

Gettier and the Program of Analysis

PATRICK RYSIEW

1. Introduction

A central task of mid-20th century epistemology was ‘the analysis of knowledge’, understood as an attempt to state necessary and sufficient conditions for the correct application of the concept. Gettier’s (1963) paper marked a crucial moment in this tradition’s history: on the usual telling (see note 5), it awoke ‘justified true belief’ (JTB) theorists from their dogmatic slumbers, showing the received analysis of knowledge to be inadequate. Just as importantly, for many, the subsequent (perceived) failure of attempts to solve the problem, with ever more complex analyses succumbing to counter-example, called into question the value and/or feasibility of analyzing knowledge. The Gettier problem has thus prompted important methodological questions about the epistemological enterprise. Most obviously, is it realistic to think that providing an analysis – or ‘definition’ -- of knowledge, understood along classical lines, is worthwhile, or that it’s even possible? If, as many now seem to think, it is not, why not? And -- just as pressing -- what might we put in its place? This chapter reviews some of the most prominent answers to such questions. Following a brief discussion of the aim and method of conceptual analysis in its classical form (Section 2), we will review (Section 3) some possible reasons for scepticism concerning either its prospects or value, at least as applied to knowledge. We will then turn (Section 4) to a consideration of some suggestions as to a plausible successor to the program of classical analysis, and some further potential reasons for pessimism about the latter. As we’ll see, some of the latter proposals are quite revisionary, and others less so, with little consensus and much lively debate about the appropriate methodological moral to draw from the Gettier problem.

2. Classical analysis

The search for classical analyses or ‘definitions’, including of central epistemic subjects, is at least as old as Plato; and, as has been said, “[t]he rules of the game have changed very little over the last 2500 years” (Stich 1992: 247). In the *Theatetus*, for example, in response to Socrates’ query, “What is knowledge?”, Theatetus tries out a number of answers, with it being pointed out in each case that there are examples running counter to the proposal. Such counter-instances lead to refinements of the proposed definition, which is then subjected to further testing.

This familiar instance illustrates the typical goal and method of conceptual analysis, as classically conceived. The goal is to identify the features that all and only cases of the target have in common. That is, where knowledge is our *analysandum* or *definiendum* (the thing we want to analyse or define), in attempting a *classical analysis* we are seeking a satisfactory completion of the following schema,

S knows that p if and only if (‘iff’) _____,

where what appears on the right-hand side (in the *analysans* or *definiens*) are conditions that are individually necessary (any instance of knowledge must satisfy them) and jointly sufficient (when they are all satisfied, the subject knows). An analysis fails when it is

found that there is a case of knowledge in which one or more of the supposedly necessary conditions does not obtain (in which case the account is too ‘narrow’ or ‘strong’), or an instance of the supposedly jointly sufficient conditions obtaining that we wouldn’t count as knowledge (in which case, the account is too ‘broad’ or ‘weak’).

As to method, the type of investigation we’re considering is usually taken to be *a priori*, not empirical. This is connected with the fact that the investigation is in some good sense conceptual: we are trying to find necessary and sufficient conditions for the correct application of the *concept* of knowledge; otherwise put, we are trying to identify necessary and sufficient conditions for someone’s knowing that *p* as determined solely by the *concept* in question, and the biconditional as a whole is meant to express a conceptually necessary truth. This, in turn, is why the method permits the use of merely hypothetical, and sometimes rather ‘far out’, cases: the conditions we are seeking are intended to be conceptually necessary and sufficient; and a merely possible case, however unrealistic, is enough to refute a claim of conceptual necessity or sufficiency (Lycan 2006: 151).

Part of the appeal of demanding a traditional analysis is that it promotes the ideals of clarity and rigor to which analytic philosophy aspires, and encourages the testing of theories. But we also want to increase our understanding of the relevant phenomena. Thus, it is because it is wholly uninformative that ‘*S* knows that *p* iff *S* knows that *p*’, while perfectly clear and testable, and while expressing a necessary truth, is a poor analysis (Neta 2002: 663; Lehrer 1990: 6-8; Ichikawa and Steup 2016: Introduction). (It is circular, of course, but its utter uninformativeness seems to be the more fundamental failing.) So the extensional equivalence of *analysandum* and *analysans* across possible worlds, and hence a necessarily true biconditional, doesn’t suffice for a good analysis.¹ Different suggestions have been made about the kind of information a good analysis provides – e.g., an account of how ascriptions of knowledge are made true by other facts (Jackson 1998: 28-30); an explanation of “how a person knows that her information is correct and how her knowledge claims are justified” (Lehrer 1990: 8); guidance in the application of epistemic terms to particular cases; or help in solving certain puzzles concerning knowledge (e.g., Neta 2002). But there is general consensus in prohibiting circularity, vagueness, *ad hoc* or ‘merely negative’ conditions, and the appeal to items more obscure or complex than what we’re trying to understand (Zagzebski 1999: 98; Earl n.d.: Section 3c).

As others have noted (e.g., Ichikawa and Steup 2016, Williamson 2011), within this general framework an analysis is sometimes undertaken as a metaphysical investigation into the *nature* or *essence* of knowledge – as a search for a ‘real definition’, in Locke’s sense (Zagzebski 1999; cf. Earl n.d.: Section 1). At other times, it is undertaken as an analysis of concepts, breaking them down into their simpler components, or of the

¹ Compare Kvanvig’s (forthcoming) concern that the search for a “real definition” -- for “whatever is fundamental to knowledge” -- “doesn’t fit well with the language of necessary and sufficient conditions”, since not everything necessary for someone’s having knowledge reveals something about its nature. Kvanvig sees Williamson (see below) as stressing this distinction between logical and metaphysical relations.

meanings or truth conditions of certain sentence types. And often, the matter is just left unclear.² ('Conceptual analysis' is itself syntactically ambiguous: it might refer to the object of analysis or to its distinctive means or mode.)

