The epistemology of neo-Gettier epistemology[1]

Robert Lockie
Psychology Department, University of West London, Brentford, Middlesex, UK
bob@lockie.f9.co.uk

The paper begins by drawing a number of 'levels' distinctions in epistemology. It notes that a theory of knowledge must be an attempt to obtain knowledge (about knowledge). It is suggested that we can make sense of much of the work found in analytic theory of knowledge by seeing three (tacit) framework assumptions as underpinning this work. First, that to have philosophical knowledge of knowledge requires us to have an analysis. Second, that much of what we require from a theory of knowledge may be obtained by developing such analyses of first-order, concrete, empirical, propositional knowledge. Third, that the final arbiter of the correctness of such analyses is to be the carefully examined intuitions of the epistemologist. The paper attacks each aspect of this framework on the ground that this methodology will precisely not give us knowledge (of knowledge). In passing, comparisons are drawn with arguments that led to the demise of phenomenalism. The paper concludes with remarks about realism/anti-realism and consensus/disagreement in analytic epistemology. The paper recommends that we seek to develop theories of knowledge rather than analyses, and defends the position that such theories will precisely not be analyses.

Introduction and aims

Within analytic philosophy, broadly conceived, doubts have been accumulating over more than a decade as to the way the theory of knowledge has been most characteristically conceived and investigated over about the last 50 years. At the least there is a suspicion that this is a research program that has reached the end of its productive life. Arguing that this is so, and diagnosing why it is so, is the aim of this paper. This paper is critical of three metaphilosophical and methodological assumptions under which much work in the analytic theory of knowledge was conducted during the last half-century, and under which a substantial body of work continues to be conducted. It is argued that over this time the theory of knowledge embraced a series of framework assumptions that shackled the discipline and increasingly led it to resemble a self-standing game. (Foley 2004 called it the 'Gettier Game'.) These assumptions are tacit, are framework assumptions, are metaphilosophical (are meta-epistemic) and are methodological.

Foley is no debunking naturalist and nor am I. It should be emphasised that from the outset that this paper's negative arguments, although informed by the investigations into knowledge that took place within scientific psychology over the same half-century, are nevertheless not motivated by any adherence to a generally deflationary, nor yet a generally naturalistic metaphilosophy. It will be argued, rather, that the methodology and metaphilosophy exemplified by neo-Gettier epistemology is specifically indefensible. Since the target of this paper is a largely tacit conceptual framework, there is always the danger of being taken to be attacking a straw man. Note though, that the purpose here is ultimately to make room for alternative theories of knowledge – theories that don't fit within the framework identified. The arguments, then, are put forward in the following spirit: if the cap doesn't fit, don't wear it.

Theory of knowledge: a 'levels' distinction

In what follows we will draw upon a 'levels' distinction. At the first-order level, the objects of knowledge may be issues of book-thief identity, what the time is, how to get from Athens to Larissa, and much else. These things could be theories – of physics or geology or whatever else – but noted critically below is the fact that in the analytic tradition they never are, they are rather doxastic items: specific states-of-affairs, as represented by propositions. When we inquire into the nature of *knowledge* though, the aim of our inquiry is to develop a theory, a theory of the state or relation an agent must be in when she knows or doesn't know such first-order things – the time, who owns this car, etc. In this case the object of knowledge involves the state or relation that obtains between a first-order knower and his object of knowledge. Epistemology is, then, already a higher-order activity: seeking to *know* about *knowledge* (which here, with this tradition, concerns things such as cars, and barns, and the time). But we can seek to understand the nature of knowledge of knowledge in turn, and that is meta-epistemology. In what, then, does the knowledge of *theories of knowledge* consist?

Any correct theory within any discipline must, when possessed in the right way (when understood), give us knowledge. This should apply to theories of knowledge no less than anything else. Searching for the correct theory of knowledge means that we are seeking to *know* about knowledge. And this offers us an important constraint on theories of knowledge. One way of testing a theory of knowledge is to ask whether possession of the theory could give us knowledge (of knowledge). That is, (1) does the theory employ a *methodology* apt to lead to knowledge (of knowledge)? (2) Is the theory in a form such that possessing it could *constitute* knowledge (of knowledge)?

The methodology that, historically, has been most used in constructing analytic theories of knowledge is to take a putative analysis of knowledge, pick ever more contrived thought experiments (albeit involving cases of highly concrete, first-order, empirical knowledge – that *Jones owns a Ford*, that *that's a barn*, etc.), then ask if the analysis classifies these cases as our intuitions hold that they should be classified. If the theory doesn't classify such cases in accordance with our intuitions, this shows that the analysis is either not necessary or not sufficient for knowledge, and must be modified or rejected accordingly – for some alternative analysis.

Tacitly, such approaches appear to be motivated by three framework assumptions.

(1) That to have knowledge of knowledge requires us to have an analysis (paradigmatically, a set of necessary and sufficient conditions) – that it is in the possession of such analyses that knowledge of knowledge must consist.

(2) That much – in fact most – of what we as philosophers require from a theory of knowledge is to be obtained by developing such analyses of a certain kind of first-order, concrete, empirical, propositional knowledge (of Ford-ownership, book-thief-identity, etc.).

(3) That the final arbiter of the correctness of such analyses is to be the (refined, carefully examined) intuitions of the epistemologist.

This paper is an attack on approaches to normative epistemology that embrace these three assumptions. In pressing this attack, the critical question will be whether theories embodying these assumptions are capable of furnishing us with knowledge (of knowledge). It will be argued that they are not.

