Intuitions in Epistemology: Towards a Naturalistic Alternative

Kristoffer Ahlstrom

Philosophy Department, University of Massachusetts at Amherst

The present paper revisits the main methodological problems with conceptual analysis and considers two attempts to rectify them in terms of prototypes and reflective equilibria, respectively. Finding both wanting for the purposes of epistemological analysis, a naturalistic alternative is then sketched that explores the positive implications of aforementioned problems for the demarcation of the respective roles of intuitions and empirical investigation within three epistemological domains, viz., the evaluation of epistemological hypotheses, the amelioration of epistemic practices, and the construction of a theory of epistemic value.

Keywords: methodology, intuitions, conceptual analysis

1. The Dominant Approach

On the dominant methodological approach in philosophy, philosophical analysis is identified with conceptual analysis and candidate analyses evaluated by determining whether they are susceptible to intuitive counterexamples. Presently, there is no need to suppose that all philosophy is conducted on this model. Some endow intuitions with a great evidential weight (e.g., Kripke 1980); some take a more moderate standpoint (e.g., Lewis 1973); some defend the use of intuitions, but argue that probing them should be construed not so much as an inquiry into concepts as an exploration of *a priori* truths (Sosa 2007); and some claim that philosophical analysis should not be concerned with intuitions at all (Kornblith 2002). For present purposes, however, it suffices to note that a heavy methodological reliance on categorization intuitions constitutes the rule rather than the exception in contemporary analytic philosophy. The Gettier discussion constitutes a particularly clear example of intuition-driven philosophy (Shope 1983, Gettier 1963), as

Corresponding author's address: Kristoffer Ahlstrom, Department of Philosophy, University of Massachusetts, Amherst, MA 01003, USA. Email: kahlstro@philos.umass.edu.

Studia Philosophica Estonica (2009) **2.2**, 15–34
Published online: February 2010
Online ISSN: 1736–5899
www.spe.ut.ee
http://dx.doi.org/10.12697/spe.2009.2.2.02

does the literature on causal theories of meaning (Kripke 1980, Putnam 1975) and discussions about personal identity (Parfit 1984, Thomson 1971). In fact, this practice of supporting and refuting philosophical analyses in terms of definitions with reference to categorization intuitions and the concepts that they are taken to reveal is so widespread that it has been referred to as part of "the standard justificatory procedure" in philosophy (Bealer 1998).

1.1 The Desiderata of Philosophy as Conceptual Analysis

On this procedure, philosophy inquires into concepts by defining the meaning of linguistic terms, most likely against the background of an assumption to the effect that what a subject means is a function of the concepts that she possesses; an assumption that, furthermore, finds support in the fact that an impressive amount of the phenomena discovered in the empirical work on conceptual and non-verbal classification also turn up in the corresponding linguistic tasks (Murphy 2002).[1] Moreover, the favored format for capturing meanings is that of *necessary and sufficient application conditions*. In fact, philosophers typically are not asking for just any necessary and sufficient conditions, but for conditions that have *clear-cut* boundaries, so as to ensure that category membership is a straightforward yes-or-no affair (Laurence and Margolis 2003). Let us refer to the desideratum of identifying such clear-cut, necessary and sufficient conditions as *exactitude*.

As already noted, common practice typically also requires that the definitions provided via such conditions do not admit of any genuine intuitive counterexamples (i.e., counterexamples that cannot be explained away). Let us refer to this desideratum as *exhaustiveness*. The assumption behind this desideratum is that the possession of a concept involves a tacit (albeit potentially incomplete) knowledge of its defining conditions, to be captured by our philosophical theories (Ramsey 1998).[2] Being tacit, the content of that knowledge is not something that may be straightforwardly produced by any competent user of the corresponding term. Still, in so far as categorization

[1] This is not meant to suggest that there is a neat one-to-one mapping between concepts and (meaningful) words—a claim that would be incompatible with ambiguous words, synonyms, and concepts that do not have a word (in some cases, not even a string of words) to go with them—but merely that there is a (non-unique) mapping from concepts to meanings. It also is not to say that conceptual contents are determined by cognitively internal factors, or that all aspects of what a word or sentence *communicates in use* is a function of concepts possessed.

[2] A similar assumption is driving the Chomskyan tradition in linguistics, where the intuitive judgments of speakers are taken as linguistic evidence—on some readings the *only* available evidence—for particular sets of rules and principles of the speaker's language. See (Samuels et al. 1999) for a discussion of the analogy between the linguistic and the conceptual analytic case.

intuitions are the products of the concepts possessed, a subject's categorization intuitions may serve as (fallible) guides to her concepts. Consequently, probing the categorization intuitions of competent users of a term may serve to elucidate the defining condition that attaches to the corresponding concept.[3]

As such, the elucidation of concepts plays a dual methodological role. On the one hand, it provides positive material for the *discovery* of candidate analyses, i.e., candidate necessary and sufficient conditions that are clear-cut, under the requirement of exactitude. On the other hand, it also serves a *justificatory* role in the evaluation of such analyses, not by necessarily contributing incorrigible or indubitable evidence, but by providing strong reasons for disqualifying analyses in so far as they either include counter-intuitive instances or fail to include intuitive instances, under the requirement of exhaustiveness. Hence, the method of contemporary conceptual analysis.