Insofar as philosophers in many cases fail to show much concern with distinguishing between the metaphysical and conceptual (or linguistic) projects, one reason may be that these are assumed to be closely related. For instance, those who focus on concepts or words tend to think that theirs isn't a purely psychological or linguistic investigation – the relevant concept/word refers to something real, and an analysis of the former reveals something about the nature of the latter (e.g., MacIver 1958: 1; Hacker 2013: 452; cf. Jenkins 2014: 105-106 and Weatherson 2003: 16-17). Likewise, those whose aims are expressly metaphysical may be seen as employing the concept, and "the method of possible cases" (Jackson 1998: 28), in an attempt to articulate truths about the natures of the relevant worldly epistemic phenomena (e.g., Ichikawa and Jarvis 2013: 228-229; Hetherington 2016: 128-129); and insofar as our judgments about whether someone knows are guided by our concepts, the former can reveal things about the latter (e.g., Goldman and Pust 1998; Jackson 1998: 31-33). (As we'll see, such a happy complementarity has been challenged: some, e.g. Kornblith, hold that we should investigate knowledge itself *rather than* our concepts, the latter being prone to error and of no real philosophical interest.)

Second, whether their stated target is metaphysical or conceptual (/linguistic), all parties within the tradition in question seem committed to providing analyses of the sort described above – i.e., lawlike biconditional statements, subject to testing by a consideration of actual and merely possible cases (Ichikawa and Steup 2016: Introduction; Shope 2004: 285).³ Thus, as Shope notes, an example showing that a statement of the form '*S* knows that *p* only if *q*' or 'If *r*, then *S* knows that *p*' is false would challenge any analysis suggesting otherwise, regardless of the specific form it took. "Perhaps," he says, "this is why so many philosophers leave ambiguous which of the above types of analyses they are pursuing" (2004: 286).

3. Analyzing knowledge: some reasons for scepticism?

As we've seen, while the specific aim and target of those attempting an 'analysis of knowledge' have varied, the shared assumption is that an analysis should provide a noncircular and informative set of necessary and sufficient conditions that survives the method of possible cases. The definition of knowledge as justified true belief is "perhaps the most famous example" of a classical analysis in this vein (Beaney 2016: Section 6). This, it is often said, is the conception of knowledge with which epistemologists since Plato had been operating, often implicitly, and in many cases without using the language of 'justification', until Gettier (1963) showed that there are situations in which the three conditions are satisfied though we balk at granting knowledge.⁴

² Such variability – and, at times, unclarity -- is evident in and across the papers collected in Roth and Galis (1970), for example.

³ Zagzebski calls this "the method of truth condition analysis" (1999: 96).

⁴ For reasons to doubt whether JTB really was *the* going view prior to Gettier, see Kaplan (1985),

In response, and just as Theatetus had done in the face of Socrates' probing, epistemologists modified the account: adding a fourth condition (e.g., a defeasibility, no false lemmas, anti-luck, safety or sensitivity requirement); less often, replacing the justification condition (causal views, some reliabilist accounts); or, more rarely still, advocating a view of justification whereby it entails truth, thereby undermining one of the assumptions made in setting up Gettier cases – namely, that one can be justified in believing a falsehood.

It's no part of the present chapter to scout and evaluate such proposals. For present purposes, the important point is that none has won universal, even general, acceptance: each has been confronted with further purported counterexamples and/or charges of vagueness, *ad hocery*, and so on. Of course, it could be that some such account is correct, or that the correct account has yet to be discovered. After all, while it's plausible that we each have practical mastery of the relevant concepts, there's no reason to think that a correct explicit analysis of them should therefore be easy to obtain or recognize. But the lack of any consensus, and perhaps even of any significant progress, after fifty-plus years of intense theorizing, has dampened many philosophers' enthusiasm for pursuing an analysis of knowledge that survives Gettier cases. Meanwhile, and for reasons beyond the just-noted recent history, a number of philosophers have questioned whether it is worthwhile or possible to provide a reasonably neat set of noncircular and informative necessary and sufficient conditions for knowing.

One reason for disaffection with the program of analysis concerns not its prospects but its value. For example, Mark Kaplan argues that “we simply have no use for a definition of propositional knowledge” (1985: 363). As an inquirer, having carefully weighed the evidence and formed the belief that *p*, there is nothing to be gained from asking, “But do I *know* that *p*?” The question does nothing to advance or clarify the proper conduct of inquiry. Hence, whatever the prospects for a solution to the Gettier problem, it's just not clear what even a widely accepted solution to it would be “good for” (*ibid.*: 363). Other, related complaints about the worth of the attempt to provide an analysis of knowledge concern its failure to engage with radical scepticism (Williams 1978) or to explain the distinctive *value* of knowledge (see Greco 2015: 429; Williamson 2000: 31). No doubt, it is in part because it is seen as failing to engage with such questions that the analysis of knowledge is sometimes derided.⁵ However, even granting the relevant points, they seem not to expose any fatal defect in the program of analysis itself. An analysis or account, of whatever form, is always relative to some objective (Lehrer 1990: 5; Zagzebski 1999). We shouldn't expect that an account that serves one epistemological aim will thereby accomplish others. Thus, that the analysis of knowledge does not address certain other issues does not mean that its typical goal of helping us better understand its target is not

MacIver (1958: 4), Shope (1983: 12-17), Antognazza (2015), and Dutant (2015).

⁵ Beaney, for example, speaks of “idle cog-spinning or epicycling, or Gettier games played for their own sake” (2013: 27). It is not clear whether the perceived detachment from such extra-definitional matters, any more than the apparent failure to have come up with a reasonably neat, non-circular (etc.) analysis that's widely regarded as satisfactory, explains the scorn often directed at “the Gettier industry”. See Lycan (2006).

worthwhile (Conee 1988). Neither is it clear that a better understanding of what knowledge is, or what the concept entails, couldn't at least help in clarifying its value, promoting its acquisition, or addressing other long-standing epistemological issues.

