Methodological Naturalism and Epistemic Internalism

Gregory Wheeler and Luís Moniz Pereira
Artificial Intelligence Center - CENTRIA
Department of Computer Science, Universidade Nova de Lisboa
2829-516 Caparica, Portugal

June 18, 2007

Abstract. Epistemic naturalism holds that the results or methodologies from the cognitive sciences are relevant to epistemology, and some have maintained that scientific methods are more compatible with externalist theories of justification than with internalist theories. But practically all discussions about naturalized epistemology are framed exclusively in terms of cognitive psychology, which is only one of the cognitive sciences. The question addressed in this essay is whether a commitment to naturalism really does favor externalism over internalism, and we offer reasons for thinking that naturalism in epistemology is compatible with both internalist and externalist conceptions of justification. We also argue that there are some distinctively internalist aim that are currently being studied scientifically and these notions, and others, *should* be studied by scientific methods.

Keywords: Statistical default logic, evidentialism, coherentism, logic programming, applied logic

This essay is dedicated to Deborah Mayo, who has long advocated using error statistical techniques to analyze and resolve epistemological puzzles in the philosophy of science. This essay follows the same spirit by advocating that computational concepts and techniques be applied within the heart of traditional, analytic epistemology.

1. Introduction

Traditionally, our pre-theoretic notion of epistemic justification is understood to involve two properties: *accessibility* and *truth-conduciveness*. Epistemic justification (hereafter, justification) is thought to be accessible in the sense that an agent S who is justified to believe a proposition p is in a position, even if only in principle, to access the item that justifies p—whether that item be a linguistic entity, precept, memory, or other belief. First-person accessibility is thought necessary for S to demonstrate or evaluate his reason for holding p, which is one role that justification is thought to play. That S is justified to believe that p appears to mean that S has a good reason for believing p, which suggests that having justification for a belief entails some capacity to view the items responsible for that justification in order to judge their bearing on that belief.

Truth-conduciveness concerns the contribution that justification appears to make to the possession of true beliefs. The value of a justified belief is not simply for someone to have a just-so story for his belief that p. Rather, its value appears to derive from a tendency for justified beliefs to also be true beliefs. The concept of justification therefore appears to also involve a belief forming procedure that typically results in true beliefs.

That truth-conduciveness and accessibility are difficult properties to reconcile within a single concept of justification is an important backstory for philosophical theories of epistemic justification. As a result, two fundamentally different conceptions of justification arise from taking each of these properties to be primary. How to investigate each type of theory is the topic of this essay, and our specific focus is the role that scientific methodology is thought to play.

It is common to frame the debate between first-person-accessible conceptions of justification and truth-conducive conceptions of justification in terms of *internalist* versus *externalist* theories of justification, respectively. But there are a variety of ways to construe 'internalism' and 'externalism'. A classical example of an internalist theory of justification is Roderick Chisholm's *access* internalism (Chisholm 1966), which holds that

(i) items that justify an agent's belief should be accessible to that agent, and

(ii) an agent may establish on reflection whether a particular belief of his is justified.

However, 'internalism' may also refer to a view that is more restricted than access internalism. *Mentalist internalism*, advanced by Rich Feldman and Earl Conee (Conee and Feldman 2004), is a version of internalism that restricts the items that may serve as justifications to mental states. Hence, *mentalist internalism* replaces condition (i) of access internalism with

(i') necessarily, only mental states justify beliefs and mental states that justify an agent's belief should be accessible to that agent.

Even so, the signature trait of internalist theories of justification is a commitment to an accessibility condition, such as (i) or (i'), and the first-person evaluation condition of (ii).

Externalism, by contrast, describes any theory that rejects either the accessibility condition, the first-person evaluation condition, or both. While the motivation for denying one or both of these conditions

typically comes from adopting a truth-conducive notion of justification, truth-conduciveness is not entailed by externalism. Even though truth-conducive theories are simply one type of externalist theory, we shall follow convention by using 'externalism' to refer to theories of justification that reject one of the tenants of internalism and embraces

(*iii*) An agent S is justified in believing p only if S's belief that p is produced by a process C such that the propensity of C to produce true beliefs is greater than the propensity of true beliefs to occur without C.