Analysis as such

Consider the idea that to have knowledge of something (here, knowledge about knowledge) we must have a set of necessary and sufficient conditions for that thing; or at least, that this is the ideal we should aspire to, with at best repeated failure leading us to settle for something weaker – such as necessary and nearly sufficient conditions.[2] After Rosch, herself following Wittgenstein, we would not today accept this, or anything like it, as a requirement on possession of much (any?) first-order knowledge concerning matters of fact (Wittgenstein 1953; Rosch & Mervis 1975; Rosch 1978; Smith & Medin 1981; Murphy & Medin 1985). Of course, not so

2   Or a recursive series of heavily 'patched' accounts that grow out of and begin with the attempt to find a single necessary and sufficient account. Or any of a number of stylistically grueling variants of such cases, in which the analytic ambition is in some core sense retained despite the repeated failure to achieve it.

long ago, things were different here – in both philosophy and psychology. Phenomenalism (in psychology: structuralism) assumed that knowing that *that's a chair* or *that's Jones* required us to have an analysis, built ultimately out of sensory simples, of that object of knowledge. Such early approaches to perceptual knowledge, though dead in psychology by the 1920s and in philosophy by the 1960s, persisted even into the early era of research into conceptual knowledge (cf. Bruner et al. 1956). These approaches are now long abandoned. Psychologists investigating such first-order knowledge reject the idea that almost any of it is analytic in form. Embarrassingly for epistemology, the psychologists explicitly derived their opposition to analysis by rigorously 'psychologising' a number of philosophical accounts – for instance, the neo-Kantian accounts of memory and perception to be mentioned below, and accounts of concept possession developing out of Wittgenstein's 'family resemblance' remarks.

Psychological research into knowledge teaches us something that common sense and philosophical experience should already have primed us for: that very little of our knowledge is in the form of definitions. Our knowledge of what a chair is, or who Jones is, is not any kind of definition – not even a compressed, tacit or elliptical definition. A sharp challenge to the first component of this framework view is then to ask: if even first-order knowledge (of furniture, faces, etc.) isn't analytic, why suppose that our knowledge of what knowledge is is any such kind of thing? (cf. Kornblith 2007). It is highly doubtful whether knowledge of knowledge consists in a definition of this concept – or even a partial (e.g. necessary but not sufficient – though almost sufficient) quasi-definition. It is doubtful whether much knowledge of anything is analytic in nature. Research into concepts constitutes one of the great triumphs of modern cognitive psychology. When it comes to theories of concept possession there are a range of live options, of which the 'classical' theory (definitional, necessary and sufficient conditions) simply isn't one.

Concrete examples and the purposes of epistemology

Epistemologists employ often fantastical thought experiments to test theories about nevertheless highly concrete objects of knowledge – concerning the ownership of motor cars, the existence of agricultural buildings, the presence of zebras. This type of knowledge is first-order, empirical, concrete-propositional. Critically, knowledge of each such proposition is pretty much self-standing;[3] it does not greatly embed within some larger framework of knowledge – as a proposition of philosophy, or some other discipline might – and knowledge of such propositions (though not, of course, the analysis of such knowledge) commonly relies on little more than basic perceptual/conceptual abilities on the part of the knower, substantially of the kind we might share with a dog or ape.

One should ask what the point is of concentrating on such knowledge – and should ask this question in two respects. First, does epistemology have something to contribute to the development of accounts of this kind of knowledge? Second, would developing satisfactory accounts of this kind of knowledge give us what we want when we do epistemology? I suggest the answer should be 'no' to both questions.

One precisely does not use epistemology in applied empirical contexts. Not since the time of the phenomenalists have epistemologists thought they could contribute to the investigation of first-order concrete object perception and the like. Rightly, one today uses scientific psychology to inquire into such matters. But consider now, what would happen if we had an exhaustive account of this kind of self-standing, first-order, empirical, concrete, propositional knowledge: would it even speak to our philosophical concerns – would it advance us philosophically any distance? To start to answer this question, let us go back to the levels distinction earlier – to deal with the question through raising an anticipated objection to that hierarchy of levels. The objection

---

3   The 'pretty much' clause here is a qualification: of course, to know that that's Jones' car isn't by any means entirely self-standing – even for Jones' dog, and much less for Jones' work-mate, Smith. Of course this proposition embeds with other things that Smith (or Rover) knows. The point, however, is that to be known, such a proposition doesn't need to embed within a framework of other knowledge to nearly the same extent that a proposition of physics, or philosophy, or computer programming, or politics, or any major discipline (or craft, or art, or game, or involved cultural practice) does. The scientific details covered by this caveat cash out in powerful connections with a very major tranche of the psychological literature; notably, the extensive literature concerning expert–novice reasoning, cited below.

anticipated is that this talk of theories of theories of knowledge is regressive. Won't we need another 'theory of ...' *ad infinitum*? Why begin going down this road? Can't a theory of knowledge apply to itself?

There is, though, every reason to think that a theory of first-order, concrete, propositional knowledge won't apply to itself. After all, our knowledge of *Jones*, and our identification of *that barn*, doesn't apply to itself – no-one seriously thinks of stopping the regress there. (Jones' dog may have that knowledge.) 'Knowledge' is a very different kind of object of inquiry than, say, *car ownership*; an adequate theory of concrete-object perception/identification is likely to be inapplicable to knowledge of theory-construction and theory-adequacy in epistemology. For example, very different types of cognition are likely to subserve each of these abilities. If we are lucky, there will come a time when we get a higher-order theory of knowledge, one that applies to other, lower-order aspects of our cognition, and is also reflexive – applies to itself. Certainly we should be striving for such a theory, but we will probably require at least one other 'level' before we have even a candidate for getting to that stage. We will need knowledge of first-order things (say, as here, of people, or of concrete objects), then a theory of knowledge of these first-order things, then a theory of knowledge of theories. Perhaps by this stage, if we are lucky, we may have a theory of knowledge that can apply to itself, since a theory of theories will also be a theory. But concentrating on the kind of first-order knowledge indicated – of *items* of information, states of affairs, as represented by largely self-standing propositions about concrete objects and the like – will not take us any distance towards this goal.

Why though, should we, as philosophers, want especially this kind of reflexive, abstract, higher-order knowledge? A short answer would be because that's just what philosophical knowledge is. Meaning, if we want to know about philosophical knowledge (and we do) it is this kind of thing we have to investigate, not something else. And if we want to have philosophical knowledge (whether about philosophical knowledge, or about anything else) we want to have this kind of thing, not something else.