Of course, this all assumes that the project of analyzing knowledge can be satisfactorily carried out. And many now doubt that it can. One argument that might seem to suggest such a conclusion is presented by Linda Zagzebski (1994). According to Zagzebski, for any proposal whereby knowledge is true belief meeting some further condition, x – justification, or justification plus some additional requirement(s) –, so long as x is logically independent of truth, we can construct a Gettier case: imagine that condition x is met, but that unluckily S 's belief is false; then alter the case slightly, so that a further element of luck results in S 's belief's being true after all. Importantly, Zagzebski herself does not conclude from this that knowledge cannot be analyzed. Instead, she infers that “there must be a necessary connection between truth and [x]” (1999: 102) – in particular, that the concept of an act of intellectual virtue, which she sees as closing the gap between (true) belief and knowledge, is truth-entailing (*ibid.*: 111). Many worry that any such solution to the Gettier problem comes at too high a cost, ruling out most ordinary cases of knowing and casting most intuitive epistemic judgments as simply mistaken (e.g., Lycan 2006: 160-161). But Zagzebski does not. Meanwhile, others have argued that incorporating an anti-luck condition into one's account, for example, short-circuits Zagzebski's argument: x need not be truth-entailing in order to avoid the existence of cases having the structure she describes (see Leite 2010: 373). It's an open question, of course, whether the imagined type of account satisfies other standard *desiderata* for a classical analysis mentioned above -- e.g, whether it avoids unacceptable vagueness or circularity.⁶ In general, however, whether Zagzebski's argument provides grounds for pessimism about the prospects for a satisfactory analysis of knowledge is controversial.

A second possible reason for scepticism about the likelihood of finding a (good) classical analysis of knowledge has more metaphysical roots (see Zagzebski 1999: Section 2). As we saw above, one common aim in pursuing such an analysis is to improve our understanding of the *nature* or *essence* of knowing. Natural kinds such as *gold* or *water* might be thought to be amendable to such Lockean real definitions. But is knowledge a natural kind? Perhaps the category of knowledge is more like *refreshing water*, or *offside* (Williams 2015), which seem not to carve nature at the joints. Or perhaps it is like *jade*, constituting no single kind of its own but different types that we happen to have grouped together. The ontological status of knowledge is an important matter to consider. However, two things bear noting here. First, it's not obvious what the correct stance on the issue(s) just posed is – this is no more clear and settled than whether we should be optimistic about the prospects for an analysis of knowledge, on which it was implied the present considerations might shed light. Second, it's not clear how neatly a given stance

⁶ Zagzebski (1999: 103) and Ichikawa and Steup (2016: Section 8), for example, raise such concerns -- concerns that might apply to Zagzebski's own proposal, which she grants “is vague and needs more extensive analysis” (1999: 112). Sorensen (1987) argues that the vagueness of typical fourth conditions supports the view that ‘know’ itself is vague, and that this raises doubts about the prospects for a definition of knowledge (1987: 770-772). See too Hetherington (n.d.: Section 14) and Earl (n.d.: Section 4c).

on the issue(s) just posed lines up with a specific view as to the prospects for a successful classical analysis anyway. Some of the staunchest critics of the latter project (e.g., Kornblith 2007) think that knowledge *is* a natural kind. Further, even if asking after the nature or essence of knowledge really does involve the presumption that this must be given by nature, as we saw above, not everyone who undertakes such an analysis views it as a metaphysical investigation. Goldman, for example, is among the most vocal defenders of the idea that epistemologists must engage in some conceptual investigation of a broadly traditional sort; any doubts he has as to the prospects for a classical analysis (see below) have their source in something other than scepticism -- which he in fact shares (e.g., 2005, 2015) -- about whether knowledge is a natural kind.

We've just fended off some reasons for pessimism about either the value of or the prospects for a classical analysis of knowledge. But is there any good reason to expect that such an analysis is possible in the first place? As Keith Lehrer observes, "the finest monuments of scientific achievement mark the refutation of claims of impossibility" (1990: 6). Even so, it's plausible that not every concept (or word) can be analyzed, on pain of regress. And, not least because of the already-noted absence of clear progress in solving the Gettier problem, current epistemology tends to be less focused on the attempt to analyze knowledge than it was even a generation ago. At the same time, a number of alternatives to the program of analysis have been suggested. In the next Section, we consider the more prominent among these, as well as some more alleged reasons for doubt about the viability or interest of a classical analysis of knowledge.

4. Some proposed alternatives

One influential recent alternative to the program of analysis is the knowledge first approach to epistemology. 'Knowledge first' does not refer to a specific thesis, but to a number of more or less closely related themes and commitments.⁷ Chief among these is the idea that knowledge should be treated as fundamental in at least a couple of senses, each of which contrasts with an assumption of the tradition in which epistemologists sought to analyze knowledge. First, on the traditional view, knowledge is seen as a hybrid state -- i.e., as a composite of mental and non-mental factors. (Belief and truth, respectively; whether and to what extent justification is regarded as mental depends one's theory.) On the knowledge first approach, by contrast, knowledge is seen both as non-composite and as mental through-and-through. Second, and relatedly, on the traditional approach knowledge is treated as a, even *the*, central *explanandum*, with things like evidence, justification, and belief being recruited to help us better understand it. The knowledge first approach reverses this priority: knowledge is a kind of "unexplained explainer" (Williamson 2000: 10); knowledge, or our concept thereof, is used to help us understand such things as justification, evidence, and belief, and/or the corresponding

⁷ For discussion of the latter, and of the relations among them, see Williamson (2010), Ichikawa and Jenkins (forthcoming), and McGlynn (2014).

concepts.^{8,9} Thus, for example, it has been suggested that a subject's evidence is constituted by what s/he knows ('E=K') (Williamson 2000: Chapter 9); that belief "aims at knowledge", with mere believing being "a kind of botched knowing" (*ibid.*: 47); and that, by 'the knowledge rule for belief', "only knowledge constitutes justified belief. ...there are no justified false beliefs." (Williamson 2010: 214) – which, as we've seen, is an idea with obvious application to the Gettier problem.

Clearly, if this approach is correct, knowledge is not analyzable in anything like the usual sense: if knowledge is not composite, it cannot be broken down into simpler elements, and so cannot be given the sort of decompositional (or 'reductive') treatment that classical analysis attempts. Thus, the lack of any clear solution to the Gettier problem, with each successive attempt to provide a satisfactory analysis of knowledge "succumbing to the same pattern of counterexamples and epicycles" (Williamson 2000: 31), is exactly what the knowledge first approach predicts. And indeed, the (perceived) failure of post-Gettier epistemology is often stressed by proponents of the knowledge-first approach and taken to provide strong inductive support for it (e.g., Bird 2007: 82; Kelp 2016: 78-79; Williamson 2010: 209).¹⁰

Importantly, however, there are alternative explanations of the (supposed) insolubility of the Gettier problem, ones that don't single out the knowledge first approach. For example, it is plausible that the program of classical analysis presumes a view of concepts that is now widely discredited (Ramsey 1992, Kornblith 2007). According to the 'classical theory of concepts', a lexical (roughly, 'word-sized') concept is a structured representation; it is composed of simpler concepts that express just what a classical analysis seeks – individually necessary and jointly sufficient conditions for something to fall under the concept's extension (see Margolis and Laurence 2014, Earl n.d.). In spite of its long history and obvious attractions (on which, see *ibid.*), the classical theory has fallen into disrepute.¹¹ One reason for this coincides with a major source of scepticism about the Gettier project: while some concepts yield easily to an analysis (a *bachelor* is an *unmarried male*), most resist the attempt. Just as important, the classical theory fits poorly with various empirical findings.¹² Prominent among the latter are *typicality effects* – i.e., the fact that in various tasks subjects treat apples, for example, as more typical or representative fruits than plums, though there's no room for one kind of fruit to *better satisfy* the presumed definition.

