Applied Epistemology: Prospects and Problems

By Søren Harnow Klausen,
University of Southern Denmark

Abstract: Surprisingly few attempts have been made to apply theories in normative epistemology to concrete questions about epistemic policies and practices, in spite of the fact that many of today's most pressing and hotly debated questions are about the production, transmission and use of knowledge. The paper discusses the prospects for an applied epistemology. Alvin Goldman's conception of a social epistemology provides an important source of inspiration. But Goldman's framework needs amendments which are likely to make the task of applied epistemology more complicated. I argue that one needs to pay more attention to different forms of justification and knowledge and other epistemic desiderata. One cannot isolate the project of epistemic evaluation from pragmatic, prudential or moral evaluation. A pure applied epistemology is a contradiction in terms. But this does not mean that the whole project must be handed over to pragmatism. By accommodating factors that are not purely epistemic, one might actually strengthen the position of truth-based normative epistemology vis-à-vis pragmatic or sociological approaches to knowledge.

The Missing Discipline

Epistemology is a normative enterprise. It is about what we ought and ought not to believe. It is concerned with the evaluation of actions, practices and policies. Contemporary epistemologists are prone to emphasise the parallels between epistemology and ethics (see e.g. Goldman 1986: 2f.; Plantinga 1993: 3ff; Zagzebski 1996 xiii; Williams 2001: 11; a rare dissenting voice is Fumerton 1995: 8ff.). The internalism-externalism debate mirrors the traditional debate between deontology and consequentialism. The influential movement of virtue epistemology has been modelled more or less directly on virtue ethics. It is generally acknowledged that epistemology should look to ethics for conceptual tools to use in solving the traditional problems (see e.g. Fricker 2003).

However, the rapprochement between epistemology and ethics has taken place mainly on the level of metaepistemology and normative epistemology. There is a fairly long tradition of analysing concepts like knowledge and justification in ethical terms (see e.g. Ayer 1956: 31f; Chisholm 1957: 4), and discussions of normative epistemology often turn on questions concerning the

© Res Cogitans 2009 no. 6, vol. 1, 217-254

merits of consequentialism and deontology in general (Alston 1986; BonJour and Sosa 2003: 145f., 175f.). But there have been few attempts to apply the normative theories to concrete questions about epistemic policies and practices. "Applied epistemology" does not denote any distinctive branch of philosophy on the current scene.

This is surprising, since the production, propagation and use of knowledge are among the most popular topics of public debate. One hears much these days of the knowledge society, the knowledge economy, knowledge management, knowledge sharing etc. Surely large parts of this talk are superficial or not really about knowledge in the strict sense. Still, there seems to be plenty of genuine epistemological questions around. What are the most efficient teaching methods and learning strategies, and how should the educational system be organised? What weight should be given to experts' judgements about various subject matters? To what extent and in what respects should public and scientific debates, mass media and communication channels be open, and to what extent should they be censored or regulated (Lessig 2001)? What kind of management and organisation is most conducive to the production and dissemination of corporate knowledge (Tuomi 1999)?

Moreover, recent developments in epistemology have lead to the formation of theories that seem well fitted to deal with such practical questions. Apart from the general tendency to model epistemology on moral theory, many theories in the externalist and naturalist movement picture knowledge as the product of concrete, real-world processes that should be open to empirical investigation and practical evaluation and improvement. Dretske has described perceptual knowledge in terms of "the flow of information" (1981). Goldman's reliabilism is meant to provide an evaluative standard that can be used to examine particular belief-forming and problem-solving processes (1986: 181). Edward Craig has suggested that we view knowledge as a concept used to "flag approved sources of information" (1990: 10). Although contemporary internalists usually remain more tightly focused on abstract epistemological problems, some of them have pointed to a wide range of specific epistemic duties, values and norms (see e.g. Pollock 1986).

So why have philosophers shun the task of applying their theories to what seems to be some of today's most pressing problems? One reason may be that the field is too messy. Whereas ethical problems emerge from social life and popular debates in a relatively clear form, epistemic questions often appear inextricably intermingled with pragmatic and ethical questions. Discussions of educational policies are not only about how to achieve epistemic goals, but also about social justice, well-being, democratic citizenship and the like. The question "what is the best kind of education?" is multifariously ambiguous. When it comes to assessing specific institutions, non-epistemic values like the social atmosphere or the institution's reputation may play a larger role than epistemic ones.[1] Parents looking for a good school for their child may have in mind primarily – or at least partly – whether the child will have a pleasant time in school or whether it will acquire the right social skills or the right ethical or religious values.

Discussions of the knowledge-based economy tend to be even messier. It does not matter much from the point of view of management theory or economics whether an economic gain depends on the epistemic qualities of the procedure by means of which it was achieved. Their interest in knowledge is pragmatic, and this is reflected in the way they tend to play fast and loose with concepts like knowledge and information (see e.g. Boisot 1998).

This may suffice to explain why philosophers have been reluctant to engage in the practical debates. But it cannot justify it. For the situation in applied ethics is not that different. Many problems in this field are equally complex and ambiguous. Problems in research ethics often take the form of a conflict between epistemic and moral values. Some of the most popular arguments for free speech, multicultural societies or the presence of members of different social groups in juries appeal to the capacity of these arrangements for promoting knowledge (the *locus classicus* of this line of thought is Mill's defence of free thought and discussion in *On Liberty* (1859: Ch. 2)). Many allegedly moral judgements may in fact be aesthetic judgements in disguise: a person saying that homosexuality is bad may actually mean that

[1] Both factors can of course also be seen as means to achieve epistemic ends. But this is just a further example of the ambiguity in question.

she finds it disgusting; and people praising a heroic deed may value it more for its originality and creativity than for its moral qualities.

The upshot of this is that applied philosophy in general is a messy business. It is more about disentangling different questions and resolving ambiguities than actually arguing for particular conclusions.

Recent work on the method of applied philosophy points in the same direction. The simple top-down conception – which views the inquiry as a one-way process of applying principles laid down in advance to a definitely described practice – has been widely criticised. The currently dominant view is that there must be an interplay between the normative and the descriptive levels (Beauchamp 1984; Pihlström 1999). This accords well with the popular idea that the proper method of philosophy is to strive for reflective equilibrium, i.e. mutually adjusting general principles and particular judgements data to one another (Rawls 1971; Daniels 1979).

As will emerge, I do not find it wise to give up the top-down conception completely. If epistemology is to retain its normative function, some principles must be given a certain priority and kept relatively immune from revision due to particular empirical findings.[2] But I reckon that investigations of particular epistemic practices and judgements may prompt modifications of the normative framework, or at least affect the way it is applied. In any case, one must take into account that allegedly neutral descriptions of particular practices can be theory-laden. A newspaper or journal may be reported to be highly reliable; but the report may merely reflect the biased views of the source. Conceptual clarification and considered judgement are needed on all levels. Neither theories nor data come in a pre-digested form.

Pointing to the complex character of the enterprise forestalls another likely objection to applied epistemology: that it is trivial. It might be thought that once a certain normative epistemology has been adopted, once the epistemic issues have been disentangled from the pragmatic and ethical issues and received a clear formulation, and once the relevant empirical information has been obtained, the answer will follow automatically. This is undeniable.

[2] See Stich 1988. This objection to a non-hierarchical notion of reflective equilibrium echoes the familiar objections to the coherence theory of justification.

But it merely describes the ideal which an applied epistemology will be aiming at. In practice, there will be plenty of unfinished business on all levels, and in spite of the division of intellectual labour, the practiser of applied epistemology will have to work on all of them simultaneously.

Role Models for Applied Epistemology

My claim that applied epistemology is a missing discipline may strike some as implausible. And there are indeed some putative counterexamples. Most of them turn out to be irrelevant, but considering them may help me to clarify what I understand by applied epistemology and provide us with some hints as to where to look for genuine sources of inspiration.

The *expression* "applied epistemology" has been around for some time. There are electronic journals and research centres for applied epistemology. Yet the expression is seldom used as a name for the application of normative epistemology to practical, real-life problems. Sometimes it means *specialised* epistemology – epistemology that concentrates on a particular source of knowledge or an aspect thereof. Sometimes it refers to the application of specialised scientific theories, notably theories in cognitive science or psychology, to epistemology. The latter enterprise, which coincides with the project of naturalised epistemology, is indeed an example of how epistemology might be made more concrete and empirically informed. Still, it is not really applied epistemology. It is applied cognitive science.

There are other activities named applied epistemology which do seem to fit my definition. For example, there is an internet site dedicated to discussions of which authorities we should trust on various subject matters. But the participants in these discussions do not appeal to any clearly formulated normative principles. Instead, they make tacit appeal to common sense or uninterpreted scientific results. The discussions are thus comparable to popular ethical debates unaided by philosophical reflection.