1.2 A Problem for Exactitude

One well-known problem for conceptual analysis understood thus is that the classificatory structures unveiled by psychologists studying human categorization display properties incompatible with the idea that concepts should be represented in terms of clear-cut conditions (Murphy 2002). The categorizations we, in fact, make reveal a taxonomy where members of a category form a *continuum*, to the effect that some members are deemed to be "better" or more typical instances than others, some less typical members are considered to not be all that different from some non-members, and there, hence, seldom will be any sharp boundaries between different categories. Consequently, the conditions called for by exactitude simply do not seem to be of the right kind if we are interested in capturing human concepts.

Why have philosophers concerned with conceptual analysis focused on *clear-cut* definitions? One explanation might be a desire for clarity. Another might be a desire to incorporate the properties defined through conceptual analysis into axiomatic systems, perhaps in terms of set theory, the sets of which have traditionally been defined in such a way that the borders that delimit sets are completely clear-cut and no distinctions are made between different members of the same set. Undoubtedly, understanding properties in terms of such determinate sets has several advantages; for one thing, it provides a nice framework for inferences between properties. However, as we have seen, the classificatory structures that arise out of actual human

[3] Common practice suggests additional desiderata, e.g., that the set of characteristics cited in the analysis should be fairly small and their relation straightforward, so as to make sure that the resulting analysans is as *simple* as possible (Weatherson 2003, Ramsey 1998). However, for present purposes, I will be focusing on exactitude and exhaustiveness.

classification do not lend themselves to such characterizations and this has direct implications for philosophical methodology. Or to put the point more bluntly, as William Lycan has done recently in a retrospective piece on the Gettier discussion:

> It is well to remind ourselves that no effort of analytical philosophy to provide strictly necessary and sufficient conditions for a philosophically interesting concept has ever succeeded. And there should be a lesson in that (Lycan 2006, 150).

But what is the lesson? Since simply rejecting exactitude does not amount to abandoning the idea of philosophy as conceptual analysis, it is too early to be pessimistic about the dominant approach as such. For this reason, the next section considers two possible rectifications of traditional methodology, highlighting yet another problem for conceptual analysis—i.e., does the armchair provide the most suitable methodological location for the analysis of concepts in light of the more rigorous methods of empirical psychology?—that, ultimately, leads us to the more fundamental question: Why (if at all) should philosophers analyze concepts in the first place?

2. Two Attempts at Rectification

It might be argued that the methodological implications of aforementioned psychological research are perfectly straightforward: Any philosopher concerned with the analysis of concepts has to take the empirical work on actual human categorizations seriously. This also highlights the sense in which the problem for exactitude ammounts to an *internal* objection to conceptual analysis, that does not discredit the project of analyzing concepts *as such*, but merely throws doubt on a particular way of doing so. Consequently, the problem might seem perfectly solvable; all we need to do is find an empirically more warranted way to characterize concepts.

2.1 Prototypical Conceptual Analysis

One way to do this is to do conceptual analysis in terms of so-called prototypes, i.e., lists of prototypical features which, according to some psychologists, provide more plausible representations of human concepts (Rosch and Mervis 1975).[4] Such an analysis—let us refer to it as Prototypical Conceptual Analysis, or PCA for short—would proceed by way of an assumption to the effect that, for any concept '*F*,' there is a prototype set, *Q*, and a threshold value for categorization, *T*, such that the weighted sum of *Q* predicts posi-

[4] Goldman and Pust (1998) suggest something along these in light of aforementioned problems with traditional conceptual analysis, as does Goldman (2007).

tive categorization for the given concept when, and only when, it coincides with or exceeds T.

Let us consider this assumption in light of the two aforementioned desiderata. Given that items may have more or less of the features in Q and the features may be assigned different weights, items that have enough weighted features to exceed T will form a continuum. Furthermore, given that T may be construed either as an absolute value or as a value satisfied in so far as it is approximated, the present model can also handle situations involving vague concept boundaries. As such, PCA rejects exactitude, in allowing for the possibility of the kind of intra-concept distinctions and non-clear-cut boundaries identified by way of the scientific findings reviewed above.

If we hold on to exhaustiveness, PCA latches on to the relevant weighted features and yields predictions of categorization judgments that do not admit of any (genuine) counterexamples. As already the above characterization makes clear, PCA would most likely yield utterly complicated analyses if conducted thus. Still, if one looks at actual philosophical practice, PCA is not terribly different from traditional conceptual analysis here. Unlike traditional conceptual analysis, however, PCA would not generate complicated analyses due to an inaccurate theory of concepts. Rather, PCA would yield complicated analyses because concepts and categorizations are, as a matter of empirical fact, governed by quite complicated mechanisms.

2.2 Analysis, the *A Priori* and Empirical Investigation

This observation also serves to highlight the fact that conceptual analysis may be construed as an explicitly empirical task—i.e., as substantially relying on sense experiential input—indeed, perhaps even a *scientific* task, which is in stark contrast to how it has been conducted by most analytic philosophers. Considering the common view that conceptual analysis is (and should be) *a priori*—i.e., that it pertains to pure thought or reason alone—it will serve us well to elaborate on the distinction between empirical and armchair inquiry in relation to the distinction between *a priori* and *a posteriori* (i.e., non-*a priori*) warrant.