That is, we ought to think about what our ends are in searching for a philosophical theory of knowledge. Epistemic inquiry is, of course, partly an end in itself; but even so, it is also in the service of other, characteristically *philosophical* ends. An example here is that many metaphysical inquiries and arguments – about the existence of God, the nature of value, the prospects for mathematical realism, the nature of mind, the reality of the external world – require epistemic premises at crucial junctures: concerning how or whether we can know, or rationally believe, in the objects of such inquiry. But philosophical inquiry requires more than isolated epistemic premises. The purposes and theories of the great historical epistemologies massively intertwine with the purposes and theories of metaphysics, philosophy of language, and value inquiry, in all kinds of theoretically powerful, richly structured and difficult-to-draw-out ways. Epistemology, where it is more than just a 'Gettier Game', is almost always at least incipiently metaphilosophical.

In the light of even such brief considerations, it should be seen as troubling how little the development of epistemic theories is shaped by an awareness of the need to apply such accounts to 'big picture' knowledge – whether philosophical knowledge, or any other knowledge involving propositions that require to be situated within the framework of a larger discipline or theory, including knowledge that is in any way abstract, or higher-order, or values-based. Notice here that we are still concerned, with the philosophical tradition, to inquire into propositional knowledge. It is just that the kind of propositional knowledge philosophy is concerned with, like so much propositional knowledge in general, needs to be situated within a larger body of top-down knowledge: of tacit or explicit theory.

Why, when we do not need and cannot employ epistemology for investigation into quotidian and concrete cases of first-order identification, recognition and recall, are epistemic theories nevertheless constructed against a testing-ground of such cases? A theory of such knowledge looks to be precisely of the kind that scientific psychology is best equipped to deliver. And the theories of such knowledge do not look to have much obvious ability to generalise to ourselves as possessors of *philosophical* knowledge – or, for that matter, knowledge of physics,

mathematics, art, literature, politics or architecture.[4] Investigation into self-standing, first-order, concrete, propositional knowledge looks to be pretty much irrelevant to our philosophical purposes,[5] and philosophical expertise looks to be pretty much irrelevant for purposes of investigation into such knowledge.

Data and intuitions: phenomenalism and analysis

We saw connections with phenomenalism in discussion of points falling under both the preceding headings. There is more that is worth saying about this movement, however – more that is relevant to the final framework assumption.

The idea of phenomenalism was that certain things just had to be the case (on pain of regress): a bottom-up associationism, sensory simples as bedrock data given to us by acquaintance/ introspection, and the requirement that we give an analysis of perceptual knowledge – of objects – in terms of these simples, whether a logically reductive analysis (phenomenalism) or a psychologically reductive analysis (psychological structuralism).[6]

Phenomenalism was rejected for at least the following two good reasons. First, neo-Kantian scientific psychologists[7] reminded the philosophers of what they should already have appreciated, that a bottom-up associationism wasn't conceptually necessary – indeed, that it was (substantially) empirically false: face-recognition, object-recognition, concept-formation, problem solving, memory, and most other aspects of cognition (especially high level cognition) involve top-down processing to a greater or lesser extent. So, complex psychological items are not typically 'glued together' out of simpler items in a logical, associative, summation or synthesis – one that might, at least in principle, have permitted a reverse-process of analysis or definition. Second, the phenomenalists' sensory 'simples' (their 'data') weren't *simple* at all. On the one hand, underlying them were unconscious things that were simpler still; and on the other hand, the simplest conscious things weren't by any means typically sensations: they were whole objects perceived, meaningful memories recalled, and persons recognised. These putatively 'simple' data, then, were typically the *more* complex outcome of artificial and convoluted processes of abstraction from genuinely simpler normal-object-recognition, remembering and the like, by philosophers in the grip of a theory, attempting an armchair reconstruction of processes that were not at all amenable to such reconstruction: processes whose psychological nature and normative-epistemic status they in fact could have no *access to* in this way. Their 'simplicity', their putative status as 'data', was illusory, an artifact.[8]

---

4 An impressive psychological literature on expertise has had much to say about the nature of knowledge within such disciplines. I believe that epistemologists should learn from this literature, with its emphasis on 'big picture' knowledge; on knowledge as theory-embedded, inferentially interconnected, domain-specific and top-down, e.g. Larkin et al. (1980), Chi et al. (1981), Schoenfield & Herrmann (1982), Weiser & Shertz (1983), Lesgold et al. (1988), Hassebrock et al. (1993), Lane et al. (2000), Gobet & Chassy (2009).

5 Excepting, possibly, where the purpose is solely an investigation into problems of skepticism. But the Gettier and post-Gettier literature in epistemology is not primarily, much less solely, concerned with refuting skepticism; it is (rightly) concerned with inquiry into knowledge as a philosophical task of importance in itself. (Though some theories of knowledge claim also to offer a solution to problems of skepticism, some don't; and even those that do, work through analyses of self-standing, concrete, first-order knowledge in addressing both tasks.) Note further, that there are problems of skepticism for abstract, higher-order, theory-embedded knowledge also; indeed, many of the respects in which problems of skepticism really matter to us, precisely concern such knowledge – e.g. atheism, moral skepticism, doubts about the reality of mental properties and agency.

6 Goldman (2007: pp. 14–15) has interesting remarks about the relation between phenomenalism and the metaphilosophical 'intuitions' program. Surprisingly, given his first-order epistemic commitments, he endorses a notion of intuitions as constitutive of metaphilosophical 'evidence' (his, revealing, term) for the presence of a concept: a constitutive relation that he himself compares with the constitutive relation phenomenalism posited between perceptual appearances and existence of an object. (He does not, of course, defend the latter – as opposed to its meta-epistemic analogue.)

7 I am thinking especially of the Gestaltists, Bartlett and Piaget, though by the time these European neo-Kantians were eventually assimilated into American psychology (via, e.g., Neisser, Bruner and Bransford) most traces of philosophical provenance had been lost.