Likewise, many discussions of education and learning touch on questions about how to achieve epistemic goals. But apart from failing to distinguish clearly between epistemic and other goals, most writers on education and learning tend to apply theories of cognition or society directly to their subject

matter, without paying sufficient attention to the normative dimension of their inquiry. For example, it is common to criticise the traditional stand-and-deliver method of teaching by citing anti-empiricist theories which view cognition not as a process of passive reception but of the active construction of concepts (see e.g. Bruner 1996). This line of thought surely has something to be said for it. Since an "ought" implies a "can", a theory of learning must be sensitive to the intellectual capacities of real-life students. Still, one can doubt whether theories about very general traits of human cognition can give detailed prescriptions for particular teaching methods. In any case, this is a question which a philosophically informed discussion of the issue would have to consider explicitly. It would also have to consider whether the cognitive theories have been chosen because of their greater plausibility or rather because they support the preconceived opinions of teachers and learning theorists. The theory of education is a field where applied epistemology might have a job to do, not a field where it is already to be found.

Still, there are a number of genuinely philosophical inquiries which seem to come closer to what I understand by applied epistemology. Take the philosophy of science. Is it not an attempt to evaluate a particular – albeit widespread – knowledge-seeking practice with respect to its epistemic qualities? If this description were fitting, I would simply have to concede that the field of science constitutes an important exception from my claim (though this would still leave me with several other fields of knowledge production). But in fact it is doubtful whether the philosophy of science really fits the description. As Sven Ove Hanssons puts it:

> The reason why we do no call the philosophy of science "applied epistemology" is that it does not constitute an application of epistemology in the same way as applied mathematics consists in the application of pure mathematics. Philosophers of science do not transfer and use theories from epistemology. Instead, they develop new theory that may be related to epistemology but is not derivable from it (Hansson 2003: 1).

It is indeed a significant fact that mainstream philosophy of science has been developed independently of normative epistemology. Concepts like knowledge and justification play a marginal role in standard works in the

philosophy of science, and the efforts of epistemologists to analyse them and formulate criteria of epistemic evaluation have been almost completely ignored.[3]

What, then, about the movement or discipline known as "social epistemology"? Again, the main reason for not counting it as a species of applied epistemology is that it does not pay sufficient attention to the normative dimension. Most defenders of social epistemology are out to *sociologise* epistemology, transforming it into a more or less neutral description of various processes of belief formation. It may be too much to demand of a genuine applied epistemology that it should adopt the traditional notion of knowledge as something which necessarily involves true belief. But at a minimum, an applied epistemology must posit some other distinctively epistemic norm, like idealised rational acceptability, long-run or otherwise qualified consensus, impartiality or coherence. Otherwise, it will not have the means to distinguish between those practices that are simply good at propagating beliefs and those that are good at propagating epistemically qualified beliefs. Most versions of social epistemology do not meet this requirement.[4]

There is, however, one notable exception.[5] The programme for a social epistemology presented by Alvin Goldman in his *Knowledge in a Social World* (1999) seems to fit my description in almost every respect. It is explicitly based on a recognition of truth as the fundamental epistemic value. It is explicitly concerned with the epistemic evaluation of practices. It can be seen as an application of Goldman's own reliabilist normative epistemology to practical questions. Last but not least, Goldman goes on to actually apply his framework to cases and topics quite similar to the ones I have envisaged.

A minor difference is that although Goldman's social epistemology is clearly an instance of applied epistemology, applied epistemology need not be exclusively social. Problems of individual knowledge-seeking might also call for an application of normative epistemology. In practice, however, the

---

[3] There are some slight exceptions in the more recent literature, see e.g. Psillos 1999: 83f.

[4] One prominent example is Fuller 1988

[5] Another, closely related example is the approach of Bishop and Trout (2005), which is less modest in its practical ambitions, but rather narrow in its scope.

domains of applied and social epistemology will overlap to a large extent. Most individual cases of knowledge-seeking, especially those which are likely to come up for discussion and evaluation, are also social, inasmuch as they are instances of common practice types and socially embedded. And Goldman himself seems to understand "social" in a rather weak sense. He is mainly concerned with finding the right social means for making individuals adopt certain practices or beliefs or making them respond appropriately to factors in their social environment.[6] This might seem disappointing to those who would like give epistemology a more strongly social orientation. But it makes the difference between social epistemology in Goldman's sense and applied epistemology in my sense negligible.

Goldman's theory can thus serve as a role model for applied epistemology. I have nothing to object to his general approach. However, his framework needs modifications and amendments which are likely to render its application more difficult than one could have hoped for. But first to some good news.

Problems to be Ignored

The complex nature of most problems of applied epistemology is something of an obstacle. But otherwise the prospects look far from bad. Many of the problems that beset normative epistemology are absent on the level of applied epistemology.

For example, applied epistemology does not have to await a resolution of the controversies in normative epistemology or metaepistemology. It can set scepticism to one side and make use of our best current theories about the nature and status of epistemic norms. This can be justified both by appealing to the division of intellectual labour and by noting that even if we had to give in to scepticism, we would still be left with surrogate notions of knowledge and belief – viz. some of those which had been found insufficient from the standpoint of rigorous basic epistemology, like, say, contextualist or

[6] An example is Goldman's treatment of testimony, which he takes to be social merely "because it operates on the reports of others" (1999: 130). In 2002b he explicitly suggests that the problem of testimonial justification might be viewed as a problem for individual epistemology.

sophisticated pragmatist or antirealist notions – which would probably be just as apt for this practical purpose. Here, too, the situation is just like that in applied ethics, where it is quite common to be agnostic about metaethics.

With sceptical challenges set to one side, the practiser of applied epistemology is free to choose among the competing theories of knowledge and justification solely on the basis of their merits as normative theories. A main objection to the relatively undemanding externalist notions of knowledge and justification is that they are hypothetical; externalism can say only that *if* the relevant cognitive processes (notably the epistemologist's own) are in fact reliable, *then* there is knowledge (Stroud 1994: 152; Fumerton 1995: 175ff.; BonJour 2002, 236). But this objection is neutralised once it is taken for granted that we do have knowledge in general and that our basic sources of belief formation are reliable. Even hardcore internalists are prone to admit that an externalist approach can be appropriate when it comes to assessing epistemic practices from a third-person-perspective (see e.g. BonJour 2002: 234; cf. Kusch 2002: 108ff; Brandom 1994: 212f.).

Another main objection to externalism is the so-called generality problem.

Conee and Feldman have argued that the notion of a reliable process is irredeemably vague, because a process token belongs to many different process types. Hence one and the same process token can be said to be more reliable if it is described as instantiating one type, e.g. as the act of looking at a maple leaf from a certain distance, and less reliable if it is described as instantiating another type, e.g. as a visual process occurring in a short-sighted person under bad lighting conditions) (Conee and Feldman 1998).

I believe that part of the answer[7] to this problem consists in simply conceding that the same process token can be described as more or less

[7] Another part of the answer may be that our basic cognitive processes are sufficiently similar to natural kinds to dispel the most troublesome ambiguities (as suggested by Alston 1995). I am aware that no version of this "psychological realist" approach has proved completely successful, but do not consider the problems raised by Conee and Feldman (1998) very severe. A realistic aim of the approach is not to dispel all ambiguities, only to narrow down the range of candidates for relevant types to an acceptable size.

reliable.[8] Sometimes we are interested in a process qua token of one type and sometimes we are interested in it qua token of another. I reckon that such an answer might sound overly concessive to some, perhaps because they think that it makes reliabilism into a kind contextualism and/or threatens to deprive it of its anti-sceptical potential. But again, the problem is much less pressing on the level of applied epistemology. Here it seems natural to assume that our descriptions and evaluations of epistemic processes will reflect our interests and the context of evaluation. (It should be noted, however, that once a particular description is chosen, it is still a determinate, objective fact how reliable the process is, and thus whether its reliability is above or below a certain threshold. Whether the threshold itself is absolute is a different question. My proposal does not commit one to relativism about the threshold for justification or knowledge and thus still differs significantly from typical forms of contextualism).

In fact I have already suggested that applied epistemology should allow of a plurality of descriptions of the same process token. One and the same process token may be described in either individualist or social terms – e.g. as an individual's adopting a belief on the basis of testimony or as part of a more comprehensive social process of belief transmission. Or to take an example that is more pertinent to the generality problem: We might ask whether Linda's habit of acquiring political information through reading and trusting a particular newspaper is sufficiently reliable. We might ask whether her habit of believing a particular commentator of that newspaper is sufficiently reliable. Or we might ask whether believing whatever the newspapers say is sufficiently reliable. That we will probably get very different answers to these questions is just as it should be. The epistemic evaluation of Linda's belief will change if the evaluator is first told simply that she got it from reading a newspaper, then learns that she got it from reading a particular newspaper and finally finds out that she adopted it from a very reliable commentator of that newspaper. It is hard to find any problem in this.