As pointed out by Alvin Goldman (2007), the two distinctions do not seem to coincide. Granted, *a priori* warrant never involves an empirical method and paradigm examples of an empirical method do, indeed, proceed by way of *a posteriori* sources of warrant (knowledge through perceptual observation being one of them). But since some warrant is neither *a priori* nor flowing in any straightforward way from sense experience, it hardly follows that *a posteriori* warrant never involves a non-empirical method. Take introspection and memory, for example, neither of which can be plausibly said to give rise to *a priori* warrant. However, since they do not relate to our

senses in any relevant way, they would still qualify as non-empirical sources of warrant under the above characterizations.

This way of construing the distinctions, Goldman points out, enables us to formulate two distinct approaches to categorization intuitions. On *the third-person approach*, the classification judgments of a subject are treated as data used in testing hypotheses about the content of the subjects' concepts. Since the evidence is distinctly empirical, the approach is an example of conceptual analysis construed in empirical, *a posteriori* terms. Still, it does not resemble conceptual analysis as it has traditionally been conducted, which brings us to the *first-person approach*, where one is primarily consulting one's own categorization intuitions. Construed thus, conceptual analysis does not involve perceptual observation. But on the above definitions, it does not, thereby, follow that the resulting warrant is *a priori*. As Goldman points out, "the process of generating classification intuitions has more in common with memory retrieval than with purely intellectual thought or ratiocination, the core of the a priori" in that the "generation of classification intuitions involves the accessing of a cognitive structure that somehow encodes a representation of a category" (Goldman 2007, 20). As such, conceptual analysis from a first-person approach is best described as an instance of non-empirical, *a posteriori* investigation.

What is the most promising approach to PCA, the first-person or the third-person approach? Here, we need to keep in mind that we, on PCA, not only need to determine the prototypical features of the concept in question but also the weight of each feature—not to mention any contextual factors that might influence the actual assignment of such weights. And make no mistake: This is all as it should be, if we want fully exhaustive characterizations of our concepts. In other words, if we are to hold on to the ambition of exhaustively characterizing concepts, it seems reasonable that conceptual analysis should go beyond the armchair exercises characterizing contemporary conceptual analysis in favor of hands-on, empirical investigations. For one thing, the methods of empirical science are better suited for coping with performance errors on part of the intuiting subject, i.e., situations in which our intuitions are off the mark due to human limitations in attention span, computational capacity, and the like. In addition, empirical science has superior resources for not only collecting and handling large sets of data but also for generating plausible generalizations about human concepts, when compared to the experimental condition of the single, intuiting subject.

This, clearly, speaks for the third-person perspective on intuition probing. As such, however, it is *not* a call for giving up on intuitions. Indeed, any inquiry into our concepts will rely, at least in part, on introspective reports regarding conceptual content. However, above considerations do suggest

that the armchair does not provide the most methodologically sound location, if philosophical inquiry is to be conducted in accordance with exhaustiveness, given the more rigorous methods of empirical psychology, which typically do not only proceed by way of a more substantial body of data—as in: data that goes beyond what is provided by the introspective reports of a single philosopher and her colleagues—but also takes seriously the methods as well as the methodological challenges that face psychological research in general and introspective reports in particular.

2.3 Why Exhaustiveness?

At the same time, we need to remember that what we are interested in here is conceptual analysis as it figures in philosophy. Exhaustiveness is not an unreasonable desideratum within the branch of psychology concerned with the empirical study of concepts, since nothing short of exhaustive accounts would reveal those subtleties that we need to understand in order to get a better grip on the exact structures of concepts and the ways in which they develop. But let us turn to epistemology, not because that field has any priority in philosophy at large, but because it is a field where I think the relevant methodological questions can be given informative answers. More specifically, keeping in mind the possibility of a difference in desiderata between psychology and epistemology, why (if at all) should epistemologists be interested in exhaustive accounts of our epistemic concepts?

First of all, it needs to be remembered that epistemologists are not interested in analyzing just any concepts. The ones of interest are the particularly *normative* concepts such as KNOWLEDGE and JUSTIFICATION, that, unlike non-normative concepts (like TABLE or DOG), are not only associated with a set of semantic norms regarding proper word use, but also with a set of particularly *epistemic* norms or hypothetical imperatives according to which you *should* be justified, in so far as you have certain epistemic goals. As such, the present normative framework is explicitly *instrumentalist*, in that it takes epistemic normativity to be a question of the extent to which something is a means to a desired epistemic goal.[5] In the following, I will refer to combinations of concepts, norms, and goals as *epistemic architectures* (without passing any judgment on their instrumentalist merits), representing different frameworks for epistemic inquiry.[6]

[5] That a subject has certain epistemic goals does *not* imply that these goals may not be overruled by non-epistemic considerations in natural settings. In other words, if it is an epistemic goal of a subject to (say) have true beliefs, that does not imply that she should have true beliefs *all things considered*.

[6] My notion of epistemic architectures bears some similarity to Goldman's (1992) notion of epistemic folkways. However, unlike Goldman, I will (*a*) not assume that there is enough

Next, note that, if we take such epistemic architectures as objects of epistemological study, we may identify a candidate rationale for why we should analyze epistemic concepts in exhaustive ways:

(S) Our epistemic concepts and norms are in full sync with our epistemic goals, in the sense that, by adhering to the norms in which our epistemic concepts figure, we are presented with the optimal way of reaching our epistemic goals.[7]

Does (S) hold? That is, can we assume that what we *should* do and what we (barring performance errors) are *prone* to do coincide? As noted by Quine (1969b) in an oft-cited passage: "Creatures inveterately wrong in their inductions have a pathetic but praiseworthy tendency to die before reproducing their kind" (Quine 1969b, 126). More specifically, consider the following (admittedly informal) *modus tollens*: Say that our epistemic concepts, norms and goals are *out of* sync, in the specific sense that, by employing our epistemic concepts and norms, we tend to *not* reach our epistemic goals. Under the common assumption that one central epistemic goal is true belief, it follows that we tend to *not* have true beliefs. Furthermore, since true (or at least approximately true) belief is an integral part of attaining most practical goals, including those involved in survival, it would follow that we tend to not survive—which is demonstrably false.