8 So, this was the Gestaltist objection to the Structuralists' idea of the 'stimulus error' (Titchener 1912: the structuralists' idea that we don't really see the banana [stimulus] *in strictu dictu* we see rather the sense data that comprise it – though trained, 'expert' introspection may be needed to uncover these fundamental data.) This objection was much later taken up by Wittgenstein, with whom analytic philosophers often believe it originated (e.g. Anscombe 1981: pp. viii–ix). Note, even conceptual analysts who oppose according special epistemic status to experts' intuitions (and most of them do oppose this – e.g. Goldman, against Kornblith – Goldman 2007: p. 21ff; Goldman & Pust 1998: p. 187) nevertheless hold that analysing the Folk's intuitions requires major philosophical expertise and training.

Phenomenalism applied to first order knowing – analyses of concrete objects, for example – whereas analytic theories of knowledge offer analyses of knowledge; nevertheless, consider the parallels: we have to offer an *analysis* (of knowledge). The bedrock data of that analysis, its *input* (Bealer 1996a: p. 131, 1996b: p. 20) has to be our *intuitions* ('what else, in the limit, could one have to go on?') Whatever 'saves' or 'captures' the greatest number of these intuitions is the most nearly adequate epistemological theory.

The intuitions that supposedly ground our concept of knowledge are normally, on this tradition of inquiry, treated as basic. Q: Why is Theory of Knowledge θ wrong? A: Because it does not yield results that accord with our intuitions when tested by these thought experiments. But what is the motivation for these intuitions? (Why should we accept these intuitions?) Such questions will be treated as category mistakes – because they must be asking for further intuitions in turn. This is, tacitly, to employ the stock regress argument for foundationalism: there can't be an infinite chain of justifying reasons, some must be bedrock; call these – whatever they are – our 'basic', 'fundamental', 'deepest', 'pre-analytic', 'pre-theoretical' intuitions.

> [M]any alleged rational or *a priori* insights are, if not necessarily *mysterious*, at least *inexplicable* in the sense of being apparently *irreducible*: they are apparently incapable of being reduced to or constituted out of some constellation of discursive steps or simpler cognitive elements of some other kind. But ... any such reduction would have to appeal to other apparent insights of a similar sort and thus ultimately, if an infinite regress is ruled out, to irreducible ones … (Bonjour 1998: p. 108, emphasis in the original).

And how do we *know* we have these intuitions and not others? This question too must be rejected as resting on a category mistake: we sit in the armchair, present ourselves with a thought experiment, dwell inwardly on it until satisfied, and then just *have* said intuitions. They're given to us, known by acquaintance, self-presenting. One's intuitions, on this conception of philosophical thought, are treated very much as the (prior) data one tests the theory against, the fundamental basis for such a theory, a 'basic evidential source' (Goldman and Pust 1998).

> Our question has been whether intuition can be understood clearly and defended adequately as a source of foundational a priori justification. Once that is accomplished, our task is completed (Sosa 2007: p. 106).

The trained philosophical introspector judges a putative analysis of knowledge against the nuanced feelings of satisfaction or nagging feelings of dissatisfaction he feels at a case being classified thus-and-so. The theory that, on balance, causes the most net satisfaction, or least dissatisfaction, wins. A curiously *philosophical* empiricism[9] this – not, as for phenomenalism, an explicit framework theory of what first-order knowledge would have to consist in; rather, a tacit framework theory of philosophical knowledge: what epistemic knowledge would have to consist in. A foundationalist meta-epistemology, a First Philosophy of Intuitions.

Given that we would no longer accept anything like such a position as a first-order account of knowledge, why should we tacitly accept this as a higher-order account – of how we must come by and represent the knowledge of what *knowledge itself* is? The earlier objection to the sensory simples of phenomenalist legend was that the only reason to suppose their existence was the undeniable prior presence of concrete-object perception, allied to the undefended and equivocal framework assumption of this being either quite impossible (psychologically) or quite unwarranted (normatively) without such simples. In turn, though, concerning the 'intuitions' of epistemology: what reason is there is to postulate these as basic, other than the undeniable prior presence of conclusions drawn by rational, philosophical, thought, and the undefended and equivocal framework assumption that this process would be either quite impossible

---

9   Not, that is, by any means a rationalism worthy of the name, despite the armchair methodology and the convictions of many who would defend it – something that makes all the more dispiriting the tendency to the view that the analytic tradition faces something like a choice between this and naturalism.

(psychologically) or quite unwarranted (normatively) without such prior intuitions? Clearly, one may inquire into the motivation for such intuitions – this is no kind of category error – and equally one may inquire into how one knows that one has these intuitions. The motivation for intuitions need not require further intuitions ('intuitions' as ontologically real items that we have a prior, more-basic handle upon). It may, for instance, require a structure of theory connected either by tacit inference or explicit argument.[10]

That is, only on certain very empiricist, atomistic, reductive and hypostatic assumptions must the search for reasons be tacitly the search for items ('intuitions'). Talk of 'intuition' may perhaps be retained in philosophy, but only as a loose, meta-epistemically unfreighted locution. Intuitions are not bottom-up 'pre-theoretical' *inputs* to be 'saved' or 'captured' in an analysis – Sosa's '*source* of foundational a priori justification' (Sosa loc. cit.).

To the extent, then, that we wish to retain talk of intuitions, it should be acknowledged that these are more the outcome of (tacit) argument/inference than the input to it; and more usually still, as holistically embedded within, or emergent out of, a (partly) tacit, framework, top-down structure of argument, inference, knowledge. In any robust, ontologically committed sense, *we do not have 'input' intuitions at all*.