⁸ As in the program of analysis, there is occasional unclarity about whether what's at issue is the relevant concepts or that to which they refer (Ichikawa and Jenkins forthcoming; McGlynn 2014: 17).

⁹ Cook Wilson provides a clear and recent precedent for this type of approach (see Marion 2016: Section 4). According to some, Plato for example constitutes a much earlier one (see the works cited in note 5 above).

¹⁰ Williamson (2000: 31, n. 4) also cites Fodor's (1998) discussion of "the demise of definition".

¹¹ Murphy, for instance, says that "[t]o a considerable degree, it has simply ceased to be a serious contender in the psychology of concepts" (2002: 38); more bluntly, he calls it "a total flop" (*ibid.*: 483).

¹² For reviews, see Murphy (2002) and Smith and Medin (1981).

The psychology of concepts is itself the focus of intense debate, and even if the classical theory is rejected it's not clear what to put in its place (Murphy 2002, Margolis and Laurence 2014). So it's not obvious how we should think of the concept of *knowledge*, if not along the lines suggested by the classical theory. In response to certain knowledge first arguments, both Adam Leite (2005: 168) and Elizabeth Fricker (2009: 44-45) have suggested that it may be an 'open texture' or 'familiar resemblance' concept, possessing semantic complexity, but none capturable by any set of (noncircular) necessary and sufficient conditions.¹³ Current theories of concepts congenial to the latter idea include the prototype and exemplar views. (Briefly, a *prototype* is a collection of properties often present in instances of the concept, weighted by frequency or perceptual salience; an *exemplar* is a specific instance of the kind in question that the cognizer has in mind (Goldman 1993: 128-129; Earl n.d.: Section 2d).) Perhaps predictably, the latter views tend to excel where their chief rivals falter, and *vice versa*; and a similar combination of strengths and weaknesses attends other views (e.g., the 'theory' or 'knowledge' approach and 'neoclassical' theories.)¹⁴ For our purposes here, however, the important point is that, as Leite puts it, "it is implausible to claim that if a concept of a certain state or condition cannot be provided with informative^[15] necessary and sufficient conditions for its application, then we do not conceive of that state or condition as a complex involving different kinds of factors" (2005: 168). And if there *is* complexity there, there is the opportunity for analysis, whether or not it is undertaken with the expectation that it will result in a set of neat, noncircular (etc.) necessary and sufficient conditions. Thus, for example, we find Goldman (2015) suggesting that the discovery of typicality effects, and the consequent implausibility of the classical theory of concepts, calls for a *refinement* of the usual approach to philosophical analysis, rather than its abandonment:

"...I see no problem in principle here. There should be ways to incorporate into the structure of concepts a weighting scheme that assigns different 'strengths' to different properties. Cognitive science can and should be used to improve the kind of investigations that traditional philosophers engage in. However, as far as I can see, this would not imply a thorough debunking or abandonment of methods that use hypothetical examples and intuitive classifications thereof. (Indeed, the finding of typicality effects relied on just such intuitive classifications.)"

While Goldman is speaking here in defense of key components of "the standard practice of conceptual analysis" (*ibid.*) – specifically, the use hypothetical examples and intuitive

¹³ Lycan (2006: 158-159) argues that *knowledge* does not possess a 'family resemblance' structure in particular. Hacker (2013: Chapter 4, Appendix) rejects the latter model for *knowledge* too, suggesting instead that it is a 'multi-focal' concept best treated via a connective analysis. (On the latter notion, more below.)

¹⁴ For instance, while prototype and exemplar theories seem well-suited to accommodating typicality effects, they are said by some (e.g., Fodor 1998) to struggle with compositionality.

¹⁵ The qualification is important. As we saw above, uninformative such conditions are all too easy to provide. In terms of the prototype theory, for example, "being sufficiently similar to the relevant prototype is necessary and sufficient for being a bird" (Jackson 1998: 61; cf. Weatherson 2003: 18-19). As Kornblith says, then, "[a]ny account of concepts will give necessary and sufficient conditions for application of the concept, including prototype and exemplar accounts" (2007: 41 n. 10; cf. Goldman 2007: 23).

classifications thereof --, it should be noted that knowledge first epistemologists employ such methods too, and that they do so in service of claims and theories that often look a lot like the products of traditional conceptual analysis. According to Aiden McGlynn, “the knowledge first literature is no less full of attempts to state informative necessary and sufficient conditions than the rest of the post-Gettier literature” (2014: 17). Relatedly, that literature contains familiar post-Gettier examples and patterns of argument. For example, Christoph Kelp (2016) argues that various other knowledge first accounts of justification fall prey to Gettier counterexample, whereas his does not; Williamson’s claim that “only knowledge constitutes full justification” (2010: 215) recalls views (like Zagzebski’s) on which what must be added to true belief to yield knowledge is truth-entailing, and is subjected to similar criticisms (e.g., Cohen and Comesaña 2013: 15-16); and familiar counterexamples, such as barn cases, are used against various knowledge first theses, prompting refinements, alternatives, and new distinctions (McGlynn 2014: 173).¹⁶