However, this modus tollens, at most, lends support to the idea that

(S*) our epistemic concepts and norms are *sufficiently* in sync with our epistemic goals to guarantee that, by adhering to the norms in which our epistemic concepts figure, we tend to have sufficient success in reaching our epistemic goals to guarantee the attainment of most practical goals.

But this, of course, is completely compatible with the idea that

(A) there might be an *alternative* set of concepts and norms such that, if we were to employ those concepts and norms instead, we would reach our epistemic goals to a greater extent than we are currently doing.

conceptual homogeneity to warrant talk about our epistemic folkways (Weinberg et al. 2001), and (*b*) not only include concepts and norms but also epistemic goals, which, I take it, is also a more proper place to assume at least some degree of homogeneity across different architectures. However, see (Goldman 2001, 477), where Goldman seems to concede the latter point.

[7] Cf. (Weinberg et al. 2001) on "epistemic romanticism".

By way of illustration, say that an elucidation of our—or, at the very least, a sufficiently prevalent—concept of justification yields the conclusion that to be justified is to reason in accordance with one's evidence in the specific sense of scrutinizing the evidential connections that hold between one's beliefs and their grounds and, ideally, only assenting to those propositions that survive such scrutiny (Conee and Feldman 2004). Assume, furthermore, that our ability to scrutinize the evidential relations between our beliefs and their grounds is, in many situations, weakened by the dual fact that (*a*) we seldom have introspective access to the grounds for our beliefs, and, (*b*) in the cases where we do, we sometimes misconstrue them in ways that may be ever so flattering to our self-images but that often make for quite unreliable reasoning tendencies (Wilson 2002). If so, there is an alternative way to conceive of justification, if only in the minimal sense of a conception that takes this empirical fact into account and, thereby, provides a more promising tool in the attainment of true belief (more on the question of axiology below).

In light of the possibility of this kind of scenario, why pursue exhaustive accounts of our concepts? Given the potential gap between our actual architectures, and the variety of architectures that might enable us to attain our goals to an even greater extent than we are presently doing, it is no longer obvious why the mere fact that an account violates our intuitive conceptions should disqualify it as a potentially promising way to attain our epistemic goals. After all, on the present instrumentalist framework, concepts and the norms they figure in are only as good as they enable us to attain our epistemic goals. Granted, epistemology needs to start somewhere and our current epistemic concepts serve the crucial role of fixing the subject matter. However, given the possibility of aforementioned gap, we no longer have any reason to believe that we are better off with comprehensive rather than with the kind of approximate accounts that may be provided by most competent users of the corresponding terms.[8] Focusing on competent users will guarantee a *prima facie* relevance of the accounts in question, while still acknowledging the fact that the real money will *not* be in knowing whether or not the approximate accounts provided capture all the idiosyncrasies of our epistemic concepts, but in determining whether the notions introduced can be said to serve us well, given our goals as inquirers (cf. Kitcher 1992).

What are the implications for PCA? By throwing doubt upon exhaustiveness, the above line of reasoning suggests that PCA's focus on experimental

[8] The notion of a "competent user" that I have in mind here is not a particularly demanding one. I simply mean to be referring to the kind of speakers that may communicate on the matters at issue with other speakers without any obvious problems about incomprehension or communicative breakdowns.

methods in elucidation is misdirected, at least as far as epistemology is concerned.[9] More than that, it also suggests that the stage at which we are going to want to invoke experimental science is not so much at the point of fixing the subject matter, as in the context of an evaluation of the extent to which the goals are actually attained by the methods and practices suggested by the potential multitude of approximate accounts in question. As such, the above raises two questions. First, given that there is no obvious rationale for pursuing exhaustive rather than approximate accounts of our current conceptions, what epistemological weight should be given to the investigation of our current epistemic norms and concepts? Second, if traditional conceptual analysis is not likely to do the trick as far as the evaluation of our epistemic ways is concerned, what exactly is the role of more straightforwardly *empirical* investigations in epistemology?

2.4 Analysis via Reflective Equilibria

In attempting to answer these questions, let us consider an influential methodological suggestion that, unlike traditional and prototypical conceptual analysis, relinquishes exhaustiveness in favor of the construction of philosophical theories that put our categorization intuitions in a *reflective equilibrium*—i.e., in a state of balanced co-existence—with any general principles that we may accept. Although first discussed by Nelson Goodman (1983) as a way to account for the justification of inductive norms, it was John Rawls (1971) who first introduced the idea of reflective equilibrium as a desired end state of philosophical analysis. On the process Rawls imagined, the intuitions associated with considered judgments about particular instances do, indeed, carry some philosophical weight. However, just like a general principle may be reconsidered in light of a conflict with particular judgments to which we are strongly attached, a conflicting categorization judgment need not in all cases indicate a flaw in the theory, but in some cases rather be rejected, given that the principle responsible for the conflict can be reasonably deemed more central, explanatory, or in any other way more important than that particular judgment. Hence, the rejection of exhaustiveness.