In case this seems a jarring assertion, note the tendency to ambiguity[11] in the various claims that philosophy 'comes down to intuitions in the end' – a tendency that may account for why this view has itself seemed 'intuitive' to so many. On the one hand, the metaphilosophical role of intuitions may be defended by reference to a particular epistemic[12] theory as to these intuitions' unique role as a terminus of justification – a specific metaphilosophical account of the kind identified and attacked in this paper. (Such a contentful, committed, sense of 'intuitions' will be what considered, published defenses of intuitions articulate and explicitly argue for.) However, particularly in the wider philosophical community, 'intuitions' tend to be tacitly operationalised as 'whatever terminates philosophical argument'. If the former disambiguation is embraced, the view that 'it all comes down to intuitions in the end' is, I have argued, an indefensible meta-epistemic foundationalism; if the latter, the view is viciously circular.

Top-down objections to the myth of 'contamination' by theory
In the literature one encounters talk of 'pre-theoretical' intuitions *passim*, with both Goldman (2007) and Bonjour (1998: p. 102) revealingly glossing these as intuitions that are 'uncontaminated' by theory. This very much recapitulates (even in the language used: 'contaminated') older debates in scientific psychology – especially between Bartlett and Ebbinghaus about the role of meaning in memory, and between Titchener and the Gestaltists on the notion of the 'stimulus error'. Theory does not 'contaminate' intuitions – any more than meaning 'contaminated' memory, or seeing a stimulus as a whole real object rather than an inference from a collection of sense data was a 'stimulus error'. The notion of an 'uncontaminated' (by meaning) memory was a Chimera, a framework error, with the Ebbinghausian bottom-up research program – into memory as a pure faculty of retention – at least a profound mistake of emphasis. In

---

10   The positive view I hold in opposition to the metaphilosophy criticised here is that philosophical understanding is top-down not bottom-up, and that 'intuitions' (if you must) are the outcome of tacit and explicit cognition (or at least, cognitive structures), not the input to it. To once again offer a glimpse of the background, positive, view: propositions, positions, insights, knowledge in philosophy, is and are situated within major, powerful, articulated, inferentially connected organised knowledge structures (neo-Bartlettian schemata). Philosophical propositions, positions, claims, insights ('intuitions' if you must) have no separate, prior, primitive existence. The schemata that embed such philosophical knowledge do not have an analytic structure; and though we may be said to have philosophical intuitions, we do not have atomistic bottom-up, 'primitive' intuitions as 'data' to be 'saved' or 'captured' by such analyses – via, say, some process of accounting, of associative summation and negation.

11   As numbers of philosophers have noted, talk of 'intuitions' in analytic philosophy is pervasive, yet chronically vague and ambiguous. For a lot more than I have space for on the equivocal nature of 'intuitions talk' see Cappelen (2012), also Williamson (2007). Often, it is a harmless locution indicating simply a philosophical belief, or sometimes a staring point for argument, or as a *pro tanto* justified position that nevertheless may require revision in the light of other considerations. I am concerned to chase to earth a more metaphilosophically committed conception of 'intuitions' than these.

12   Sometimes not only an epistemic theory. If intuitions are constitutive of meta-epistemic justification we will have also a metaphysical (hence metaphilosophically anti-realist) theoretical framework here. (The distinction between 'constitutive' and 'non-constitutive' accounts of metaphilosophical justification is due to Goldman [2007] – who himself defends a constitutive theory.)

philosophy *rational insights* are developed by philosophers who work hard to situate their thought within large, organised, theoretical structures (schemata). The philosophers' top-down knowledge does not 'contaminate' their otherwise pristine input intuitions.[13]

This framework of assumption – as to the 'contaminating' role of theory and the consequent need for 'pre-theoretical' (Bonjour 1998: p. 107 'direct and immediate') intuitions – is carried over wholesale into experimental philosophy, with the sole difference being that experimental philosophers seek to measure these supposedly bottom-up intuitions with Likert scales rather than $n = 1$ armchair introspection. Such work typically uses naïve (non-philosophical) subjects. Asking questions of *philosophers* is, though not unheard of, usually seen as problematic: philosophers will seek to make their responses internally self-consistent and consistent with the constraints of prior theory. What the experimental philosophers are seeing to avoid here is known to the constructors of attitude surveys as a *consistency bias*. These 'pre-theoretical' intuitions are seen as our *raw* data and the effect of an effort after consistency here in non-naïve subjects is seen as introducing measurement error – much as, for the structuralists, reporting one's supposedly prior sensory data as *interpreted* (stating one sees a banana, which this tradition of thought saw as an *inference* out of said data) was seen as a *stimulus error*. To one with an awareness of the history of scientific psychology, the psychological naivety of these approaches – armchair or experimental – is painful to behold. Theories – elaborately articulated, stored (including culturally stored) knowledge-structures driven from the first by *deep efforts after consistency* – these are at the core of all philosophical thought: our outcome 'intuitions' (if you must) are constitutively bound up with such structures. Theory, effort after consistency, stored prior top-down knowledge, these are not *contaminants* of a collection of 'immediate', 'prior', 'pre-theoretical' intuitions: such intuitions only make sense within (are emergent out of) the structure of such theories. Here, Daniel Kahneman makes a methodological recommendation (in the context of the distinct heuristics and biases literature) that makes this 'consistency bias' assumption explicit. He situates it within Two Systems Theory – though he and others had made this point before Two Systems Theory existed.

> Different research designs are appropriate for the study of System 1 and of System 2. If the goal of the research is to study intuitive judgment, the design should minimize the role of deliberation and self-critical reflection. Intuitive judgments and preferences are therefore best studied in between-participants designs and in short experiments that provide little information about the experimenter's aims. … Within-participant factorial designs are particularly undesirable because they provide an unmistakable cue that any factor that is varied systematically must be relevant to the target attribute (Kahneman 2003: p. 712).