[8] Kappel (2006) argues that the externalist can embrace the "no determination view", i.e. accept that there is sometimes no fact of the matter as to which one of two conflicting epistemic evaluations of the same process token is correct.

Internalism Revived

It does indeed seem that the natural choice of a normative basis for applied epistemology would be some form of externalism, probably something akin to Goldman's reliabilism. Externalists frame epistemic questions in a way that makes them susceptible to empirical study and quantitative measurement, using concepts like process reliability, channels, transmission and information, which lend themselves easily to application to social and practical matters.

Yet it would it be too rash to conclude that internalist notions of knowledge and justification have no role to play in applied epistemology. For one thing, it is possible to assess an epistemic practice according to both externalist and internalist standards. Such a pluralist approach is quite popular in applied ethics, where it is common to ask what utilitarianists and deontologists, respectively, should say about a certain practice. And there need not be anything schizophrenic about it. Consequentialism and deontology tend to converge on the classification of particular judgements. Externalists and internalists are likely to recommend the same practices, externalists because they are truth-conducive, internalists because they are evidence-based (and because we *know* them to be truth-conducive). There may of course be cases where the two theories yield markedly different results. But that can itself be considered a significant result. If a practice meets only externalist criteria, this could be taken to indicate that its value is mainly pragmatic (or that we do not know enough about how it actually works). If a practice meets only internalist criteria, this could raise doubts about whether it is of real social significance. In both cases, conflicting results may prompt modifications of the normative framework.

More importantly, even if the general normative framework is kept purely externalist, internalist notions and requirements are likely to re-emerge on a more subordinate level. They may not be necessary conditions for knowledge in general, i.e. in its most basic sense. But they are nevertheless conditions for particular kinds of knowledge, some of which we take to be of great social significance. We do value believing on and gathering evidence, not believing on insufficient evidence, seeking and having reflective knowledge of

various kinds etc. We do distinguish between people who simply know certain things and people who are perceptive, reflective, enlightened etc. It is of course possible to provide an externalist explanation of these preferences of ours: evidence-basing and reflectivity are likely to enhance the reliability of our cognitive processes. Still, the upshot is that our concrete investigations will focus on precisely those factors that have been highlighted by internalists.

Initially, an externalist might be tempted to adopt a *black box approach*: to treat the whole complex practice or institution which is up for evaluation as an epistemically unstructured entity, focusing exclusively on the relation between its input, output and external conditions. Yet if the practice or institution is found to be epistemically defective, we would like to explain the unreliability and locate its source. And if it is found to be good, we would like to explain what makes it good in order to be able to improve it further, copy it or apply it to new domains. In principle, this might perhaps be done by treating it as a purely causal mechanism.[9] But in practice, the only feasible strategy will be to view it as composed of sub-processes which are themselves described in epistemic terms and seen as connected by epistemic relations. We will have to see whether the individuals involved fulfil their epistemic obligations and follow correct rules of inference and argumentation, whether at least some of them have higher-order knowledge, whether the transmission of beliefs among them meet the criteria for testimonial justification and so on.

In this way internalist criteria will serve as proxies for assessing the epistemic value of a process or institution – save for the rare cases in which it might turn out that a certain internalist criterion is not really an indicator of a truth-conduciveness. More realistically, insofar as the internalist criteria are only *prima facie*-principles or rules of thumb, they might sometimes be overruled. For example, breaking the standard rules of good argumentation may be the epistemically best thing to do in some contexts. Yet it is worth noting that when it is actually argued that e.g. the traditional ban on appealing to authority is not always to be respected, this is done not simply by pointing out that appeals to authority meet pure externalist criteria, a claim for which

[9] I am here ignoring the general problem of providing a reductive analysis of epistemic concepts, though it could certainly also set limits to the black box-approach.

sufficient empirical evidence would also be extremely hard to obtain, but rather by arguing that it is an inductively correct argument, i.e. that it is really an instance of another practice which internalists *do* recommend (Salmon 1963: 64; cf. Goldman 1999: 150).

Accommodating both externalist and internalist notions[10] will no doubt make the task of applied epistemology more laborious, since it becomes necessary to prioritise different forms of knowledge, justification and warrant and explore their interrelations. But this is something which needs to be done anyway. Many of the current debates over educational and research policies are about the relationship between and value of different forms of knowledge. What is most important: practical skills, factual knowledge, theoretical knowledge or reflective knowledge? When is it good enough merely to have or transmit a true belief, when should we take steps to secure that it be justified, and when are we required to have or provide knowledge of the belief-forming process itself, including its epistemic properties? According to an influential view on liberal education, enculturating children into the world of ideas (i.e. providing them with metaknowledge) will enable them to find out for themselves the lower-order knowledge and skills they will ultimately need (Bereiter 2002). This is the sort of claim an applied epistemology should be able to assess, and it can hardly be done by sticking closely to externalism.[11]

The Value of Justification

That the difference between externalism and internalism turns out to be rather insignificant in practice is confirmed by the way Goldman actually proceeds. By employing his veristic framework he is led to recommend the application of Bayesian reasoning, seriously considering adopting inductivist or coherentist approaches to testimony, defending a version of the total evidence principle (1999: 145), demanding that premises-conclusion relationships be displayed perspicuously in argumentative discourse, and so on. His whole

[10] Alston (2005) defends a similar pluralist approach to epistemic evaluation.

[11] Morton (2006) makes a similar point, albeit in a different context: "… in general one will not know if one's belief is formed by an acceptable process; we're better off considering simply how adequate the evidence is" (127).

discussion is pervaded by an acute sense of the importance of having and finding evidence. This is particularly vivid in his treatment of law (300ff.).

Still, Goldman does attempt to keep his framework very simple. The most striking example of this is perhaps his suggestion that in the context of social epistemology one might ignore the justification requirement entirely and adopt a weak notion of knowledge as *true belief*. He defends this view by pointing out that people's dominant epistemic goal is to obtain true belief, that is, to be *informed* (1999: 24; cf. also 2002b: 185ff.).

At a first glance, this seems intuitively right. I sometimes want to be informed about the sport results or a colleague's phone number. I do not care whether I will be justified in my belief – I just want to know the truth, i.e. to be aware of what is actually the case. In many situations it would be a waste of time to provide justification for true beliefs people are inclined to adopt anyway. A particularly vivid example of a system of belief transmission that deliberately ignores justification requirements for the sake of efficiency is the typical military system of communication. But there are countless other examples.

Yet closer reflection reveals that the matter is not that simple. First, as I pointed out above, there are also many situations in which we *do* care about justifying our own beliefs and those of others, about reflectivity and critical thinking. Hence at a minimum, the weak concept of knowledge will have to be supplemented by a stronger, justification-involving one, perhaps by several different concepts of knowledge, some of which are likely to be of an internalist sort.

Secondly, even in those cases where we do apparently only care about being informed, it is likely that our caring really embodies a tacit concern for justification. After all, according to most theories, the whole point of justification is that it is a means to truth (Goldman 2002a: 53ff.). Because I want to be informed about the sport results, I turn to the sport pages of a newspaper which I take to be a reliable source, instead of, say, merely flipping a coin in order to decide which team won the game.

It is not just that Goldman cannot dispense with the justification requirement (Le Morvan 2005). There is no good reason why he should try to.

If a certain practice is apt at making people informed, then it will meet Goldman's own reliabilist requirements and thus suffice to make the beliefs it produces justified. This is not to deny that there is may be point in distinguishing between (more or less) *pure information practices* and (more or less) *qualifying practices,* i.e. practices that aim not only at generating or transmitting true beliefs but of generating or transmitting justification for them as well as well. But even the purest of information practices cannot avoid conferring producing justification of the basic externalist sort if it is good at all from a veristic point of view.

These considerations show the relevance to applied epistemology of what has come to be known as the *value problem* (Zagzebski 2003: 13; it also referred to as the *Meno* problem): What is it about knowledge that makes it better than true belief *simpliciter*? The problem is that there seems to be no significant difference: true belief is as good a guide to action as true justified belief (*Meno* 98c), and whatever might be inferred from the latter could also be inferred from the former. But few have been willing to conclude that the justification requirement is really superfluous. A brief survey of the various responses to the value problem is instructive, since it highlights some of the reasons why justification deserves a more prominent place in applied epistemology than Goldman is willing to admit.