This provides at least the beginnings of *an* answer to the question of what role our current concepts should play in epistemology, even if demoted from a position where they have the first *and* final word: They provide some relevant data, but only in so far as they are deeply entrenched and do not conflict with any general principles deemed more important. However, when considering whether the method of reflective equilibria might be successfully

[9] Note that the same conclusion would follow on any other incorporation of a scientific theory of concepts into conceptual analysis, such as the exemplar view. See (Murphy 2002) for an excellent overview of the relevant experimental literature.

applied in epistemology, we need to remember the point made in the previous section. That is, while bringing our categorization judgments and general principles into greater coherence certainly may be a good thing, the goal of analysis is not merely to prune our conceptions for coherence; it is also to make sure that the resulting analysans serves us well in the attainment of our goals. This point is, by no means, lost on the defender of reflective equilibria.[10] For example, Michael DePaul (1998) notes that the process of seeking a reflective equilibrium, properly construed, cannot be a question of mere coherence, but rather "an even wider equilibrium" wherein the analyst also must factor in any relevant theories that she accepts (cf. Daniels 1979). Still, the extent to which incorporating theories thus may make a reflective equilibrium approach both applicable within epistemology and able to handle the kind of suboptimal combinations of concepts, norms, and goals discussed above depends on whether the theories are responsive to the ends that we are striving for. More specifically, a pursuit of reflective equilibria in epistemology would not only need to take into account our current concepts and norms, but also incorporate *explicitly* empirical theories, in order to leave room for an evaluation of the extent to which the concepts and norms utilized provide cognitive creatures like us with promising tools in the attainment of our goals—all in accordance with what was argued in the previous section.

That being said, the problem with a characterization along these lines is not that it is on the wrong track but that so much more needs to be said (cf. Williamson 2008). As it stands, the reflective equilibrium approach seems no more than a call for conscientious, philosophical inquiry, grounded in our best empirical theories about the world. As such, it, clearly, has a place in a sound, epistemological methodology. Still, the main methodological challenge—as it has been demarcated above—does not lie in noting that philosophers need to take into account not only our intuitions and general principles but also our best theories about relevant empirical matters in constructing their analyzes, but in providing a more detailed story about the exact role that our categorization intuitions and norms should (or should not) play in philosophical methodology, and how such intuitions and norms should be weighed against empirical evidence, so as to yield an analysans that serves us well. To this issue I now turn.

[10] See, e.g., (Rawls 1974/1975, 8; Rawls 1971, 49), where Rawls himself shows some sensitivity to this point.

3. A Naturalistic Alternative

The present section explores the positive implications of the largely negative points made above for the possibility of sketching a naturalistic alternative to the dominant approach by demarcating the respective roles of categorization intuitions and more straightforward empirical investigations within three epistemological areas: the evaluation of epistemological hypotheses, the amelioration of epistemic practices, and the construction of a theory of epistemic value.

3.1 Epistemological Evaluation and Amelioration

When stripped of its metaphysically saturated historical heritage, what are categorization intuitions? The result of psychological mechanisms that, for all we know, vary in systematic ways with cultural context (Weinberg et al. 2001) and often are influenced by (philosophically) irrelevant considerations (Swain et al. 2008), but that, when properly utilized, nevertheless may provide us with interesting data about the concepts and norms that have been handed down to us from our epistemic ancestors. As noted above, spelling out those concepts and norms is a crucial step for epistemology in general and amelioration in particular. More specifically, intuitions about such concepts and norms may serve to fix the subject matter of epistemology, by providing the basic material for the approximate accounts supplied by competent users of the corresponding terms.

As such, the role of intuitions would have to be considered fairly modest. At the very least, the stage of conceptual elucidation does not amount to as extensive an investigation as past epistemological excursions (if not immersions) into the domains of conceptual analysis may lead one to believe. This is to be expected since the use of intuitions, on the present suggestion, is stripped of the justificatory powers it has been endowed with within the dominant approach, and instead restricted to the context of epistemological discovery. More specifically, the point of the relevant elucidatory exercises, I maintain, is not to evaluate or justify epistemological *theories* but to provide epistemological *hypotheses*, that (given the rejection of exhaustiveness) are properly evaluated not so much with reference to whether or not they clash or mesh with our intuitions, as by way of an empirical inquiry into the question of whether the concepts and norms postulated serve us well, given our epistemic goals.

This, of course, just serves to underscore the fact that epistemology must go beyond—indeed, *far* beyond—the data provided by intuitions. As we have seen, this has, to some extent, been appreciated already by those defending a reflective equilibrium conception of epistemological methodology, although that conception failed to provide a sufficiently informative ac-

count of the exact relationship between conceptual questions and empirical investigations. Beyond that, there is, of course, a long tradition in philosophy of a fraternization (and, before that, an identification) with the sciences, halted momentarily during the last century by a process of methodological purification inspired by Frege's anti-psychologism (Kitcher 1992), re-awakened through the works of W. V. O. Quine (1969a), and developed further by more recent naturalistic epistemologists, perhaps most prominently by Goldman (1999, 1992, 1986, 1978).[11]

In this tradition, two areas of empirical contribution to epistemological theory play a particularly pivotal role. First, in order to properly evaluate the merits of different epistemological hypotheses about our current architectures, we need to pay attention to empirical research on how creatures like us actually go about forming belief. Relevant empirical research include the variety of research programs in cognitive psychology (widely construed) concerned with investigating the heuristics used in reasoning under uncertainty (Gilovich et al. 2002) and understanding how the reliability of such often surprisingly simple heuristics can be understood in terms of considerations about ecology and adaptivity (Gigerenzer et al. 1999). Second, and on a related note, to the extent that such evaluations suggest that we rely on suboptimal belief-forming strategies, we need to attend to empirical research in identifying ways in which they may be improved, typically with reference to the implementation of reasoning strategies that either cater to our cognitive strengths (Gigerenzer et al. 1999), or simply bypass some of our more serious cognitive failings, e.g., by relying on predictive models rather than intuitive judgments (Dawes et al. 2002).