All of this might lead one to wonder whether knowledge first epistemologists aren’t themselves engaged in something *like* old-fashioned analysis. It is important to note, then, that “rejection of the project of analyzing knowledge in no way suggests that there are not interesting and informative necessary or sufficient conditions on knowledge. The traditional ideas that knowledge entails truth, belief, and justification are all consistent with the knowledge first project. And Williamson (2000, 126) is explicit in endorsement of a safety condition on knowledge” (Ichikawa and Steup 2016: Section 10). Knowledge firsters hold that such conditions reflect important “structural features” of knowledge (Williamson 2009: 306) and contribute to a “reflective understanding” of the concept (Williamson 2000: 33). What they deny is that agreement on such conditions represents progress towards a reductive account of knowledge, wherein it is broken down into simpler components -- something, they think, that cannot be done.¹⁷ Still, as we’ve seen, the energies of those working within the program of analysis tend to be focused on such conditions themselves – getting *those* right -- rather than on the question of whether knowledge, or our concept thereof, is composite. At least in that sense, then, there may be greater kinship of goal and method between that program and the knowledge first approach than one might have thought – especially since, as we’ll see shortly, some traditional, ‘belief first’ theorists adopt the methods of classical analysis while sharing the knowledge first epistemologist’s scepticism that it will yield a neat, non-circular and exceptionless *analysans*. So too, those who regard knowledge as composite, but who are nonetheless sceptical about the prospects for providing a classical analysis of it, are of

¹⁶ McGlynn writes: “If the pattern of counterexamples found in the post-Gettier literature were a symptom of the pursuit of a bad research programme – the pursuit of analyses – then we would expect to find that abandoning that research programme would result in a break in the pattern. [But...] this hasn’t happened. On the contrary, the approach has been plagued by apparent counterexamples, often running into problems with the very examples that cropped up in the Gettier literature” (2014: 172-173).

¹⁷ A favorite example of Williamson’s in support this point is the following: “Although being coloured is a necessary but insufficient condition for being red, we cannot state a necessary and sufficient condition for being red by conjoining being coloured with other properties specified without reference to red. Neither the equation ‘Red = coloured + X’ nor the equation ‘Knowledge = true belief + X’ need have a non-circular solution” (2000: 3; cf. 2010: 210, and n. 1 above).

course free to investigate its structural features, as Williamson calls them, and/or to pursue a reflective understanding of the concept.

Hilary Kornblith has put forward another alternative to the program of analysis. According to Kornblith, not only does traditional philosophical analysis require a thoroughly discredited theory of concepts (*viz.*, the classical theory), but the *concept* of knowledge is of virtually no philosophical interest anyway: it is no more a worthy target of epistemological theory than the *concept* of aluminum is a worthy target of inquiry for a metallurgist. Kornblith himself thinks that knowledge is a natural kind. But even if it is not – even if it is, say, a socially constructed category -- there is no assurance that our ‘folk’ concept of knowledge, or even the epistemologist’s, characterizes it correctly. And when we realize that our epistemic concepts “may fail to characterize the categories they are concepts of, the philosophical interest of our concepts thereby wanes” (2007: 37). Instead, we should simply study knowledge itself.

As to methodology, Kornblith’s proposal is that epistemology should be thoroughly empirical: taking their cue from our imagined metallurgist, epistemologists should proceed by examining apparently clear cases of knowledge, looking to find what they have in common. Very likely, this process will involve reclassifying examples, deciding that some aren’t cases of knowledge after all, and that other, previously excluded examples are. Out of this process will emerge a picture of the true nature of knowledge. Specifically, Kornblith thinks, what will emerge is a picture of knowledge as reliably produced true belief, with the latter being “instrumental in the production of behavior successful in meeting biological needs and thereby implicated in the Darwinian explanation of the selective retention of traits” (Kornblith 2002: 62).

In response to the foregoing line of argument, one concern is that it moves too quickly from the possibility that our concepts incorporate error to their being of no theoretical interest. “The fact that there is a ‘gap’ between ‘concept and category’ is no more worrying than the fact that there is a ‘gap’ between how things *look* or *sound* to us and how they actually are” (Jenkins 2014: 106-107). A second concern is that without *some* reliance on our pretheoretic concepts the empirical study Kornblith recommends could not get off the ground: “a prior method is needed to pick out which set of extra-mental events in the world should be the target of a Kornblithian empirical investigation” (Goldman 2015; compare Jackson 1998: 30-31, 41-42; MacIver 1958: 2; and Weatherson 2003: 15-16).

According to Goldman, the needed prior method is an analysis of our concepts, as grounded in our intuitive classifications of various cases, actual and hypothetical. In a couple of respects, however, the type of analysis Goldman advocates departs importantly from the classical model. First, while in earlier work (e.g., 1986) Goldman seemed to regard conceptual analysis and the consulting of intuitions as an *a priori* method, more recently (Goldman 1999; 2005; 2007; Goldman and Pust 1998) he has suggested that the conceptual work characteristic of epistemological theorizing is a form of *a posteriori*, empirical investigation: a sample case is presented, and the subject’s classificational intuitions constitute quasi-observational data revealing features of his/her epistemic

concepts (2005: 408-409).¹⁸ Second, in part because of typicality effects and the consequent discrediting of the classical theory of concepts, Goldman doubts that a subject's concepts should be represented as specifying informative and non-circular necessary and sufficient conditions (e.g., 1986: 38-39; 2015; 2007: 23).¹⁹ (As we saw above, he suggests that those engaged in conceptual analysis should find a way of including a weighting scheme that assigns different 'strengths' to different properties.) At the same time, however, Goldman insists that, in addition to including all (philosophically significant) necessary conditions, an adequate account of a given concept should at least be *tested* in the usual way for sufficiency. If it fails, then "either there is some *sufficient* condition that has not yet been identified or...one or more *necessary* (but insufficient) conditions have been *omitted*" (2009: 75). In this way, the methods of classical analysis are retained, and even regarded as essential tools of the epistemologist, though there is no presumption that the investigation will yield a neat, non-circular and exceptionless classical *analysans* (cf. Beaney 2016: Section 6), or one expressing a necessary truth (see Zagzebski 1999: 96, including n. 12).²⁰

Still another alternative to the program of analysis is defended by Edward Craig (1990). While Craig is sceptical about the prospects for solving the Gettier problem, his more fundamental concern with the program of analysis is that even if we had an extensionally adequate account of knowing, it may not connect with the concept's "intuitive intension" – i.e., with our intuitions about *why* certain cases do, or do not, qualify as knowledge (1990: 1). More generally, even a successful analysis may fail to engage with such questions as why knowledge is worth caring about, and why the concept enjoys such widespread use (1990: 2).