The answer given by Plato in the *Meno* is that a justified true belief is likely to be more *stable* than one that is merely true: justification makes it "abiding". This sounds very plausible: When I have no justification for a belief, I will be prone to give it up easily. If I have merely been told that *p*, merely being told by someone else that non-*p* may suffice to make me change my mind; whereas if I had also been told *why* it is the case that *p*, I would probably show more resistance. Stability thus proves to be an important, but complicating factor. A practice is to be judged not only by its propensity for producing true beliefs, but by its propensity for producing true *stable* beliefs.[12] It may be better to transmit fewer true beliefs if one can thereby increase their stability. Good teachers adhere to this principle.

---

[12] Dretske makes a similar point when he describes how a piece of information might help to *causally sustain* a belief (1981: 88f.)

Plato's answer is apt, but it cannot be the whole story. As Craig has pointed out, many of our beliefs are so closely related to a particular occasion (e.g. my belief that it is now raining) that there is no point in wanting them to be stable (1990: 7). A more recent answer to the value problem has been suggested by Colin McGinn, who claims that knowing *p* involves being good at getting true beliefs on a range of associated questions (1984: 540). There is also something to this. We often take care to justify a particular belief even though we are already convinced of its truth, because it enables us to acquire *other* true beliefs. Students are not only taught the correct solutions to arithmetical problems, but also general methods which they can use for solving further problems themselves and thus increase their cognitive flexibility. Though this aspect of justification – its *fruitfulness* or *adaptability* – is distinct from its stabilising power, the two properties can have a similar function. If a true belief should be lost, possession of a general method of justification might help one to quickly regain it.

There are counterexamples to McGinn's proposal, for example that one might know his own name without being able to find out anyone else's (Craig 1990: 55). However, since I am not trying to single out a defining property of justification, but exploring the various kinds of value which (various sorts of) justification might have, this is irrelevant to my purpose.

A third response to the value problem likewise suggests a widening of the evaluative focus: according to virtue epistemology, knowledge is better than true belief *simpliciter* because it reflects the virtuous features of the agent's belief-forming activity (Zagzebski 2003; Kvanvig 2003, 76ff.). Again, I can ignore the question whether the justifiedness of an individual true belief can really be derived from the virtuous character of the believer. The point is merely that our epistemic evaluations often aim not at the particular belief (or set or type of beliefs) produced in the actual case, but rather at the overall quality of the process, system or agent that produced it.

My own complementary response to the value problem consists in noting that it is actually misleading to ask whether a justified true belief is preferable to true belief *simpliciter*. For we have no way of establishing the truth of a belief other than justifying it. The handy expression "justified true

belief" gives the false impression that justifiedness and truths are parallel properties of a belief that can be accessed independently of another, whereas in reality, justification is the means (or road) and truth the end.

In a social context, however, it certainly can make good sense to treat truth and justification as distinct. An expert might judge a social system of belief transmission by its propensity for transmitting true beliefs which she has herself acquired in advance, by quite different methods. And she might distinguish systems that are apt at propagating true beliefs from systems that are apt at transmitting justification (cf. the distinction between pure information practices and qualifying practices I made above). But then again, if an information practice is reliable, the beliefs it induces in the otherwise ignorant recipients should still possess at least externalist justification. There is no way to dispense with the notion of justification in applied epistemology, and no way to keep the justification issue simple.

Truth and Adequacy

I have been assuming that the basic epistemic value is truth, or something closely analogous to that. But it is unlikely that a full-fledged applied epistemology can actually do with positing only one basic normative factor. One is pulled towards a moderate kind of value pluralism, not just because various kinds and functions of justification have to be taken into account, but also because there are other relevant forms – or *formats* – of *cognition* than propositional knowledge. For example, *understanding* is usually considered an epistemically good thing, and it is least an open question to what extent it can be analysed in terms of propositional knowledge. Moreover, some non-propositional forms of knowledge are clearly relevant to applied epistemology. Studies of professional practice in domains like law, medical practice and management have focused on the interplay between tacit and non-tacit knowledge (Sternberg and Horvath 1999). A much-debated topic in contemporary theory of education is the use of media that are not purely linguistic, like pictures, films and music.

The good news is that we do not have to move far away from the truth-based approach in order to accommodate non-propositional knowledge.

There are various proposals on offer as to how it might be done. Stanley and Williamson (2001) have recently argued that there is no fundamental distinction between knowing how and knowing that. Knowing how can be accounted for in terms of knowing that: when we say that Hannah knows how to ride a bicycle, what we really mean is that she knows that there is a certain way *w*, which is a relevant way for her to ride a bicycle, and that she knows it under a special – practical – mode of presentation. Yet I think this proposal goes too far in its attempt to assimilate practical to propositional knowledge. As Rosenfeld (2004) has pointed out, the central notion of a practical mode of presentation is left unexplained by Stanley and Williamson; and the most plausible explanation is that knowing how to ride a bicycle under a practical mode of presentation is the same as having an ability to ride a bicycle. But then it seems that the proposed reductive account still contains an element which cannot be reduced to knowledge that. One lesson to be learned from this is that the mere fact that the linguistic expressions "knowing that" and "knowing how to" are often interchangeable (*mutatis mutandis*) – that they can be used to describe the same state of affairs – does not tell against the view that they do at root designate two fundamentally different kinds of knowledge.

Craig's "good informant"-approach seems more promising. Craig does not attempt a reduction, but points to a close similarity of structure between knowing how and knowing that. He notes that both notions involve a success-clause plus a clause indicating that the success is no accident (1990: 161). And both notions reflect our need for good informants on various questions. Just as we need true beliefs, we need capacities to act; and in order to acquire the latter, we often turn to people who can show us how to do the thing in question. Apart from being able to do it themselves, these people must be prepared to display their capacity and have some property which indicates their possession of it (159). Hence "know how to" is not synonymous with "can"; it is not a pure ascription of capacity. I take this to be an attractive feature of Craig's account, since it matches the intuition that knowing how to X is not merely being able to X (and thus explains that we sometimes say things like "How did I do that!?"). The expression "knows" is,

in this as well as in other domains, a predicate referring to a particularly *qualified* state or activity.[13] On the other hand, Craig's account cannot be charged with "intellectualism", since it does not imply that in order to know how to X, the agent must be able to *explicate* how to do X or possess propositional knowledge of her ability to X.

My own suggestion, which I take to be in line with Craig's account, is that all forms of knowledge involve a property that is similar to truth – that is, they all involve (i) an *adequacy* requirement (what Craig calls a "success-clause"). They moreover involve a property similar to justification, i.e. (ii) a *qualification* requirement. And though they are not all belief states, they must still be forms of (iii) *representation.* This last requirement is probably the most controversial part of the three, and it is surely negotiable. But it seems to me that if the agent does not have some kind of access to her capacity (or rather to the corresponding task, i.e. to X), some kind of awareness of it, then we would not ascribe to her any knowledge of how to do X.

Whereas truth is a binary notion, the notion of adequacy is gradual. But the difference is not that big. It should be remembered that for concrete purposes – say, when it comes to arguing for the truth of scientific theories – the gradual notion of *truth-likeness* is often preferred to the strict binary notion of truth (see e.g. Devitt 1991: 125; Psillos 1999: 277). And when it comes to non-propositional knowledge, the appropriateness of a gradual notion of correspondence is obvious. One's practical knowledge – i.e. one's know-how – can be more or less adequate to the corresponding task. I know how to play some easy piano pieces, but much less adequately than the professional musician. My seven-year old son already knows how to write, but his knowledge is further improving day by day. Adequacy is particularly pertinent to *pictorial* knowledge (i.e. knowing how something looks): I know how my wife looks, how George W. Bush looks and (presumably) how Julius Caesar looked, but to very different degrees. I may know how something looks in certain respects but not in others, e.g. because I only know it from black-and-

[13] Though I do not want to deny that the locution "knows how to" is sometimes used in a pure "skill sense" (see Hintikka 1975: 11).

white photographs or abstract sketches.[14] *Maps* and *diagrams* are also representational formats which an applied epistemology should be able to assess, and this again will have to be done terms of adequacy and accuracy (cf. Kitcher 2001: Ch. 5.). Even *understanding*, that most intractable of all epistemic notions, may admit of an analysis along these lines (Elgin 2006: 215).

Substituting the more general notion of adequacy for truth thus makes possible a unified treatment of different forms of knowledge, and helps making epistemology applicable to a range of socially important questions which it would otherwise have nothing to say about. And my proposal is still in the spirit of veritism. Like truth, adequacy is a matter of correspondence, of fitting the facts. One may even imagine it to be analysable, at least in principle, in terms of partial reference or truth (or, in the case of practical knowledge, in terms of a categorical ability to perform certain sub-tasks, i.e. write certain words or play certain parts of a piano piece).