3.2 Epistemological Axiology

Since these points have received thorough attention elsewhere, I will not dwell on them further here.[12] Instead, what needs to be stressed is that both aforementioned tasks need to be conducted against the background of a plausible theory of our epistemic goals, which brings us to the issue of constructing a theory of epistemic value. What role can intuitions play here? Here is a suggestion: By probing our intuitions about *why* we employ the concepts and norms that we do, we may find out something about what it is that we are at all trying to attain within epistemic inquiry, in the sense of what kinds of states, processes or properties are taken to be of epistemic

[11] See (Kornblith 1994) for additional reference and a nice anthology of relevant literature in and on the naturalistic tradition.

[12] See earlier references to Goldman's work. See also (Bishop and Trout 2005) for a more recent elaboration on the naturalistic project.

value.[13] What will such an investigation reveal? Most likely, a concern with *getting things right*, to use Andrew Latus' (2000) phrase. That might not seem very impressive a result. Still, it is a start and it seems to capture something central (albeit not necessarily exhaustive) about the epistemic point of view: As epistemic inquirers, we strive to attain true belief and avoid false belief. Let us refer to this as the *minimal conception of epistemic value*.

Since we are concerned with amelioration, the desideratum most relevant in the fleshing out of this conception is that the resulting theory be sensitive to the epistemic priorities of actual inquirers—the targets of ameliorative advice and the closest we get in reality to epistemic inquirers.[14] One way for an axiological theory to fail this desideratum is by working with a too abstract notion of epistemic value. In particular, any axiological account in terms of inquirers aiming at truth or knowledge *per se* will be too far removed from the concerns of actual inquirers to make for a fruitful basis of ameliorative advice. This is not to deny that inquirers value truth (or knowledge, etc.), but merely to say that different inquirers value *different* truths (or pieces of knowledge and justified belief), because they are trying to answer different questions. As such, we would want our axiological theory to show at least some sensitivity to what kinds of questions, results, explanations, and so on, are actually deemed *significant* within different areas of epistemic practice (Bishop and Trout 2005, Kitcher 2001).

At this point, it might seem that the path of axiological investigation would be straightforwardly empirical: All we need to do is find out what matters inquirers deem significant. However, it should be noted that another way to fail the above desideratum is by taking a proper axiology to flow directly from such an empirical investigation. Such an axiology would not fail by being out of touch with inquirers' priorities, but by not being sensitive to the distinction between the epistemic and the non-epistemic vectors underlying those priorities. After all, the choices inquirers make regarding how to invest their limited resources are a function of a variety of factors, including those pertaining to availability of funds, popular support, and so on. These and similar factors that are of real and legitimate concerns for scientific practice. However, they are arguably of no great epistemological relevance, and should, as such, not figure prominently (if at all) in an epistemological axiology.

[13] This point bears some resemblance to Edward Craig's (1990) notion of a "practical explication". However, unlike Craig, I do not take the relevant form of explication to be unveiling concepts, as much as it reveals the goals with which they are associated. Since I have elaborated on this point elsewhere (Ahlstrom 2007), I will not dwell on it here.

[14] This is not to say that there might not be other desiderata for epistemological axiology at large. See, for example, the variety of theoretical demands and desiderata discussed in relation to the so-called value problem (Pritchard 2007).

So, what factors *should* figure in an epistemological axiology? Here is a suggestion: The truth and falsity of belief, as it pertains to matters deemed significant. For all its simplicity, this suggestion negotiates the above points by (*a*) focusing on truth and falsity of belief, in order to exclude non-epistemic factors from epistemic evaluation, yet (*b*) restricting the scope to truths and falsities pertaining to matters deemed significant, in order to avoid aforementioned disconnect between epistemological axiology and epistemic practice. That being said, the suggestion certainly does not amount to a full-blown theory of epistemic evaluation. Rather, it is a general schema of evaluation that remains silent on two important points. First, what specific distributions of true and false belief are (more) epistemically desirable (than others)? Candidate distributions include but are not exhausted by *bodies of belief with no false belief, bodies of belief with large numbers of true belief, bodies of belief with favorable truth-falsity ratios*, and *large bodies of belief with favorable truth-falsity ratios* (as well as any combination thereof). Second, what matters are, as a matter of fact, deemed significant? Differently put, what *kinds* of truths (and falsities) are relevant to inquirers?