Condition (*iii*) states that C, a psychological feature of belief formation, is causally relevant to yielding true beliefs, and that C is a necessary property of justified beliefs.[1]

Now to methodology. *Epistemic naturalism* holds that the results or methodologies from one or another of the cognitive sciences are relevant to epistemology, and some maintain that scientific methods are more compatible with externalist theories of justification than with internalist theories.[2] Condition C in (*iii*) describes a cognitive process with a testable outcome, namely a purported propensity for C to yield true beliefs, and facts about cognitive psychology should inform what epistemologist imagine C to do. This alliance between psychology and externalism has left traditional internalists saddled, sometimes willingly, with defending a theory of justification that appears immune to scientific inquiry.

The question that interests us is whether a commitment to naturalism really does favor externalism over internalism. Cognitive psychology does not exhaust the cognitive sciences, after all. With this point in mind we argue here that naturalism in epistemology is compatible with either conception of justification, since (*a*) the content of an internalist theory of justification is compatible with naturalism, and (*b*) at least some distinctively internalist notions of justification are open to scientific study. One consequence of embracing these two theses is that the issue of epistemic naturalism itself is no longer a relevant factor to weigh in discussing the merits of how one should view the notion of epistemic justification.

But we propose to go one step further by arguing that (*c*) at least some distinctively internalist notions *should* be studied scientifically rather than by traditional methods of conceptual analysis and informal descriptions. With respect to (*c*), our proposal is to study the mathematical structure of epistemic support relations using both theoretical

[1] If one replaces 'only if' in (*iii*) by 'if and only if', then one has a basic description of process reliabilism.

[2] See (Kornblith 2001) and (Bonjour and Sosa 2003).

and experimental methods used by the logical artificial intelligence community. That one *can* do this is sufficient reason to accept that at least some distinctly internalist notions are open to scientific study (Wheeler and Pereira 2004). That one should do this rests upon practical advantages from applying modeling and verification techniques to similarly structured problem domains. We think that scientific methods have a constructive place within traditional internalist epistemology, and we base this conclusion upon the reasons we have for accepting (a), (b), and (c). We turn to those reasons in the remainder.

2. Substantial Naturalism and Internalism

Consider the first claim,

(*a*) The content of an internalist theory of justification is compatible with naturalism.

This claim concerns the relationship between internalism and *substantive* naturalism.[3]

Substantive naturalism in epistemology is the view that, to be meaningful, epistemic terms must denote natural facts—either directly, or by reducing suspect terms to natural terms. The motivation behind substantive naturalism is the sensible advice that one should be wary of analyses that appeal to properties or relations in the *analysans* that are more obscure than the *analysandum*. Substantive naturalism proposes to bar non-natural terms from appearing in theoretical accounts of justification since there is no way, in principle, for the notions they denote to be investigated by scientific methods.

A reason to think that internalist theories of justification are incompatible with substantive naturalism is that internalist theories of justification tend to rely upon terms that appear to be purely *epistemic*. Alvin Goldman has raised a version of this objection (1979, 1986), claiming that an analysis of epistemic justification should not rely upon other epistemic terms, such as "'justified", "warranted", "has (good) grounds", "has reason (to believe)", "knows that", "sees that", "apprehends that", "is probable" (in an epistemic or inductive sense), "establishes that", and "ascertains that"' (Goldman 1979) since the account would fail to give us purely factive conditions for the justificatory status of a belief. On Goldman's account, terms such as "'believes that', "is true", "causes", "it is necessary that", "implies", "is deducible

[3] For an overview of naturalism in epistemology, see (Kornblith 1999) and the edited collection (Carruthers, Stich and Siegal 2002).