We encountered such concerns above (Section 2), and denied that they exposed any

¹⁸ A further reason for scepticism about conceptual analysis stems from work within 'experimental philosophy' ('X-phi') purporting to show that there is significant diversity in the epistemic intuitions on which such analyses rely. Such intuitions are said to vary both within individuals (Swain *et al.* 2008) and between groups, including "along such epistemically scary fault lines" (Nagel 2012: 495) as ethnicity (Weinberg *et al.* 2001) and gender (Buckwalter and Stich 2011). The interpretation and significance of the relevant data are controversial. Further, more recent studies (Nagel *et al.* 2013, Seyedsayamdost 2015, Kim and Yuan 2015) have failed to replicate the results in question. And in his most recent work on the subject, Stich – along with his coauthors (see Machery *et al.* 2015) – has argued for the cross-cultural *robustness* of certain epistemic intuitions, including about Gettier cases.

¹⁹ Compare Chalmers and Jackson's view that conceptual analysis "proceeds at least in part through consideration of a concept's extension within hypothetical scenarios, and noting regularities that emerge" (2001: 321), where it is not assumed that such consideration involves definitions – that is, "finite expressions in the relevant language that are a priori equivalent to the original terms, yielding counterexample-free analyses of those terms" (*ibid.*: 320).

²⁰ Goldman's views here resemble Zagzebski's. Zagzebski has doubts about whether it's possible to satisfy all the requirements for a good definition simultaneously, and herself favors "retaining the method of truth condition analysis but without letting the aim to make the definition counterexample-free dominate the list of desiderata" (1999: 104). Lycan (2006: 158) and Plantinga (1993: 20) are two others who employ the methods of classical analysis while granting that central epistemic concepts may not fully yield to them.

problems with the project of analyzing knowledge *per se*. Even so, Craig argues, given the difficulties of post-Gettier epistemology, it's surely worthwhile to approach the concept of knowledge via a different method altogether (1990: 1). Specifically, he suggests that instead of fixing directly on the concept and trying to analyze it, we should ask "what knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like" (1990: 2). According to Craig, a conspicuous general fact about the human situation is that we must rely on others as sources of information (1990: 11). So we need some way to pick out *good* informants, and Craig's hypothesis is that our concept of the knower is the highly 'objectivized' notion of a good informant, where the latter is someone who is likely to have a very high degree of reliability on the matter in question (1990: 91).

Such a "practical explication"²¹ of the concept of knowledge might help to explain some of the attractions of the JTB conception, since a typical good informant may satisfy those three conditions (1990: 69-70). (At the same time, that there may be good informants who *don't* satisfy one or more of those conditions is no objection to Craig's account, since it doesn't aim at supplying necessary and sufficient conditions for knowing.) So too, that epistemologists have proposed 'truth-tracking', causal, and reliabilist analyses of knowledge is hardly surprising, since these focus on features plausibly associated with a person's being likely to be correct on some matter. However, Craig argues, there is no property, *X*, such that its satisfaction by the subject guarantees that *s/he* has a high probability of being right about *p*. For "there will always be something else which we could come to believe (call it *Y*), such that (*X* and *Y*) doesn't lend much probability to '*S* is right as to whether *p*'" (1990: 52); as with other proposed analyses, sufficiency can be achieved only by requiring something too strong to be necessary (1990: 81). So it is not surprising either that the Gettier problem remains unsolved.

Craig's practical explication of *knowledge* can be seen as one instance of a more general alternative to the program of classical analysis – what P.F. Strawson calls the method 'connective analysis' (1992: Chapter 2). When tasked with elucidating some concept, one option is indeed to employ the "reductive or atomistic" method (1992: 21) of 'dismantling' it, identifying its simpler elements and specifying necessary and sufficient conditions for its application. A "more realistic and fertile" approach (1992: 19; see too Hacker 2013: Appendix), Strawson suggests, is to undertake a *connective analysis*, which involves articulating the connections between a given concept and others with which it is

²¹ A Craigian 'practical explication' differs from Rudolph Carnap's method of explication. Whereas a practical explication, like a classical analysis, aims for descriptive adequacy (Craig 1990: 8), Carnapian explication is normative – it involves "making more exact a vague or not quite exact concept used in everyday life or in an earlier stage of scientific or logical development, or rather of replacing it by a newly constructed, more exact concept" (Carnap 1947: 7-8). While this new, more precise concept should be similar to the original, it should also be simple and theoretically fruitful, connecting easily with extant concepts and laws and facilitating the formulation universal statements. A recent application of Carnap's method to *knowledge* and the Gettier problem is Olsson (2015). According to Olsson, failure to give the right intuitive result in Gettier cases can be offset by other theoretical virtues, and so does not itself threaten a given explication. For related ideas and arguments, see Weatherson (2013).

associated. (In Craig's case, between the concept of *knowledge* and that of a *good informant*.) There is no presumption that, as we elucidate a complex concept by tracing its connections with others, there will be a movement towards either greater simplicity or anything resembling a classical *analysans* (1992: 22-23). Nor should we be surprised or troubled if, "in the process of tracing connection from one point to another of the network [of concepts], we find ourselves returning to, or passing through, our starting point" (1992: 19). Such a 'circular' path may well be illuminating, improving our understanding of the target concept.²²

According to John Greco, Craig's approach also illustrates a more general shift in epistemological methodology that has occurred post-Gettier. As Greco sees it, 'Gettier-era epistemology' was driven by "intuitionism" – the idea that epistemological theories should be evaluated "according to how well they preserve our pretheoretical intuitions about particular cases" (2015: 426). "[T]he method of counterexample" – and, more generally, of possible cases – that's characteristic of classical analysis "is in fact the method of intuitionism" (*ibid.*). There are worries about such a method, however, some of which we encountered above -- for example, that it tends to promote merely extensionally adequate theories, rather than genuinely illuminating ones (2015: 426); relatedly, that two (or more) theories might do equally well in accommodating intuitions about particular cases, in which case we'd need some further basis for choosing between them (2015: 426-427); and, once again, that the method has failed to generate any significant consensus post-Gettier (2015: 423).²³

Yet, in spite of all this, "epistemology in the early part of the twenty-first century is alive and well, even flourishing" (2015: 434). What explains this? The fact that epistemologists have moved away from mere 'intuitionism', incorporating new methodological constraints on an adequate theory of knowledge. Specifically, they've begun to demand that such a theory explain such things as the distinctive value of knowledge, its close connection with both assertion and action, the social role of the concept, and the point and purpose of epistemic evaluations more generally (Craig 1990, Henderson and Greco 2015).