But introducing adequacy as an evaluative dimension does make applied epistemology more complicated. It can no longer be a matter of merely judging processes by their propensity for producing true beliefs (or, as Goldman suggests, the ratio of true to false beliefs). We might also have to judge the adequacy or accuracy of the representations they produce. And this is likely to engender new problems. Should we prefer crude, but mostly correct representations? Or do we want representations that capture more subtle features of the object, even at the price of increasing the risk of error? Elgin (1988) has pointed out that both externalism and internalism favour the employment of crude categories: if cognitive excellence were only a matter of producing truth and avoiding error (or, in the case of internalism, of forming beliefs that are supported by one's evidence), we had better keep to forming very unspecific beliefs, like "x is a bird", rather than "x is a starling". But this is a of course *reductio* of the view that only the truth-ratio matters. Urging the use of cruder categories is hardly the way to improve our epistemic practices.[15]

---

[14] For an insightful discussion of the selectivity of pictorial representation, see Lopes 1996

[15] Although adopting a more general category may sometimes count as an epistemic improvement. A method for estimating the national unemployment rate may be deemed superior to a method for estimating the unemployment rate in the neighbourhood or among 31-year old redheads. This points to the role of interests in our epistemic evaluations, a question to be addressed below.

Interests and Interaction Principles

The above considerations show that even an applied epistemology which is basically externalist and veristic will have to recognise a wide range of normative factors: truth, other kinds of representational adequacy, error avoidance, various kinds of justification, belief stabilising and fruitfulness – and this is probably just the top of the iceberg. The recognition of such a plurality of normative factors engenders a need for ranking or *interaction* principles (Kagan 1988: 183) that can regulate their application; and these principles must be considered normative factors in their own right. What are we to do about this?

Goldman takes the easy way out by focusing almost exclusively on local processes with a very specific goal – e.g. weather forecasting done by amalgamating individual forecaster's predictions of rain versus non-rain (1999: 81). This is an easy case because the goal can be described in simple veristic terms (i.e. being right about whether it is going to rain, a simple yes-or-no question), and only the amalgamating process – the way individual judgements are merged – is evaluated. That is, Goldman frames the question in a way that leaves out of consideration alternative strategies like, say, letting the forecasters form their judgements in a collective research and discussion process. Nor does he consider the possibility that the less reliable forecasters might be better at avoiding falsehood (e.g. that they are more prone to withhold belief in difficult conditions), that they would be perform better if they were not forced to give a definitive answer, or that their errors, though more frequent, are less serious, since they are confined to cases where it is only going to rain a little, etc.

I am perfectly aware that real-life cases are usually much more complex than the schematic examples discussed by philosophers, and that the whole point of philosophical analysis is to abstract from this bewildering complexity, breaking it down into bits that can be handled more easily. I do not advocate a position akin to particularism in ethics. It is central to the notion of an applied epistemology that there are general norms of epistemic goodness and that they can be illuminatingly applied to real-life cases. Nor do I wish to deny that

the weather forecaster example can function as a model for important real life-investigations and improvements. I use the example to underline my point that an applied epistemology needs interaction principles and that it must deal with the interplay between epistemic and non-epistemic values. Otherwise it will have limited application and, ironically, leave the important decisions to purely pragmatic considerations.

Goldman is not quite clear on the relationship between epistemic and non-epistemic values. But his general maxim is to keep things simple by avoiding or minimising the introduction of non-epistemic values and staying away from problems of metaevaluation. He explicitly declines to offer any ranking of epistemic values and evaluative dimensions (1986: 181). And he seems even more reluctant to comparing epistemic with other values, something which he seems to think is a job for moral philosophy only.[16] He even claims that whenever epistemic value conflicts with moral value, epistemic value must give way (2002a, 68). This sounds plausible, inasmuch as moral values are usually taken to be more fundamental and moral norms more genuinely compelling than epistemic ones. Still, is it not possible that a great epistemic value could trump some slightly negative moral value (see [author 2009])? At any rate, it seems that such conflicts call for a kind of metaevaluation that can hardly be the exclusive province of moral philosophy. Even if it is maintained that whatever real (i.e. practically compelling) value an epistemic practice or its outcome may have will *eo ipso* be some kind of moral or prudential value, the translation from epistemic to non-epistemic value will require epistemic expertise. In order to decide how much weight should be assigned to, say, a certain state of enlightenment or understanding, one must know something about its epistemic credentials.

At some places Goldman does seem willing to accommodate non-epistemic values. In *Epistemology and Cognition*, he identifies three evaluative dimensions: reliability, power and speed. Only the first of these is purely epistemic. The notion of power is a mixed one, as it links reliability to a concern for *relevance* or "responsiveness to current goals" (1986, 124). And the notion of speed should be considered irrelevant from a strictly epistemic point

[16] At one place he calls it an SEP ("somebody else's problem) (1995: 172-73)

of view. Why should it be epistemically better to reach the truth in ten rather than in twenty minutes? I think however that Goldman is right to insist that

> … many cognitive goals, or tasks, can be conceptualised as finding an answer to a given question by a certain time (*loc cit*).

But this suggests that cognitive or epistemic goals are *generally* of a mixed sort, that they should always be conceptualised as finding an answer to a given question relative to certain further constraints, for example as finding an answer which is more or less precise, adequate, stable, accessible etc. It is not that these constraints are purely pragmatic; I think they are mainly epistemic. But they cannot be derived from veritism alone; and at least some of them are likely to contain a practical component as well. For example, adequacy is obviously an epistemic notion; but which particular degree and kind of adequacy is the best in a particular situation is likely to depend on our practical interests.

Goldman does admit interests into his framework (89). But he assigns to them merely the marginal role of a threshold condition: in order to be a candidate for veristic evaluation, a practice has to produce answers to questions that are of some interest either to the agent or an institution (95f.). As soon as this requirement is met, interests drop the out of the picture and the evaluation proceeds in an exclusively veristic manner.

Accordingly, Goldman hardly mentions interests in his discussions of specific practices. They begin to surface towards the end of the book where he turns to the topic of education. Goldman notes that there might be a conflict between student's actual interests, their potential interests (e.g. what would interest them if they were to learn some additional facts) and the interests of institutions (e.g. a school or society as a whole) (350). But apart from making these – admittedly useful – distinctions, Goldman does not seem to consider the resolution of interest conflicts a job for applied epistemology. He does suggest that one way to do it might be to appeal to differences in *veristic significance*, that is, to distinguish between more or less *important* truths (369). But because he finds it doubtful whether epistemic importance can be

quantified in an interest-independent way, the issue is handed over to political or pragmatic decision procedures.

Goldman defends his view by noting that if interests were given a more prominent role, the "specialised, veristic mission" of epistemology would be abandoned in favour of a more pragmatic enterprise (95). This is surely a legitimate worry. It can, however, be argued that by insisting too stubbornly on keeping epistemology pure, Goldman actually diminishes its domain of application.

According to Goldman's conception, applied epistemology is a process involving two discrete steps. First, a set of target practices – of a very specific sort – are selected on purely pragmatic grounds. We then go on to evaluate them veristically and, on that basis, to criticise, abandon or modify them. Whenever a certain practice has been deemed good, acceptable or unacceptable, we move back to the pragmatic stance and do all further prioritising and selection from there. The good thing about this division of labour is that it manages to keep the realms of epistemic and pragmatic evaluation apart. The bad thing is that it does so at a very high price: the interesting decisions are all left to pragmatic considerations. Yet if one thinks that there is such a thing as genuine epistemic value and that epistemic evaluation should play more than a marginal role in our practical decision-making, then one should insist that our selection and prioritising should also be guided at least in part by epistemic considerations. Goldman's strict division prevents this. He is keeping applied epistemology pure by keeping its scope very narrow and reducing its practical impact.

But is there a viable alternative? I think so, even though I am not able to spell it out in complete detail. What I have in mind is this. The general framework for an applied epistemology should be uncompromisingly veristic. Truth, and more generally adequacy, should be maintained as the fundamental factor in all epistemic evaluation. But it should be supplemented with a set of conceptually related, but distinct evaluative dimensions and, as far as possible, with interaction principles that can regulate their application and resolve

tensions between them. These interaction principles should themselves be conceptually related to the ideal of truth or adequacy.[17]

I have already suggested several supplementary criteria, though they may be in need of further clarification and the list is far from complete. Finding some plausible and useful interaction principles presents a more fundamental difficulty. One should not be overly ambitious. If such principles are to be found at all, they are likely to reside in the grey area between the epistemic and the non-epistemic and to be less clear-cut than one could have wished for. And if we look to ethics, it turns out that although the need for interaction principles has been more widely recognised here, there have been few attempts to state them explicitly. Some have argued that the important thing is to recognise the need for and possibility of a meta-evaluation – that is, to ensure that the various values and factors are somehow commensurable – whereas one does not need a set of distinctive meta-principles. Here is how Griffin puts it in his seminal work on the value of well-being:

> We can sometimes rank pleasures as to intensity, duration, and number. We can also measure degrees of one component of happiness (say, contentment with one's lot) and separately measure degrees of another, radically different component (say, achievements of one's ambitions). We seem therefore to have several partial scales – scales for certain dimensions of happiness – but no super-scale. But it does not follow from there being no super-value that there is no super-scale (1986, 90).