from", and "is probable (either in the frequency sense or the propensity sense)"' are not evaluative, so may serve this purpose (Goldman 1979). It should be noted that Goldman's primary aim is to avoid having epistemic terms appear in the analysis of justification on the grounds that doing so would yield an incomplete or circular analysis, not on the grounds that purely epistemic terms are non-natural. However, substantive naturalism may be construed as saying that the epistemic terms that appear in traditional accounts of justification fail in virtue of failing to denote (or be reduced to terms that denote) natural facts. So, for instance, the evidentialist analysis of justification,

> S is justified to believe p iff S has evidence supporting the belief that p,

and deontological analysis of justification,

> S is justified to believe p iff S has the right (duty) to believe p,[4]

would both fail because each analysis appeals to epistemic terms that are, necessarily, non-natural.

One can find internalists that appear to endorse this non-naturalistic view. For instance, Roderick Chisholm held that epistemic properties and epistemic relations are *irreducible*, meaning that they are of a kind that simply cannot be defined by a complex of psychological or familiar logical operations (Chisholm 1966). Chisholm thought that the aim of epistemology was to study these relations, which, given his internalism, were first-person accessible. He thought that we could devise epistemic principles by reflecting upon these states, proposing principles that appear to capture their structure, and then test those principles by considering purported counter-examples.

The question before us is whether internalists must be anti-naturalists. Rich Feldman has addressed half of this question by considering whether there is any reason to think epistemic terms—such as those terms used by evidentialism—must denote non-natural facts (Feldman 2001a). Feldman notes that there isn't anything about the epistemic terms used in evidentialism—'evidence' and various terms used to describe evidential support relations—to suggest that they are ontologically mysterious. One may well conjecture that there is a relation called 'evidential support' that holds between a belief and a precept, a memory, or another belief. Assuming that the *relata*—beliefs, memories, precepts—are naturalistically acceptable, there doesn't appear to be *prima facie*

[4] It is worth remarking that these these two conceptions of epistemic justification are distinct, even though one will find arguments against evidentialist theories in the literature that presume that evidentialism is committed to a deontological notion of justification. For example, consider Sosa in (Bonjour and Sosa 2003).

grounds for regarding this proposed relation to be more objectionable on substantive naturalist grounds than the relations of causation or entailment. While traditional accounts of epistemic properties typically use epistemic terms and do not provide definitions in purely naturalistic terms, this does not entail 'that epistemic relations are not themselves natural relations or that naturalistic definitions of them are ruled out in principle' (Feldman 2001a).

Feldman's argument only addresses the ontological status of evidential facts and evidential support relations: the epistemic terms used by evidentialism are purported to pick out natural facts or natural relations. It remains to be seen whether there *are* such facts and relations, but one shouldn't doubt the kind of items these terms are thought to denote. Evidentialism is the claim that the world includes evidential support relations and evidential facts; it is not the claim that evidential facts or evidential support relations exist as mysterious entities or relations.

Feldman's endorsement of substantive naturalism stops short of embracing methodological naturalism, however. *Methodological naturalism* holds that methods and results from the cognitive sciences are relevant to epistemology. We consider now the case for evidentialism going whole hog.

3. Methodological naturalism and internalism

Consider:

(*b*) Some distinctively internalist notions and relations are open to scientific study.

In contemporary epistemology the debate surrounding (*b*) has centered on the role of *cognitive psychology* within epistemology. Among internalists, there is considerable resistance to the view that cognitive psychology can make contributions to a theory of justification.[5] By the internalists' lights, facts about any cognitive process responsible for belief formation are irrelevant to an agent's first-person assessment of his epistemic position, thus irrelevant to epistemology.