Among theorists working within the reasoning and biases literature, *efforts after consistency* are seen as a System Two (rational) product, whereas more immediate responses made without such efforts are seen as a System One product. Many of the most famous putative *failures* of rationality in this literature are seen as being the product of insufficient slow, effortful searches after consistency ('Linda the feminist bank teller' etc.). Whatever your thoughts about Two

---

13   The idea that theory does not 'contaminate' our otherwise pure intuitions is an insight that, commendably, is endorsed by Kornblith (2007), though Kornblith does assimilate this to a separate point as to a greater evidential weight needing to be attached to *experts'* intuitions. (There is a sense of 'expertise' – see fn. 3 & 4 above – in which I would agree with that claim, but I am doubtful whether this is the sense Kornblith intends and am certain it is not the sense in which the greater philosophical community will interpret his remarks.) Bartlett (1932: p. 4 ff), in regard to memory, made the following 'top down' criticisms of Ebbinghaus and his 'contamination' (/'transmission', /isolated faculty of retention) bottom-up model of cognition. Among his three 'stimulus side' and three 'response side' criticisms Bartlett noted that:

(1)   It is impossible to rid stimuli of meaning (here: rid 'intuitions' of theory).

(2)   The effort to do so creates artificiality (compare the Gettier literature).

(3)   What counts as a 'simple' response – an 'uncontaminated' memory (here: a 'pre-theoretical intuition') is a highly loaded issue. For one thing, it varies with the context of inquiry. For another, and of central importance: isolation of the response in a highly controlled experimental (here: thought experimental) situation cannot be assumed to operationalise simplicity of that response (here: fundamentality of those 'intuitions') in any independently motivated sense of 'simple'.

Bartlett's overall conclusion is that there is no reason to think complex responses are just built up out of simpler. Bartlett's insights utterly transformed cognitive psychology; these insights need also to be understood, then reflexively applied, by epistemologists, I would contend.

Systems Theory (and I am not an uncritical admirer) ask yourself whether, in the language of this theory, you think philosophical conclusions should be, and generally are, more the outcome of System One or System Two. That is, do you think that philosophical reasoning should and does associatively preserve the more or less immediate outcome of fast, non-rational, automatic, computationally economical, evolutionarily old, heuristic, parallel processes (processes that 'minimise the role of deliberation and self-critical reflection')? Or is and ought philosophical reasoning to be based on slow, analytic, computationally expensive, evolutionarily new, distinctively human, culturally and educationally laden, rational, serial processes? Merely to ask the question is to answer it.

A blunter form of attack against the ontological (and in this context *also* epistemic) status of the *elements* of this foundationalism can be inspired by the attitude of the behaviorist J.B. Watson to the Psychological Structuralism inflicted on him as a student. Watson said simply that he never could see the things (the basic, isolated, 'input' sensations) his introspections-based teachers required him to see – in their experiments involving the perception of concrete objects. In turn, I cannot see the things (the 'fundamental', 'basic', 'pre-theoretical' intuitions) the intuitions-based meta-epistemologists require me to see in their thought experiments – and I don't think the things I'm supposed to see are there.[14] In constructing a theory of knowledge, one should not have to be answerable to a tacit meta-epistemology of preserving one's audience against having their self-declared basic, foundational, bedrock intuitions violated. There are no such basic intuitions – any more than there were basic sense data in perception. Making this denial no more impugns realism about the conclusions of our philosophical reasoning – our 'rational insights' – than the equivalent position regarding sense data impugned realism about the objects of our perception. In fact, repudiating such a [meta] epistemology is likely a precondition of realism in each case.

Meta-epistemic conclusions
Realism/anti-realism
There is knowledge – what it *is* – and then there is our concept of knowledge (representation of, knowledge of). We do not gain dependable access to what knowledge is from our concept thereof. And we do not even have immediate access to what our concept of knowledge itself is, merely by examination of our intuitions and testing these with thought experiments. This is basic realism in any area: a separation of constitutive from epistemological issues (here, in regard to epistemology – in regard to knowledge).

There are parallels between the nature of our target and certain characteristic features of anti-realism as found in other areas of philosophy over the last century – notably, as seen already, as regards empiricist idealism. Although there may be other ways to derive one's metaphysics from one's epistemology, a very common way is by reduction (to some conception of the epistemic data) and definition (in terms of these data). Regarding the latter, we saw that there is no reason to suppose that our concept of knowledge will be represented analytically (as a set of necessary and sufficient conditions, or even something close to this) – indeed, everything we know about concepts suggests it is not. Regarding the former, there is no reason to think that there will be either a justificatory or a psychological terminus of basic intuitions to serve as simples in such an analysis – indeed, everything we know about fallacious regress arguments for foundationalist atomism in philosophy/psychology suggests there are no such simples.

The constitutive question as to the correct theory of knowledge is not answerable to a jury of what we conceive/judge/perceive/'intuit' knowledge to be, and this in turn – our conception of

14 The situation here is in marked contrast with plausible candidates for the status of genuinely *a priori* knowledge. (I am certainly no enemy of the *a priori*.) In contrast, numbers of the most egregiously contrived Gettier cases approach parody – with even their defenders acknowledging that intuitions in response to such cases are divided and hard to fix. The same is true of certain metaphilosophically similar areas (highly involved Frankfurt cases, for example, some of which are hard for even subject-specialists to comprehend, much less adjudicate). But to emphasise, I don't think even the simpler Gettier cases generally and with modal force indicate any lack of knowledge on the part of the true believer they stipulate into existence: that is, I don't see that they do; and I don't think what I'm supposed to see is there.

knowledge – is not something that can just be 'read off' pellucidly from the armchair. It too is a substantial constitutive question.[15]

Consensus and disagreement

Why doesn't philosophy lead to consensus? The metaphilosophical 'intuitionist' will have the same sort of resources available to answer this question as did first-order empiricists such as Ayer in regard to issues of disagreement in ethics, or Hume in regard to disagreements about taste (Ayer 1936: ch. 6; Hume 1987). On the one hand, either or both of the dissenting philosophers may have failed to work through the analysis from their metaphilosophical intuitions [/emotions, /tastes] rigorously enough. Or, the philosophers may just not share some core intuitions. Proponents of this neo-Humean viewpoint tend toward the optimistic empirical assumption that our core intuitions are more widely shared than we think; and that very much of the disagreement may be resolved by one or both parties going back to inquire into the logical edifice constructed out of these intuitions – say, employing 'rephrasal strategies' (Bealer 1996a, 1996b).[16] Nevertheless, some disagreement will be stark: genuinely competing core intuitions with nothing more to say.