Griffin's point is that we may work out trade-offs between different dimensions of happiness without assuming that there is any single factor they have in common. We may weigh them against each other directly, on the basis of their own nature. Something like this may also apply to the case of epistemic value. I do not consider the case quite analogous, because epistemic values are arguably less primitive than the basic prudential values, like different kinds of pleasure and pain, and thus the prospects for analysing them further and extracting more general principles appear somewhat better. Still, what we can hope for is surely not to find a simple algorithm that will always yield a definite and incontestable ranking of epistemic values. What we

[17] Goldman himself makes a similar suggestion in his 2002a.

can do is rather to pin down a few *prima facie* principles that may serve as guidelines for the meta-evaluation, but which are themselves neither purely epistemic (or context-independent), nor regulated by any higher-order principle, apart from their conceptual connection to the value of truth or adequacy.

Epistemic Significance

As Goldman himself suggests, a plausible criterion of epistemic *significance* or *importance* would enable us to rank different procedures that would otherwise appear epistemically indiscernible or incommensurable. But many philosophers share his pessimism about finding criteria that are sufficiently interest- and context-independent (cf. Helm 1994: 116f.; Kitcher 2001: 113). Of course they are right that whatever criteria we may hope to find will be interest- and probably also context-dependent. But their pessimism still strikes me as exaggerated. There are many different kinds of interest-dependence, and not all of them need to threaten the mission of veristic epistemology. Moreover, the pessimism may rest on a mistaken assumption that epistemic significance must be a strongly compelling value which is capable of trumping other values. It is quite possible that epistemic significance can be – and is often – outweighed by pragmatic values *without thereby being eliminated.*

A good place to look for inspiration is the philosophy of science, where there is a long tradition of trying to lay criteria for the "goodness" of theories and hypotheses. Philosophers of science also tend to be less anxious about bringing together epistemic and pragmatic considerations.

Isaac Levi suggested that a fundamental epistemic aim is to acquire *new information* (Levi 1969). This aim may justify *risking error*. Though it is no super-value above truth, the demand for new information does thus seem capable of regulating conflicts between our narrow interest in truth and other epistemic and semi-epistemic desiderata: we may trade off reliability against power by assessing the novelty of a piece of potential information and compare it with the risk of being wrong.

The concept of novel information is complex, as it has both a "pragmatic" and a more "epistemic" dimension. *Novelty* is a highly context-

sensitive property: the novelty of a certain piece of information is relative to a believer, a society or a time. Still, novelty is *not* a pragmatic property in the sense that the information needs to be useful in any way. And it should be noted that what novelty is relative to is, basically, other instances of *knowledge*, i.e. a thoroughly epistemic factor! Besides, it is not to be seen as an independent value. The concept of novel information is multidimensional and capable of analysis, but it cannot be split into separate components. Only as a property of some piece of *information* does novelty contribute to epistemic significance. Contrary to certain "postmodern" assumptions, novelty *alone* has no epistemic value whatsoever.

The information requirement is closely connected to the *adequacy* requirement and thus adds a more "veristic" and epistemic flavour to the criterion. By gaining more information about something, one acquires a more adequate or "fitting" description of it. And it seems quite intuitive that gaining more information is a good thing. Knowing the complete phone number of a new acquaintance is clearly preferable to knowing only the first three digits.

But novel information cannot be the sole criterion of epistemic significance. We are pulled in the opposite direction by considering how we often value knowledge of *general* states of affairs. Indeed, according to some traditional views on epistemic significance, a truth is to be considered more valuable the more general it is. For example, Aristotle equates "the most precious truths" with judgements about things that are "universal and necessary" (Aristotle 1976: 1140b33-1141b8). It is easy to dismiss this as being merely an intellectualist prejudice. And surely generality *alone* cannot be the relevant criterion. There are general truths which appear trivial and un-illuminating, e.g. "All yellow things are coloured". And knowledge of *higher-order* states of affairs is not always valued more highly than knowledge of their lower-order counterparts. An example of this is the truth-regress: from my knowledge of *p* I may infer that it is true that *p*, go on to infer that it is true that it is true that *p*, and so – but this does not count as any real epistemic accomplishment.

Nevertheless, the old philosophers were on to something. Knowledge of deep mathematical truths is still considered very valuable, despite the fact that

these truths are assumed to hold in all possible worlds and thus contain almost no information (since they do not serve to eliminate any real possibilities).[18] Deep natural laws fare only slightly better, but knowledge of such laws is also valued very highly, usually much more highly than knowledge of particular phone numbers or the precise location of rocks and buildings.

I think the traditional view of epistemic significance is very much alive as a part of contemporary common sense. We may be paying more attention to the pragmatic aspects of knowledge, but we continue to distinguish routinely between epistemic practices and results which are somehow good in themselves and practices whose goodness depends on their practical context. Basic research is still considered the finest kind of science, despite the fact that applied science appears more useful and is often more rewarding in terms of social recognition and financial support. (If it should be objected that our instinctive preference for deep and general truths have no rational foundation, I will answer that the same might hold for, say, many of our basic moral values. Ethical theorists also accept widespread intuitions as *prima facie* evidence for or against their theories). This accords with the traditional view of epistemic significance, which does not imply that epistemic values should trump the pragmatic ones. Aristotle reckoned that because we are humans rather than gods, and because our actions take place in particular circumstances, we will often have an overriding interest in particular states of affairs (1976: 1141b3-21).

Hence there is no reason to dismiss the traditional view of epistemic significance as a remnant of Platonism. It points to a genuine evaluative dimension apart from the novel information requirement. But as we saw, not all kinds of generality or necessity are relevant. A popular and useful suggestion along the lines of the traditional view is that significant truths are truths which are *explanatorily basic* (Moser 1989: 223ff.; Goldman 1999: 369). Natural laws figure prominently in the deep explanations of many states of

[18] Dretske (1981: 3ff.) gives a good introduction to the mathematical theory of information. Of course it is possible to use the concept of information in a wider, less formal sense that also covers the case of mathematical knowledge. But it should be kept in mind that the notion of epistemic significance covers different – and even potentially conflicting – qualities; and this is made easier by using the concept of information only in the more restricted sense.

affairs, whereas un-illuminating generalisations do not. But again, this cannot be the sole criterion. There may be general truths which are far from trivial, though they can hardly be considered explanatorily basic. It is good to know that all men are mortal; and of course it may also enable me to explain that some particular man is mortal. But such an explanation is not deep, and it seems contrived to say that its goodness derives from its potential explanatory function. Simply forming a belief about the general state of affairs that all men are mortal appears to be an epistemic accomplishment in is own right.

Indeed, on some of the currently dominant views of explanation, only causal factors can be genuinely explanatory (see e.g. Psillos 2002), and this might rule out mathematical truths.[19] I am not sure if it is possible to formulate a precise criterion which is broader than the demand for explanatory basicality but still captures the intuitive distinction between important and unimportant generalisations. But again, we are not completely left in the dark. Further guiding principles naturally suggest themselves. Helm considers that a measure of epistemic importance may be "degree of entrenchment" (1992: 135), a notion which covers both explanatory and (purely) logical impact – the latter may account for the value of some of those generalisations that do little explanatory work. Elgin likewise suggests that we look to the organic structure of scientific knowledge (2006). We could thus take epistemic importance to be a matter of how central an item of knowledge is to one's web of belief. One can lend inspiration from such a view without buying into the radical holism which at least Elgin seems to advocate. Even if mathematics is considered an autonomous branch of knowledge, its almost universal applicability might suffice to justify its claim to epistemic importance.

A shift towards holism has the advantage of both making room for and shedding light on the notion of *understanding*. Understanding usually consists in a grasp or cognitive command of complex semantic, epistemic or causal relationships. And though it can be subjected to the adequacy requirement, it is not straightforwardly connected to the norm of truth. One may well

[19] On the other hand, there are some physical explanations that cite mathematical rather than genuinely causal factors (Lipton 2004: 31), so the causality requirement might be too narrow

understand a piece of fiction, like a novel or an ingenious but false scientific theory. This might be taken to mean that understanding is not an epistemic desideratum at all, but only a semantic one. But such a reaction is both unnatural and unnecessary. Intuitively, understanding *is* an epistemic desideratum; at any rate, it is one of the factors which an applied epistemology will have to trade off against reliability and other distinctively epistemic values.