Since the topic of the present paper is methodology—not axiology—I will not attempt to answer these questions here, but instead make a methodological point about the respective roles that intuitions and empirical investigation may play in answering them. Consider, first, the issue of candidate distributions. What facts could be appealed to in choosing between different candidates here? We have already considered a problem with simply looking to the practices of inquiry here—practices conducted against the background of too rich a set of non-epistemic factors for any straightforward investigations to reveal anything substantial about *epistemic* value. Alternatively, we could turn to our intuitions, which also seems the most common strategy in the literature.[15] However, it is not immediately obvious that we do have any particularly detailed intuitions about epistemic value beyond the general framework provided by the minimal conception. Moreover, to the extent that intuitions play a role in any refinement of this framework, their role could, most likely, be best explained not so much in terms of a sensitivity to some axiological truth, as with reference to the fact that some of us simply have gotten used to thinking about epistemic value in a particular way, whereupon our intuitions follow suit.

This is not to say that what epistemological axiology we opt for is a matter of a completely arbitrary choice. For one thing, it would seem that any account that is incompatible with the minimal conception is not an account

[15] See (Grimm 2008), (Alston 2005), and (Latus 2000) for three recent and representative examples of the nature of the discussion, and the way in which empirical (not to mention scientific) considerations play a minimal role in the arguments considered.

of *epistemic* value.[16] But beyond the constraints imposed by that notion, there does seem to be a constructive element to the development of a theory of epistemic value, if not even one of stipulation. This is in contrast to the second question above, about what matters are deemed significant. Here, it would seem that there not only is a matter of fact to identify, but that it could be identified by straightforwardly empirical means. For example, we could start by identifying certain paradigmatic instances of significant matters, and then analyze these instances to see whether they share certain features and can be grouped into more general *kinds* of issues.[17] For present purposes, it does not matter so much exactly what such an investigation would yield. What is important is that its results would serve to anchor our epistemological axiology in the concerns of actual inquirers, and ensure that the kind of ameliorative advice considered above will be relevant to the concerns of inquirers. Moreover, if combined with an account of epistemic evaluation that is compatible with the general framework of the minimal conception, such an axiology would be restricted to epistemic factors, and, thereby, also preserve the specialized nature of epistemic evaluation, as concerned (at the very least) with the generation of true belief and the avoidance of false belief.

4. Conclusion

The dominant approach to philosophy as conceptual analysis suffers from a series of methodological problems, at least within epistemology. As has been suggested by many philosophers, one of the main problems pertains to an inaccurate theory of concepts, but further investigation suggests that the problem runs deeper than that, and ultimately motivates a close alliance with the empirical sciences and an abandonment of the idea of philosophical analysis as exclusively or even primarily concerned with concepts. The alternative methodological picture painted above—and I say "picture," since I do not pretend to have presented anything like a complete methodological theory—was an attempt to not only account for the nature of that alliance, but also elaborate in more detail on the interaction between relevant forms of empirical investigations and aspects of philosophical inquiry that do (and in some cases perhaps even have to) rely to some degree on the probing of intuitions.

As such, the picture suggests a certain methodological continuity with both nineteenth century analytical epistemology and historically prominent as well as more recent naturalistic attempts to bridge the (perceived) gap

[16] Notice that this claim is compatible with the proper account of epistemic value being one that includes truth as one epistemic goal among many (Kvanvig 2005).

[17] See Kitcher's (2001) significance graphs for what would, most likely, be a helpful conceptual tool in such an investigation.

between philosophy and the sciences. Needless to say, further refinements of this picture are both desirable and necessary. Still, it would seem that it at least has the outlines of the kind of methodological approach that, albeit thoroughly naturalistic, not only illustrates the respect in which any complete epistemology may legitimately invoke both empirical and non-empirical modes of inquiry and, thereby, make good of the original Quinean promise of a naturalized epistemology as one involving a *reciprocal* containment between epistemology and the natural sciences (Quine 1969a), but also identifies a kind of epistemological methodology that is empirically informed and beneficial to the practices of inquiry without, thereby, being dialectically alienating for non-naturalistic philosophers.

Acknowledgements

The author would like to thank Michael Bishop, Hilary Kornblith, Helge Malmgren, Alexander Sarch, Åsa Wikforss, and an anonymous reviewer for this journal for valuable comments on earlier drafts of this paper.

Bibliography

Ahlstrom, K. (2007). *Constructive Analysis: A Study in Epistemological Methodology*, PhD thesis, University of Gothenburg.

Alston, W. (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation*, Cornell University Press, Ithaca, NY.

Bealer, G. (1998). Intuition and the autonomy of philosophy, *in* M. R. DePaul and W. Ramsey (eds), *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Rowman & Littlefield Publishers, Inc., Lanham, MD, pp. 201–240.

Bishop, M. and Trout, J. D. (2005). *Epistemology and the Psychology of Human Judgment*, Oxford University Press, Oxford.

Conee, E. and Feldman, R. (2004). *Evidentialism: Essays in Epistemology*, Oxford University Press, New York, NY.

Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Clarendon Press, Oxford.

Daniels, N. (1979). Wide reflective equilibrium and theory acceptance in ethics, *The Journal of Philosophy* **76**: 256–282.

Dawes, R., Faust, D. and Meehl, P. (2002). Clinical versus actuarial judgment, *in* T. Gilovich, D. Griffin and D. Kahneman (eds), *Heuristics and*

Biases: The Psychology of Intuitive Judgment, Cambridge University Press, Cambridge, pp. 716–729.

DePaul, M. R. (1998). Why bother with reflective equilibrium?, *in* M. R. DePaul and W. Ramsey (eds), *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Rowman & Littlefield Publishers, Inc., Lanham, MD, pp. 293–309.