Nevertheless, one point that internalist and externalist epistemologists appear to agree on is that logic offers little analytical insight into

[5] Methodological naturalism in epistemology is traced back to Quine's 'Naturalized Epistemology', appearing in (Quine 1969). See (Kim 1988) for a critical reply to Quine's program and (Kornblith 1999) for a contemporary defense of the Quinean view.

the structure of relations and conceptions mentioned in each theory. As we said before, Chisholm thought that both logic and psychology were irrelevant. The debate since Chisholm has been over the role that psychology plays within epistemology, since externalists have been happy to follow Chisholm's line on logic and epistemology. Consider, for example, remarks from Hempel, Goldman, and Harman:

> Formal logic tells us that if a given set of statements is true then such and such other statements are true as well; but it does not tell us what statements to believe or to act on. Indeed, the notion of accepting certain statements, like the notion of total evidence, is pragmatic in character and cannot be defined in terms of the concepts of formal deductive or inductive logic (Hempel 1965, p. 66).

And Goldman, remarking on the claim that a system of rules derivable from logic and probability theory may determine an acceptable set of rules for epistemic justification (J-rules), writes that

> This [claim] accords with a widespread assumption that logic provides us with proper methods. Since 'proper method' is easily construed as 'justification conferring method', it is natural to assume that J-rules—at least those governing reasoning—can be derived from logic. This assumption is false (Goldman 1986, p. 81).
>
> There is no way, then, in which J-rules are literally *derivable* from, meaning entailed by, truths of formal logic (Goldman 1986, p. 82).

Finally, Gilbert Harman thinks that there is a sharp distinction between the psychological process of drawing an inference and the logical relation of implication, and that we court serious confusion by thinking that logic yields 'laws of thought':

> . . . to call deductive rules 'rules of inference' is a real fallacy, not just a terminological matter. It lies behind attempts to develop relevance logics or inductive logics that are thought better at capturing ordinary reasoning than classical deductive logic does, as if deductive logic offers a partial theory of ordinary logic (Harman and Kulkarni 2006, 560).

Be this as it may, it is a mistake to think that logic can play *no* role in representing and testing epistemic relations. The AGM theory of belief revision is no more a category mistake than Newton's theory of mechanical motion. The mistake that Harman, Goldman, Hempel and Chisholm are pointing to is to imagine theorems of logic to reveal epistemic principles, rationality constraints, or instructions for how we

ought to settle upon a view. On this point we may all agree (Wheeler, forthcoming).

But one may concede that logic is not the language of thought and still see that modern logic offers methods for articulating, studying, and testing hypotheses about epistemic relations and their properties. What epistemologists do, particularly internalist epistemologists, is offer descriptions of mathematical structures when they propose a solution to Gettier cases, a route to avoid the lottery paradox, an account of the basing relation or accessibility, or a theory of coherentist justification. There is then a need to understand the structure of these proposals and how they actually behave within a theory. These are issues that traditional introspection and thought experiments are not suited to address.[6]

We recommend the use theoretical results and experimental methods from logical artificial intelligence for the study of epistemic relations. One example of our work in this area is the study of the class of non-monotonic logics that best represent the form of classical statistical inferences (Kyburg, Teng, Wheeler, 2007), while another, which is influenced by Deborah Mayo's work, is how to represent individual classical inference forms in terms of statistical default logic (Wheeler 2004) and implement a testable model within logic programming (Wheeler and Pereira 2004, Wheeler and Damásio 2004). Let's consider this idea in some detail.

Statistical default logic offers both an analysis of the logical structure of individual statistical inference forms and provides a basis for computing error-probabilities for arguments composed of a sequence of statistical and deductive inference steps. Default logic is a non-monotonic logic formed by augmenting first-order classical logic with non-monotonic inference rules, called *defaults*, that appear in the object language. If α, γ and β_i are formulas in a first-order language, then defaults are inference rules of the form

$$\frac{\alpha : \beta_1, ..., \beta_n}{\gamma}. \tag{1}$$

Schema (1) is interpreted roughly to mean that given α and the absence of any negated β_i's, conclude γ by default. The β_i's in (1) correspond to conditions the *absence* of which, when α holds, allows γ to be derived. The non-monotonic behavior of defaults rests in the possibility that one of the default justifications that permits the rule to be applied may be

[6] Note that this point is recognized by Goldman, who concedes that knowing the semantic properties of such rules 'is undoubtedly *relevant* to belief forming principles' (1986, p. 82), but would seem to be resisted by Harman, Hempel, and Chisholm.

triggered by new information, thus blocking the applicability of that rule.