And it does not take us far away from veritism: in order to understand something – say, a novel – adequately, you need must represent to yourself the semantic relations more or less as they really are.

Another way to reconcile the – otherwise internalist- and antirealist-flavoured – holistic view of epistemic significance with veritism is to notice that important truths are often those which are *ontologically* basic. Explanatory and ontological basicality goes hand in hand. Our total worldview or web of belief determine which elements of reality which are to be considered more central or fundamental, and this in turn sets an external standard for the evaluation of our specific belief-forming practices. This is also in keeping with the traditional view of epistemic significance. Knowledge of essential features has been considered more important than knowledge of its accidental features. Such a view could be acceptable even to those who are not metaphysical essentialists. It could be admitted that what counts as essential depends somehow on our interests, but that we have a more or less stable and pre-scientific understanding of it, perhaps something akin to the phenomenological notion of the "life-world" (Husserl 1962). Or it could be held that our conception of the world is constantly changing due to the impact of scientific inquiry. In both cases, practical interests would affect the underlying metaphysics, but not the epistemic norms themselves.

Finally, though I have treated the demand for *novel information* as a potentially conflicting requirement, it may nevertheless supply us with a further criterion for distinguishing significant from insignificant generalisations. For the bad thing about un-illuminating generalisations seems to be precisely that they do not provide us with any new information. They just give us more of the same. Good generalisations bring into focus patterns or properties of their subject matter which would otherwise have gone

unnoticed. This holds for real mathematical and logical discoveries – as opposed to spelling out obvious implications – as well as for real philosophical insights, as opposed to a mere recording of platitudes.[20]

Debating Descriptions

It thus seems that we have at our disposal a variety of principles that can guide us when trading off different epistemic desiderata and interests. Surely the application of these principles is not regulated by any more general principle and thus remains irreducibly *ad hoc*. But this is just what could be expected. Applied philosophy is never a simple top-down affair. As researchers in applied ethics remind us, we have to look to the facts of the particular subject matter. So even if we were to reach a consensus on the normative factors and interaction principles (perhaps something along the lines of my proposal), there would still be ample room not only for empirical investigations of the target practices, but also for more philosophical discussions about the significance of the empirical findings.

Though we cannot hope to resolve these debates in advance, we can make a conjecture about their general form. Discussions in contemporary epistemology are often about choosing between a local and a more global point of view, or between different levels of generality: what should be the primary evaluative focus – individual beliefs, belief-forming processes or the general capacities of epistemic agents or of social groups? Across how wide a range of actual or possible situations should an epistemic mechanism or practice be assessed? In how general terms should an epistemic mechanism or practice be described? I have already recommended a pluralist stand on these matters: there is no principled way of making such choices, and the different standards should be able to co-exist as different epistemic desiderata. But then the conflict will have to be resolved at the level of concrete judgements. This means that debates in applied epistemology are likely to be about finding the right scale – the right level of generality or locality. They will be debates over the most appropriate way to describe the target practices.

---

[20] An *explication* of platitudes or an elaboration of their constitutive role may of course be highly illuminating.

No doubt one will very often – perhaps always – have to appeal to non-epistemic factors. Maybe certain big questions about science and education policies are best viewed as political problems and should be tackled by devising suitably qualified democratic decision procedures (as suggested by Kitcher 2002). But one should not overlook that there is a large and important domain in between these very broad and complex questions and the narrow problems of optimising local truth-conducive practices on which Goldman focuses. Here, epistemology still has a major role to play, even if it does not reign supreme.

Consider the following examples. You might criticise the epistemic practice of an agent – say, an office worker – by pointing to its disadvantageous effect to the epistemic (i.e. veristic) output of the social entity (viz. the office or firm) or practice of which it is part, while recognising that considered in isolation, it is epistemically impeccable. Though this judgement surely depends in part on a pragmatic interest in the office or firm's functioning well, it is still mainly epistemic, as the normative factor on which it is based is the epistemic output of the social entity (and not, say, its economic success).

Or recall the example with Linda's habit of acquiring political information through reading and trusting a particular newspaper. Most would probably agree that it is unfair to criticise this habit on the ground that newspapers are not generally to be trusted, if the particular newspaper in question is highly reliable. But most would also agree that it is no good to defend Linda's habit by pointing out that the newspaper or commentator recently happened to be right on some marginal points. In this way we may move in on the appropriate level of generality. And while it is perhaps not able to give a general (sic) solution to the generality problem, there seems to be in concrete cases a fact of the matter as to what kind of description is most appropriate. At least we are able to discuss it rationally and independently of our particular interests.

Elgin is surely right that "phenomena do not dictate their own descriptions" (2006, 204). But as she recognises herself, this does not mean that our choice of description is arbitrary. As a matter of fact, phenomena *do*

to a certain extent dictate their own descriptions, but only insofar as they have already been apprehended or described in a certain way. As suggested above, our background beliefs about the essential and accidental features of objects and situations may affect our assessment of the corresponding epistemic practises. If these beliefs are correct, they may provide an objective (i.e. relatively interest-independent) basis for our epistemic judgments.

Nancy Cartwright has given a particularly striking example of how questions of scale are both important and susceptible to rational discussion.[21] The graduate school at Berkeley was accused of discrimination against women, because the probabilities of acceptance for men were much higher than for women. But this compromising pattern disappeared when the data were partitioned by department. It turned out that women tended to apply to departments with high rejection rates. Department by department, women were admitted in the same ratios as men (Cartwright 1983: 37). The moral of this example is certainly not that you can get any result you want by choosing an arbitrary level of generality. It shows that there are more or less appropriate descriptions, and that you can argue for the relevancy of a certain description by pointing to facts about the target practice.

We refine our description of the target practice with an eye to what kind of evaluation it is going to be subjected to. We are able to distinguish reliably and convincingly between more or less relevant descriptions. And when an appropriate description has been found, we have at hand a set of epistemic criteria, intimately connected to the basic value of truth, the application of which will yield a definite and arguably objectively result.

Moral Constraints?

I have conceded that epistemic values may often be outweighed by non-epistemic values. Indeed, it is not unreasonable to suppose with Goldman that epistemic value should always give way to moral value. But I think that Räikkä (2004) goes too far when he contends that epistemological evaluation should always be morally and politically constrained. Räikkä's own example is instructive: we might find out that the use of torture makes a legal practice

[21] An example which Elgin herself cites with approval.

more truth-conducive (2004: 66). According to Räikkä, this is uninteresting, because the use of torture is morally unacceptable. But even though I condemn the use torture just as strongly as Räikkä, I do not think we can say *a priori* that such a result would be completely insignificant. As a matter of fact, one of the most popular and convincing arguments in the current discussions about the possible use of torture in, say, cases of kidnapping or bomb threats is that it is highly unreliable.

On a consequentialist view of morality, a practice's capacity for producing potentially useful truths must be taking into account as one factor among others. And even on a deontological view, it would still be worth knowing whether a constraint like the ban on torture can be justified solely by appeal to our moral intuitions, or whether it also compatible with our epistemic interests. Moderate deontologists tend to accept that constraints have a threshold: if enough good is at stake, it may be permissible to act against them (Kagan 1998: 78). One could imagine a situation where the existence of a large part of our pragmatically useless knowledge (assuming that there is such a thing) were at stake and had to be traded off against something that would cause considerable harm to a few innocent persons. It is not clear that a moderate deontologist would embrace the constraint against doing when faced with such a choice.

Räikkä overlooks the possibility that investigations of actual practices may prompt revisions in our normative framework, even across different normative domains. The case of torture actually shows quite the opposite of what he claims. It is not that applied epistemology is made superfluous by antecedent moral reflection. It is rather that a potential moral dilemma can be resolved in advance by applied epistemology: if torture cannot deliver the epistemic goods, we need not speculate further about its moral status – at least if all we want to know is how we ought to act. But we may of course still have an intellectual interest in finding out whether it would also be wrong from a purely moral standpoint, just as we may have an intellectual interest in speculating about the epistemic potential of practices that are morally prohibited.

Applied epistemology overlaps with moral, political and practical evaluation. And there is or ought to be a real interplay between them, even if a moral or political overruling or correction of an epistemic judgement seems to be the more typical case. As mentioned at beginning of the paper, the classical defence of freedom of speech is a clear and uncontroversial case of arguing for a moral and political right by pointing to its epistemic efficacy. And there are probably quite a few other examples to be found.