> [W]here there is such a diversity in the internal frame or external situation as is entirely blameless on both sides, and leaves no room to give one the preference above the other; in that case a certain degree of diversity in judgment is unavoidable, and we seek in vain for a standard, by which we can reconcile the contrary sentiments (Hume 1987: p. 244).[17]

Weinberg entitles this a lack of 'dialectical robustness', which property requires 'that the evidence we cite to each other be evidence that we can recognize the force of. But intuitions are damnably subjective'[18] Weinberg (2006: p. 29). In an obviously question-begging move, when pressed hard on their metaphilosophical views, this species of intuitionist will commonly come to the conclusion that we differ on core intuitions about how philosophy is done and have it that further discussion is fruitless. Often they do this, however (this drawing of a line under discussion), whilst giving the impression that they doubt whether we have honestly

15 This is to say, in a forced choice between experimental philosophy and conceptual analysis, I would choose the former every time. Fortunately, we are not in such a cleft stick. The ground shared by these two opposed positions is huge – and is precisely what should be called into question. Each is a species of reductive, foundationalist naturalism, in which 'all we have to go on' in doing philosophy is a metaphilosophical foundation of 'what [fundamentally] seems right to us'. Experimental philosophy then goes scientifically realist on finding out what this is. I am heartened to discover that I am not alone in my judgment here: 'Methodological rationalists (such as Bealer, Pust, and Bonjour) and experimental philosophers (such as Stich, Weinberg and others) are, as I see it, joined at the hips by a commitment to [the centrality of intuitions]' Cappelen (2012: p. 19). I demur somewhat from Cappelen's judgment that Bealer, Pust and Bonjour's use of intuitions as evidence makes them rationalist, however – notwithstanding their self-conceptions – cf. fn. 9 above.

16 Bealer is notoriously optimistic that when it comes to conflicting philosophical intuitions 'most apparent conflicts can be reconciled by standard rephrasal techniques' (Bealer 1996b: p. 9). He thus denies Weinberg's claimed datum (below) that in philosophy there is a lack of 'dialectical robustness' – or in Bealer's own terminology, he maintains that most of the concepts philosophy deals in are 'semantically stable'. I agree with Bealer that rephrasal may diminish disagreement in philosophical disputes, and that intuitions are often less 'damnably subjective' than appears at first sight. However, Bealer's degree of optimism is simply false to the facts. I also think that 'rephrasal' is far too small a notion for what takes place in philosophical reasoning; but above all that Bealer's intuitions – described as 'inputs' (Bealer 1996a: p. 131, 1996b: p. 20) to a process of systematisation – are vastly too bottom-up. I think quite a lot of *intra schematic* disagreement may yield to Bealer's 'standard rephrasal techniques'. Some extra schematic disagreement may yield, to a degree, also: but much extra schematic disagreement will lead to a Kantian 'dialectic' at best – and perhaps incommensurability at worst (n.b. I am not circularly defining intra- and extra-schematic disagreement in terms of dissent or assent). *We do not have the analytic philosopher's 'input' intuitions at all: philosophical and metaphilosophical views ('intuitions', if you must) are the sophisticated outcome of structures and processes of reasoning, argument, inference.*

17 I was prompted to re-acquaint myself with this striking passage by Aaron Meskin's use of it (in a talk) in its original, aesthetic, context.

18 Both the philosophical and psychological behaviorists would recognise instantly Weinberg's 'subjectivity' objection, having used its equivalent passim against their structuralist (phenomenalist) opponents, whether in highly sophisticated or crudely vigorous form. The behaviorists' assumption that one must accept pretty much the whole bottom-up framework as introspectionist structuralism, to disagree solely on the issue of whether the psychological elements of one's atomistic account (one's data) were subjective (e.g. qualia) or objective (e.g. S-R connections) is not something that fared at all well in the longer-term of course (*vide* the cognitive revolution's rediscovery of the more rationalistic, structured, top-down, anti-associative tradition of Bartlett, Piaget and the Gestaltists). I suggest that we should not expect this assumption's epistemic equivalent to fare any better. It is the bottom-up nature of this meta-epistemic framework that is wrong: accepting intuitions as foundational basic givens to be 'saved' in our philosophical accounts. This latter assumption is not something we should leave unquestioned only to dispute these data's status as subjective (mainstream conceptual analysis) or objective (experimental philosophy).

scrutinised our intuitions (here: about the role of intuitions) as rigorously as perhaps we ought. If *what knowledge is* is a question that must be constrained by our conception of knowledge, and if this in turn really is answerable to our intuitions, and if these in turn really are basic – the bottom-up, unanalysable, input, foundation of an analysis – then this approach is the best we can do. But none of this metaphilosophical framework is defensible, and by rejecting it we can aspire to a maturer, more optimistic and more realistic conception of disagreement and consensus in philosophy.

Diagnosis

Early analytic epistemology was assumed to have necessary application to first-order cognition – e.g. concrete object perception. It was empiricist, foundationalist, associationist, internalist (that is, mentalist/introspectionist-internalist) and idealist. The (tacit) meta-philosophical framework within which it was constructed would naturally have tended to be congruent with its first-order theories – though not deliberately, reflectively, so. Then things changed. This program was subjected to sustained attack; stock arguments developed against the cluster-position that was classic empiricist, foundationalist, associationist, internalist and idealist epistemology.[19] Only the meta-epistemology remained tacit, as meta-philosophy usually is. It therefore didn't change. No-one noticed that counter-arguments that were common currency in first-order epistemology might apply to the metaphilosophy also.