No doubt, the general shift in methodology Greco describes is real. It is important to stress, however, that it is a *shift*, rather than the swapping-in of an old method for a new one.²⁴ For example, the above-noted demand (Section 2) that a good analysis be *informative* is an acknowledgement within the program of analysis of the insufficiency of mere extensional adequacy, and so of the importance of considerations beyond mere conformity to intuitions about particular cases. Likewise, the broader methodology Greco

²² Strawson's notion of a connective analysis is anticipated in A.C. Ewing: see Beaney (2016: Section 8). It also recalls Ryle's (1949) notion of a 'logical geography' of concepts. On the similarities and differences between connective analysis and Williamson's knowledge first approach, see Cassam (2009: 24-25) and Williamson (2009: 285-290).

²³ Greco also cites experimental philosophy and the doubts it might raise about "the evidential quality of our intuitions about cases" (2015: 428). As he notes, the relevant results and their interpretation are contentious, even within X-phi; see n. 18 above.

²⁴ I don't mean to suggest that Greco sees things otherwise.

describes includes epistemologists' making liberal use of our epistemic intuitions about particular cases and the method of counter-example – and so, in Greco's terms, of the method of intuitionism. In general, whatever the prospects for an analysis of knowledge – in its classical form, or in the modified form that Goldman, say, envisages – there is no reason to think that the methods employed therein can't be peacefully and productively used alongside, or within, the broader methodological framework Greco describes. The same point applies to Strawsonian connective analysis, and – once again -- to the knowledge first epistemologist's similarly non-reductive focus on 'structural features' of knowledge: they too afford useful means of approaching the relevant concepts while promoting clarity and rigor and improving our understanding of epistemic subjects. Employing them does not require that one take the concept of *knowledge* not to be analyzable, or knowledge itself not to be composite; it no more requires this than using the method of possible cases and attempting to state informative conditions on knowing requires that one think that they *are*.

5. Conclusion

The Gettier problem is often seen as casting a shadow over recent epistemology. But it has had the welcome effect of encouraging greater reflection on epistemological methodology. As we've seen, there is little consensus about either the value of seeking an analysis of *knowledge* or whether there are principled reasons why the attempt is bound to fail. More generally, there is plenty of lively debate about the appropriate methodological moral to draw from the Gettier problem. Whatever the correct position to take here, an increased awareness of the methodological issues needing to be addressed, and of the options available to us moving forward, is surely a good thing.

References

- Antognazza, M. R. (2015) "The Benefit to Philosophy of the Study of its History", *British Journal for the History of Philosophy* 23(1): 161-184.
- Beaney, M. (2013) "What is Analytic Philosophy?", in *The Oxford Handbook of The History of Analytic Philosophy*, edited by M. Beaney. Oxford: Oxford University Press, pp. 3-29.
- ___ (2016) "Analysis", *The Stanford Encyclopedia of Philosophy* (Summer 2016 Edition), E. N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/sum2016/entries/analysis/>](https://plato.stanford.edu/archives/sum2016/entries/analysis/).
- Bird, A. (2007) "Justified Judging", *Philosophy and Phenomenological Research* 74: 81-110.
- Buckwalter, W. and S. Stich (2011) "Gender and the Philosophy Club", *The Philosophers' Magazine* 52: 60–65.

- Carnap, R. (1947) *Meaning and Necessity*. Chicago: The University of Chicago Press.
- Cassam, Q. (2009) “Can the Concept of Knowledge be Analysed?”, in Greenough and Pritchard, pp. 12-30.
- Chalmers, D. J., and F. Jackson (2001) “Conceptual Analysis and Reductive Explanation”, *The Philosophical Review* 110(3): 315-360.
- Conee, E. (1988) “Why Solve the Gettier Problem?”, in *Philosophical Analysis: A Defense by Example*, edited by D. Austin. Dordrecht: Kluwer, pp. 55-58.
- Craig, E. J. (1990) *Knowledge and the State of Nature*. Oxford: Clarendon Press.
- Dutant, J. (2015) “The Legend of the Justified True Belief Analysis”, *Philosophical Perspectives* 29: 95-145.
- Earl, D. (n.d.) “The Classical Theory of Concepts”, *Internet Encyclopedia of Philosophy*: <http://www.iep.utm.edu/conc-cl/>.
- Fodor, J. (1998) *Concepts*. Oxford: Oxford University Press.
- Fricker, E. (2009) “Is Knowing a State of Mind? The Case Against.” in Greenough and Pritchard, pp. 31-60.
- Gettier, E. (1963) “Is Justified True Belief Knowledge?”, *Analysis* 23: 121-123. Reprinted in Roth and Galis (eds.), pp. 35-38.
- Goldman, A. (1986) *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- ___ (1993) *Philosophical Applications of Cognitive Science*. Boulder, CO: Westview Press.
- ___ (1999) “A Priori Warrant and Naturalistic Epistemology”, in J. E. Tomberlin (ed.), *Philosophical Perspectives* 13. Cambridge, UK: Blackwell.
- ___ (2005) “Kornblith’s Naturalistic Epistemology”, *Philosophy and Phenomenological Research*, 71(2): 403-410.
- ___ (2007) “Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status”, *Grazer Philosophische Studien*, 74: 1-26.
- ___ (2009) “Williamson on Knowledge and Evidence”, in Greenough and Pritchard, pp. 73-92.
- ___ (2015) “Review of Hilary Kornblith, *A Naturalistic Epistemology: Selected Papers*”,

Notre Dame Philosophical Reviews: <http://ndpr.nd.edu/news/a-naturalistic-epistemology-selected-papers/>.

____, and J. Pust (1998) “Philosophical Theory and Intuitional Evidence”, in M. DePaul and W. Ramsey (eds.), *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Lanham, MD: Rowman and Littlefield, pp. 179-197.

Greco, J. (2015) “Post-Gettier Epistemology”, *Veritas* 60 (3): 421-437.