Conclusion

I have tried to show how an applied epistemology can accommodate a variety of normative factors without abandoning its veristic mission. And I have argued against the attempt to keep applied epistemology pure by distinguishing sharply between epistemic evaluation and the pragmatic selection of practices to be evaluated. Selection procedures are crucial and should not be handed over to purely pragmatic considerations. The practical impact of an applied epistemology can be greatly improved by taking into account factors which are not purely epistemic, but not purely pragmatic either, as they do not depend on particular practical interests. There are two sides to this result. On the one hand, it means that there is room for significant investigations, since the application of normative epistemology need not be restricted to narrowly described practices of the sort Goldman considers. Applied epistemology can take on many of the hotly debated questions the use and production of knowledge in society. And it can do so without making major concessions to pragmatism. On the other hand, it also means that applied epistemology is going to be a very tricky business. Investigations will be multi-facetted, comprising formal and informal, conceptual and empirical studies; and they will be holistic, as the interplay between the various factors is crucial. The results obtained will often be controversial. Applied epistemology may improve the quality of the debate, but is unlikely to shorten it or cool it down. But I have no doubt that it will be worthwhile.

References:

Alston, William P. 1986. "Internalism and Externalism in Epistemology", *Philosophical Topics* 14: 179-221

Alston, William P. 1995. "How to Think About Reliability", *Philosophical Topics* 23: 1-29

Alston, William P. 2005. *Beyond 'Justification'*. Ithaca: Cornell

Aristotle. 1976. *Ethics*. Trans. J. A. K. Thomson. London: Penguin

[Autor]. 2009. "Two Notions of Epistemic Normativity". *Theoria* (in press)

Ayer, Alfred J. 1956. *The Problem of Knowledge*. London: Pelican

Beauchamp, Tom. "On Eliminating the Distinction between Applied Ethics and EthicalTheory", *Monist* 67 (1984): 514-531

Bereiter, Carl. "Liberal Education in a Knowledge Society", in Smith (ed.) 2002: 11-33

Bishop, Michael A. & Trout, J. D. 2005. *Epistemology and the Psychology of Human Judgment*.
Oxford: Oxford University Press

Boisot, Max H. 1998. *Knowledge Assets. Securing Competitive Advantage in the Information Economy*. Oxford: Oxford Univ. Press

BonJour, Laurence. 2002. *Epistemology. Classic Problems and Contemporary Responses*. 2002

Brady, Michael & Pritchard, Duncan (eds.). 2003. *Moral & Epistemic Virtues*. Oxford: Blackwell

Bruner, J. 1996. *The Culture of Education*. Cambridge, MA: Harvard

Cartwright, Nancy. 1983. *How the Laws of Physics Lie*. Oxford: Clarendon

Chisholm, Roderick. 1957. *Perceiving*. Ithaca: Cornell

Chisholm, Roderick. 1977. *Theory of Knowledge* (2nd ed.). Englewood Cliffs, N. J.: Prentice-Hall

Conee, Earl & Feldman, Richard. 1998. "The Generality Problem for Reliabilism", *Philosophical Studies* 1998: 1-29

Craig, Edward. 1990. *Knowledge and the State of Nature*. Oxford: Clarendon.

Daniels, Norman. 1979. "Wide Reflective Equilibrium and Theory Acceptance in Ethics": *Journal of Philosophy* 76: 257-282

Devitt, Michael. 1991. *Realism and Truth*. Oxford: Blackwell

Dretske, Fred I. 1981. *Knowledge and the Flow of Information*. Cambridge, Mass.: MIT Press

Elgin, Catherine Z. "The Epistemic Efficacy of Stupidity", *Synthese* 74: 297-311

Elgin, Catherine Z. "From Knowledge to Understanding", in Hetherington 2006: 199-215

Fricker, Miranda. 2003 "Epistemic Injustice and a Role for Virtue in the Politics of Knowing", in Brady & Pritchard (eds.) 2003: 139-158

Fumerton, Richard. 1995. *Metaepistemology and Skepticism*. Lanham, MA: Rowman & Littlefield.

Fuller, Steven. 1988. *Social Epistemology*. Bloomington, In.: Indiana Univ. Press

Goldman, Alvin I. 1986. *Epistemology and Cognition*. Cambridge, Mass.: Harvard

Goldman, Alvin I. 1995. "Social Epistemology, Interest and Truth", *Philosophical Topics* 23: 171-87

Goldman, Alvin I. 1999. *Knowledge in a Social World*. Oxford: Oxford Univ. Press

Goldman, Alvin I. 2002a. "The Unity of the Epistemic Virtues", in *Pathways to Knowledge*. Oxford: Oxford University Press: 51-70

Goldman, Alvin I. 2002b. "What is Social Epistemology? A Smorgasbord of Projects", in *Pathways to Knowledge*. Oxford: Oxford University Press: 182-204

Griffin, James. 1986. *Well-Being. Its Meaning, Measurement and Moral Importance.* Oxford: Clarendon
Hansson, Sven Ove. "Applying philosophy", *Theoria* 69: 1-2
Helm, Paul. 1994. *Belief Policies.* Cambridge: CUP
Herington, Stephen (ed.). 2006. *Epistemology Futures.* Oxford: Clarendon
Hintikka, Jaakko. 1975. "Different Constructions in Terms of the Basic Epistemological Verbs", *The Intentions of Intententionality*. Dordrecht: Reidel: 1-25
Husserl, Edmund. 1962. *Die Krisis der europäischen Wissenschaften und die transzendentale Phänomenologie.* Husserliana Vol. 6. The Hague: Martinus Nijhoff.
Kappel, Klemens. 2006 "A Diagnosis and Resolution to the Generality Problem", *Philosophical Studies* 127: 525-560
Kitcher, Philip. 2001. *Science, Truth and Democracy.* Oxford: Oxford University Press
Kvanvig, Jonathan L. 2003. *The Value of Knowledge and the Pursuit of Understanding.* Cambridge: Cambridge University Press
Le Morvan, Pierre. 2005. "Goldman on Knowledge as True Belief". *Erkenntnis* 62: 145-155
Lessig, Lawrence. 2001. *The Future of Ideas.* New York: Random House
Levi, Isaac. 1969. "Information and Inference", *Synthese* 17: 369-91
Lipton, Peter. 2004. *Inference to the Best Explanation.* 2. ed. London: Routledge
Lopes, Dominic. 1996. *Understanding Pictures.* Oxford: Clarendon
McGinn, Colin. 1984. "The Concept of Knowledge", *Midwest Studies in Philosophy* 9: 529-54
Mill, John Stuart. 1859. *On Liberty* (ed. C. V. Shields), London: Prentice-Hall (1997)
Morton, Adam. 2006. "Epistemology Meets Theory Choiche", in Herington 2006: 111-130
Moser, Paul K. 1989. *Knowledge and Evidence.* Cambridge: CUP
Pihlström, Sami. "Applied Philosophy: Problems and Applications", International Journal of Applied Philosophy 13: 121-133
Pollock, John. 1986. *Contemporary Theories of Knowledge.* Lanham: Rowman & Littlefield
Psillos, Stathis. 1999. *How Science Tracks Truth.* London: Routledge
Psillos, Stathis. 2002. *Causation and Explanation.* Chasham: Acumen
Räikkä, Juha. 2004. "Knowledge Promotion as a Criterion for Evaluating Social Institutions", *Sats. Nordic Journal of Philosophy* Vol 5, No 2: 59-68
Rawls, John. 1971. *A Theory of Justice.* Oxford: Oxford University Press
Rosenfeld, Tobias. 2004. "Is Knowing-how Simply a Case of Knowing-that?", *Philosophical Investigations* 27: 370-379
Salmon, Wesley. 1963. *Logic.* Englewood Cliffs, NJ: Prentice Hall
Smith, Barry (ed.). 2002. *Liberal Education in a Knowledge Society.* Chicago: Open Court
Stanley, Jason & Williamson, Timothy. 2001. "Knowing How", *Journal of Philosophy* 98: 411-444
Stich, Stephen. 1988. "Reflective Equilibrium and Cognitive Diversity", *Synthese* 74 (1988): 391-413
Stroud, Barry. 1994. "Scepticism, 'Externalism', and the Goal of Epistemology", *Proceedings of the Aristotelian Society (Supplementary Volume)*, 1994. Reprinted in *Understanding Human Knowledge*. Oxford: Oxford University Press: 139-154
Sternberg, Robert J. & Horvath, Joseph A. (eds.): *Tacit Knowledge in Professional Practice*
Williams, Michael. 2001. *Problems of Knowledge.* Oxford: Oxford University Press
Tuomi, Ilkka. 1999. *Corporate Knowledge. Theory and Practice of Intelligent Organisations.* Helsinki: Metaxis
Zagzebski, Linda T. 1996. *Virtues of the Mind.* Cambridge: Cambridge Univ. Press

Zagzabski, Linda T.:"The Search for the Source of Epistemic Good", in Brady & Pritchard (2003):13-28