Our meta-epistemology should not be assumed to be foundationalist, empiricist and analytic if we have, as most of us think we have, knock-down arguments against our first-order epistemology being this entwined cluster of things. We need a maturer meta-epistemology. This paper has been intended to clear a space for its development.[20]

References

Anscombe, G.E.M. 1981. *Metaphysics and the Philosophy of Mind (The Collected Philosophical Papers of G. E. M. Anscombe)*, Vol. 2. Minneapolis: University of Minnesota Press.

Ayer, A.J. 1936. *Language, Truth and Logic*. London: V. Gollancz.

Bartlett, F. 1932. *Remembering*. Cambridge: Cambridge University Press.

Bealer, G. 1996a. '*A priori* knowledge and the scope of philosophy', *Philosophical Studies* 81(2–3), 121–142.

Bealer, G. 1996b. 'On the possibility of philosophical knowledge', *Noûs* 30(Supplement: *Philosophical Perspectives* 10, Metaphysics), 1–34.

Bonjour, L. 1998. *In Defense of Pure Reason*. Cambridge: Cambridge University Press.

Bruner, J., Goodnow, J., Austin, G. 1956. *A Study of Thinking*. New York: Wiley.

Cappelen, H. 2012. *Philosophy Without Intuitions*. Oxford: Oxford University Press.

Chi, M.T., Feltovich, P.J., Glaser, R. 1981. 'Categorisation and representation of physics problems by experts and novices', *Cognitive Science* 5(2), 121–152.

Foley, R. 2004. 'A trial separation between the theory of knowledge and the theory of justified belief', in Greco, J. (ed.), *Ernest Sosa and His Critics*. Malden: Blackwell, pp. 59–71.

Gobet, F., Chassy, P. 2009. 'Expertise and intuition: a tale of three theories', *Minds and Machines* 19(2), 151–180.

Goldman, A. 2007. 'Philosophical intuitions: their target, their source, and their epistemic status', *Grazer Philsophische Studien* 74, 1–26.

Goldman, A., Pust, J. 1998. 'Philosophical theory and intuitional evidence', in DePaul, M.,

---

19 The stock arguments I am referring to, of course apply against this classic cluster-position as a whole, not to the components of this position considered separately.

20 This paper is concerned to pave the way for a positive, and hopefully *reflexive*, position that violates the framework identified above – in the teeth of a tradition that would prevent this position from having a right to be heard. That is the sole aim of this paper, not obtaining a terminological victory over 'analytic epistemology' to the extent that this term describes something which leaves the right to develop such an alternative theory untouched. The three tacit presuppositions identified above are what have to be got out of the way in order for that new theory to get a hearing, and I would maintain that they are still embraced as framework assumptions by a large majority of analytic philosophers constructing theories of knowledge within normative epistemology. If some don't embrace these assumptions, then fine. If a given metaphilosophy sanctions the three assumptions above, it is challenged by the arguments given; if it does not, it is no target of these arguments.

Ramsey, W. (eds), *Rethinking Intuition: the Psychology of Intuition and Its Role in Philosophical Inquiry*. Lanham: Rowman & Littlefield, pp. 179–200.

Hassebrock, F., Johnson, P.E., Bullemer, P., Fox, P.W., Moller, J.H. 1993. 'When less is more: representation and selective memory in expert problem solving', *American Journal of Psychology* 106(2), 155–189.

Hume, D. 1987. 'Of the standard of taste', in *Essays, Moral, Political, and Literary* (revised edn), Miller, E. F. (ed.). Indianapolis: Liberty Classics.

Kahneman, D. 2003. 'A perspective on judgment and choice: mapping bounded rationality', *American Psychologist* 58(9), 697–720.

Kornblith, H. 2007. 'Naturalism and intuitions', *Grazer Philosophische Studien* 72, 27–49.

Lane, P., Gobet, F., Cheng, P. 2000. 'Learning-based constraints on schemata', in Gleitman, L.R., Joshi, A.K. (eds), *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society.* Mahwah: Erlbaum, pp. 776–781.

Larkin, J.H., McDermott, D.P., Simon, H.A. 1980. 'Expert and novice performance in solving physics problems', *Science* 208(4450), 1335–1342.

Lesgold, A., Rubinson, H., Feltovich, P., Glaser, R., Klopfer, D., Wang, Y. 1988. 'Expertise in a complex skill: diagnosing X-ray pictures', in Chi, M.T.H., Glaser, R., Farr, M.J. (eds), *The Nature Of Expertise.* Hillsdale: Lawrence Erlbaum.

Murphy, G., Medin, D. 1985. 'The role of theories in conceptual coherence', *Psychological Review* 92, 289–316.

Rosch, E. 1978. 'Principles of categorization', in Rosch, E., Lloyd, B. (eds), *Cognition and Categorization*. Hillsdale: Erlbaum.

Rosch, E., Mervis, C. 1975. 'Family resemblances: studies in the internal structure of categories', *Cognitive Psychology* 7(4), 573–605.

Schoenfeld, A.H., Herrmann, D.J. 1982. 'Problem perception and knowledge structure in expert and novice mathematical problem solvers', *Journal of Experimental Psychology: Learning, Memory and Cognition* 8(5), 484–494.

Smith, E., Medin, D. 1981. *Categories and Concepts*. Cambridge, Massachusetts: Harvard University Press.

Sosa, E. 2007. 'Experimental philosophy and philosophical intuition', *Philosophical Studies* 132(1), 99–107.

Titchener, E. 1912. 'The schema of introspection', *American Journal of Psychology* 23(4), 485–508.

Weinberg, J. 2006. 'What's epistemology for? The case for neopragmatism in normative meta-epistemology', in Hetherington, S. (ed.), *Epistemology Futures*. Oxford: Oxford University Press, pp. 26–47.

Weiser, M., Shertz, J. 1983. 'Programming problem representation in novice and expert programmers', *International Journal of Man-Machine Studies* 19(4), 391–398.

Williamson, T. 2007. *The Philosophy of Philosophy.* Oxford: Oxford University Press.

Wittgenstein, L. 1953. *Philosophical Investigations*, transl. Anscombe, G.E.M. Oxford: Blackwell.