74: 27–49.

Koslicki, Kathrin (1999) "The Semantics of Mass-Predicates," *Noûs* 33: 46–91.

Lehrer, Keith (2000) *Theory of Knowledge*, 2nd edn, Dimensions of Philosophy, Boulder, CO: Westview Press.

Ludwig, Kirk (2007) "The Epistemology of Thought Experiments: First Person versus Third Person Approaches," *Midwest Studies in Philosophy* 31: 128–59.

——(2010) "Intuitions and Relativity," *Philosophical Psychology* 23(4): 427–45.

Ludwig, Kirk and Greg Ray (2002) "Vagueness and the Sorites Paradox," *Philosophical Perspectives* 36(16): 419–46.
Machery, E., R. Mallon, S. Nichols and S. Stich (2004) "Semantics, Cross-cultural Style," *Cognition* 92: B1–B12.
McGinn, Colin (2012) *Truth by Analysis: Games, Names and Philosophy*, Oxford: Oxford University Press.
Myers, C. Mason (1971) "Moore's Paradox of Analysis," *Metaphilosophy* 2(4): 295–308.
Nagel, Jennifer (2007) "Epistemic Intuitions," *Philosophy Compass* 2(6): 792–819.
——(2012) "Intuitions and Experiments: A Defence of the Case Method in Epistemology," *Philosophy and Phenomenological Research* 85(3): 495–527.
Nichols, Shaun, Stephen Stich and Jonathan M. Weinberg (2003) "Meta-Skepticism: Meditations in Ethno-epistemology," in S. Luper (ed.), *The Skeptics: Contemporary Essays*, Aldershot: Ashgate.
Peacocke, Christopher (1992) *A Study of Concepts*, Cambridge, MA: MIT Press.
Putnam, Hilary (1975) "The Meaning of 'Meaning,'" in *Mind, Language and Reality: Philosophical Papers*, Cambridge: Cambridge University Press.
Quine, W.V. (1986) *Philosophy of Logic*, 2nd edn, Cambridge, MA: Harvard University Press.
——(1953) "Two Dogmas of Empiricism," in *From a Logical Point of View*, Cambridge, MA: Harvard University Press.
Ramsey, William (1998) "Prototypes and Conceptual Analysis," in M. DePaul and W. Ramsey (eds), *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, New York: Rowman & Littlefield.
Sosa, Ernest (2008) "A Defense of the Use of Intuitions in Philosophy," in M. Bishop and D. Murphy (eds), *Stich and His Critics*, Oxford: Blackwell.
Stich, S. (1988) "Reflective Equilibrium, Analytic Epistemology and the Problem of Cognitive Diversity," *Synthese* 74(3): 391–413.
Strawson, Peter (1992) *Analysis and Metaphysics: An Introduction to Philosophy*, New York: Oxford University Press.
Swain, Stacy, Joshua Alexander and Jonathan Weinberg (2008) "The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp," *Philosophy and Phenomenological Research* 76: 138–55.
Vogel, Jonathan (2000) "Reliabilism Leveled," *Journal of Philosophy* 97(11): 602–23.
Weinberg, Jonathan (2007) "How to Challenge Intuitions Empirically without Risking Skepticism," *Midwest Studies in Philosophy* 31: 318–43.
Weinberg, Jonathan, Chad Gonnerman, Cameron Buckner and Joshua Alexander (2010) "Are Philosophers Expert Intuiters?," *Philosophical Psychology* 23(3): 331–55.
Weinberg, Jonathan, Shaun Nichols and Stephen Stich (2001) "Normativity and Epistemic Intuitions," *Philosophical Topics* 29: 429–60.

Williamson, Timothy (2007) *The Philosophy of Philosophy*, Malden, MA: Blackwell.
——(2011) "Philosophical Expertise and the Burden of Proof," *Metaphilosophy* 42(3): 215–29.
Wittgenstein, Ludwig (1961) *Tractatus Logico-Philosophicus*, ed. D.F. Pears and B.F. McGuinness, London: Routledge & Kegan Paul. Original edn, 1921.