Gettier, E. (1963). Is justified true belief knowledge?, *Analysis* **23**: 121–123.

Gigerenzer, G., Todd, P. M. and the ABC Research Group (eds) (1999). *Simple Heuristics that Make Us Smart*, Oxford University Press, Oxford.

Gilovich, T., Griffin, D. and Kahneman, D. (eds) (2002). *Heuristics and Biases: The Psychology of Intuitive Judgment*, Cambridge University Press, Cambridge.

Goldman, A. (1978). Epistemics: The regulative theory of cognition, *The Journal of Philosophy* **75**: 509–523.

Goldman, A. (1986). *Epistemology and Cognition*, Harvard University Press, Combridge, MA.

Goldman, A. (1992). Epistemic folkways and scientific epistemology, *Liaisons: Philosophy Meets the Cognitive and Social Sciences*, The MIT Press, Cambridge, MA, pp. 155–175.

Goldman, A. (1999). *Knowledge in a Social World*, Oxford University Press, Oxford.

Goldman, A. (2001). Replies to the contributors, *Philosophical Topics* **29**: 461–511.

Goldman, A. (2007). Philosophical intuitions: Their target, their source, and their epistemic status, *Grazer Philosophische Studien* **74**: 1–26.

Goldman, A. and Pust, J. (1998). Philosophical theory and intuitional evidence, *in* M. R. DePaul and W. Ramsey (eds), *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Rowman & Littlefield Publishers, Inc., Lanham, MD, pp. 179–197.

Goodman, N. (1983). *Fact, Fiction, and Forecast*, 4 edn, Harvard University Press, Cambridge, MA.

Grimm, S. R. (2008). Epistemic goals and epistemic values, *Philosophy and Phenomenological Research* **77**: 725–744.

Kitcher, P. (1992). The naturalists return, *The Philosophical Review* **101**: 53–114.

Kitcher, P. (2001). *Science, Truth, and Democracy*, Oxford University Press, Oxford.

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

Kornblith, H. (ed.) (1994). *Naturalizing Epistemology*, 2nd edn, The MIT Press, Cambridge, MA.

Kripke, S. (1980). *Naming and Necessity*, Harvard University Press, Cambridge, MA.

Kvanvig (2005). Truth is not the primary epistemic goal, *in* M. Steup and E. Sosa (eds), *Contemporary Debates in Epistemology*, Blackwell Publishing, Malden, MA, pp. 285–296.

Latus, A. (2000). Our epistemic goal, *Australasian Journal of Philosophy* **78**: 28–39.

Laurence, S. and Margolis, E. (2003). Concepts and conceptual analysis, *Philosophy and Phenomenological Research* **67**: 253–282.

Lewis, D. (1973). *Counterfactuals*, Blackwell, Oxford.

Lycan, W. (2006). On the gettier problem problem, *in* S. Hetherington (ed.), *Epistemology Futures*, Oxford University Press, Oxford, pp. 148–168.

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

Parfit, D. (1984). *Reasons and Persons*, Clarendon Press, Oxford.

Pritchard, D. (2007). Recent work on epistemic value, *American Philosophical Quarterly* **44**: 85–110.

Putnam, H. (1975). The meaning of 'meaning', *Mind, Language, and Reality: Philosophical Papers*, Vol. 2, Cambridge University Press, Cambridge, pp. 215–271.

Quine, W. V. O. (1969a). Epistemology naturalized, *Ontological Relativity and Other Essays*, Columbia University Press, New York, NY, pp. 69–90.

Quine, W. V. O. (1969b). Natural kinds, *Ontological Relativity and Other Essays*, Columbia University Press, New York, NY, pp. 114–138.

Ramsey, W. (1998). Prototypes and conceptual analysis, *in* M. R. DePaul and W. Ramsey (eds), *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, Rowman & Littlefield Publishers, Inc., Lanham, MD, pp. 161–178.

Rawls, J. (1971). *A Theory of Justice*, Harvard University Press, Cambridge, MA.

Rawls, J. (1974/1975). The independence of moral theory, *Proceedings and Addresses of the American Philosophical Association* **47**: 5–28.

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

Samuels, R., Stich, S. and Tremoulet, P. D. (1999). Rethinking rationality: From bleak implications to darwinian modules, *in* E. Lepore and Z. Pylyschyn (eds), *Rutgers University Invitation to Cognitive Science*, Basil Blackwell, Oxford.

Shope, R. K. (1983). *The Analysis of Knowing: A Decade of Research*, Princeton University Press, Princeton, NJ.

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

Swain, S., Alexander, J. and Weinberg, J. (2008). The instability of philosophical intuitions: Running hot and cold on "truetemp", *Philosophy and Phenomenological Research* **76**: 138–155.

Thomson, J. J. (1971). A defense of abortion, *Philosophy and Public Affairs* **1**: 47–66.

Weatherson, B. (2003). What good are counterexamples?, *Philosophical Studies* **115**: 1–31.

Weinberg, J. M., Nichols, S. and Stich, S. (2001). Normativity and epistemic intuitions, *Philosophical Topics* **29**: 429–460.

Williamson, T. (2008). *The Philosophy of Philosophy*, Wiley-Blackwell, Malden, MA.

Wilson, T. (2002). *Strangers To Ourselves: Discovering the Adaptive Unconscious*, Harvard University Press, Cambridge, MA.