Greenough, P. and D. Pritchard (eds.) (2009) *Williamson on Knowledge*. New York: Oxford University Press.

Henderson, D., and J. Greco (eds.) (2015) *Epistemic Evaluation: Purposeful Epistemology*. Oxford: Oxford University Press.

Hetherington, S. (n.d.) “Gettier Problems”, *Internet Encyclopedia of Philosophy*: <http://www.iep.utm.edu/gettier/>.

____ (2016) *Knowledge and the Gettier Problem*. Cambridge: Cambridge University Press.

Ichikawa, J. J., and C. S. I. Jenkins (forthcoming) “On Putting Knowledge ‘First’”, to appear in J. A. Carter, E. C. Gordon, and B. Jarvis (eds.), *Knowledge First: Approaches to Epistemology and Mind*. Oxford University Press.

Ichikawa, J. J. and M. Steup (2014) “The Analysis of Knowledge”, *The Stanford Encyclopedia of Philosophy*, (Spring 2014 Edition), E. N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2014/entries/knowledge-analysis/>>.

Jackson, F. (1998) *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon Press.

Jenkins, C. S. I. (2014) “‘Intuition’, Intuition, Concepts and the A Priori”, in A. Booth and D. Rowbottom (eds.) *Intuitions*. Oxford: Oxford University Press, pp. 91-115.

Kelp, C. (2016) “Justified Belief: Knowledge First-Style”, *Philosophy and Phenomenological Research* 93(1): 79-100.

Kim, M., and Y. Yuan (2015) “No Cross-Cultural Differences in the Gettier Car Case Intuition: A Replication Study of Weinberg et al. 2001”, *Episteme* 12(3): 355-361.

Kornblith, H. (2002) *Knowledge and Its Place in Nature*. New York: Oxford University Press.

____ (2007) “Naturalism and Intuitions”, *Grazer Philosophische Studien* 74: 27–49.

Kvanvig, J. L. (forthcoming) “Lessons from Gettier”, in R. Borges, C. de Almeida, and P. D. Klein (eds.), *Explaining Knowledge: New Essays on the Gettier Problem*. Oxford University Press.

Lehrer, K. (1990) *Theory of Knowledge*. Boulder, CO: Westview Press

Leite, A. (2005) “On Williamson’s Arguments that Knowledge is a Mental State”, *Ratio* 18(2): 165-75.

____ (2010) “Fallibilism”, in J. Dancy, E. Sosa and M. Steup (eds.). *A Companion to Epistemology*, 2nd ed. Malden, MA: Blackwell, pp. 370-375.

Lycan, W. (2006) “On the Gettier Problem Problem”, in Stephen Cade Hetherington (ed.), *Epistemology Futures*. Oxford: Oxford University Press, pp. 148-168.

Machery, E., S. Stich, D. Rose, A. Chatterjee, K. Karasawa, N. Struchiner, S. Sirker, N. Usui, and T. Hashimoto (2015) “Gettier Across Cultures”, *Nous* (early view): doi:10.1111/nous.12110.

MacIver, A. M. (1958) “Knowledge”, *Proceedings of the Aristotelian Society*, Supplementary Volume 32: 1-24.

Marion, M. (2016) “John Cook Wilson”, *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), E. N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2016/entries/wilson/>.

McGlynn, A. (2014) *Knowledge First?* Houndmills, Basingstoke, Hampshire: Palgrave Macmillan.

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

____, V. San Juan, and R. A. Mar (2013) “Lay Denial of Knowledge for Justified True Beliefs”, *Cognition* 129(3): 652–661.

Murphy, G. L. (2002) *The Big Book of Concepts*. Cambridge MA: The MIT Press.

Plantinga, A. (1993) *Warrant and Proper Function*. Oxford: Oxford University Press.

Ramsey, W. (1992) “Prototypes and Conceptual Analysis”, *Topoi* 11: 59-70.

Roth, M. D. and L. Galis (eds.) (1970) *Knowing: Essays on the Analysis of Knowledge*. Lanham, New York, London: University Press of America.

Ryle, G. (1949) *The Concept of Mind*. London: Hutchinson’s University Library.

- Seyedsayamdost, H. (2015) "On Normativity and Epistemic Intuitions: Failure of Replication", *Episteme* 12(1): 95-116.
- Shope, R. K. (1983) *The Analysis of Knowing: A Decade of Research*. Princeton, New Jersey: Princeton University Press.
- ____ (2004) "The Analysis of Knowing", in I. Niiniluoto, M. Sintonen, and J. Wolenski (eds.). *Handbook of Epistemology*, Dordrecht, Netherlands: Kluwer Academic Publishers, pp. 283-329.
- Smith, E., and D. Medin (1981) *Categories and Concepts*. Cambridge, MA: Harvard University Press.
- Sorensen, R. A. (1987) "The Vagueness of Knowledge", *Canadian Journal of Philosophy* 17(4): 767-804.
- Stich, S. (1992) "What is a Theory of Mental Representation?", *Mind* 101(402): 243-261.
- Strawson, P. F. (1992) *Analysis and Metaphysics: An Introduction to Philosophy*. Oxford: Oxford University Press.
- Swain, S., J. Alexander, and J.M. Weinberg (2008) "The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp", *Philosophy and Phenomenological Research* 76(1): 138-155.
- Weatherson, B. (2003) "What Good are Counterexamples?" *Philosophical Studies* 115(1): 1-31.
- Weinberg, J. S., S. Nichols, and S. Stich (2001) "Normativity and Epistemic Intuitions", *Philosophical Topics* 29(1): 429-460.
- Williams, M. (1978) "Inference, Justification, and the Analysis of Knowledge", *The Journal of Philosophy* 75(5): 249-263.
- ____ (2015) "What's So Special about Human Knowledge?", *Episteme* 12(2): 249-268.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- ____ (2009) "Replies to Critics", in Greenough and Pritchard, pp. 279-384.
- ____ (2011) "Knowledge First Epistemology", in S. Bernecker, & D. Pritchard (eds.), *The Routledge Companion to Epistemology*. New York: Routledge, pp. 208-218.
- Zagzebski, L. (1994) "The Inescapability of Gettier Problems", *Philosophical Quarterly* 44: 65-73.
- ____ (1999) "What is Knowledge?", in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology*. Oxford: Blackwell, pp. 92-116.