It turns out there is a structural similarity between the workings of default rules and a class of standard statistical inference forms. In making a statistical inference the aim is to select a sample that represents the population with respect to some specified parameter. Often this is achieved by a series of tests designed to detect bias in the sample. It was first noticed in (Kyburg and Teng 1999) that in making a statistical inference, some conditions are satisfied explicitly, like premises of a default, while other conditions behave like default justifications. Typically a sample is regarded representative of a population when a few explicit conditions hold (like that the sample be drawn from the target population and the distribution of error is normal) and when there is no reason to suggest that the sample is biased, which translates to the absence of information that would suggest a biased sample. Absence of evidence for bias expresses weaker assumptions than evidence for representativeness, for if we had direct evidence that a sample was representative of the population from which it was drawn then we would not need to perform statistical inference.

Default logic provides only half of the structure of a statistical inference, however, since there isn't a capability within the logic to distinguish between rules that rigorously probe for error and rules that let nearly any sample count as a 'good' statistical inference form. Another important feature of standard statistical inference is its emphasis on the control of error (Mayo 1996). Following Kyburg, we say that when making a statistical inference one *accepts* a conclusion along with a warning that there is a small, preassigned chance that the conclusion is false. A statistical inference controls error to the extent that its advertised frequency of error corresponds *in fact* to the chance one faces in making that inference and its conclusion being false. What is problematic about representing inferential statistical forms in terms of defaults is that there is no means to represent the error-probabilities of each statistical inference.

S-defaults differ from defaults by explicitly representing the *upper limit* of the s-default's probability of error.[7] Call a default in the form of

$$\frac{\alpha : \beta_1, ..., \beta_n}{\gamma}\epsilon, \tag{2}$$

[7] A trivial corollary of the probability of error $\hat{\alpha}$ for a statistical inference is the upper limit of the probability of error, denoted by ϵ. So, if $\hat{\alpha} = 0.05$ is understood to mean that the probability of committing a Type I error is 0.05, then $\epsilon = 0.05$ is understood to mean that the probability of committing a Type I error is no more than 0.05.

an ϵ-*bounded statistical default* and the upper limit on the probability of error-parameter ϵ an ϵ-*bound* for short, where $\frac{\alpha:\beta_1,\ldots,\beta_n}{\gamma}$ is a Reiter default and $0 \leq \epsilon \leq 1$. The schema (2) is interpreted to say that provided α and no negated β_i's, the probability that γ is false is no more than ϵ. (A Reiter default is a limiting case of a statistical default, namely when $\epsilon = 0$). A statistical default is sound just when the upper limit of the probability of error is *in fact* ϵ. An s-default is a good inference rule if it is sound and ϵ is relatively small, typically less than 0.05.[8]

One objection to inductive logic is the claim that there are no inductive arguments per se but that there is inductive reasoning, so attempts to capture this reasoning within an inductive logic must be predicated on endorsing 'logical psychologism', the thesis that logic is the 'language of thought'. It is one thing to propose to apply some logic to represent a feature of reasoning since this proposal is then subject to the same type of constraints that attend any other proposal to apply mathematics, one of which is the evaluation of how well the formal representation fits the problem domain. It is entirely another matter to interpret these formalisms to tell us something directly about reasoning. Harman (1986, 2002, 2006) makes much of the differences between the psychological act of making a rational inference and the syntactic conditions of applying an admissible inference rule within a formal system, as we've already mentioned. And there are several features of belief fixation and belief change that give pause to viewing logical methods to be the right branch of mathematics from which to construct a formal model of rational belief fixation and rational belief change. Even so there is no category mistake involved in advancing such a project. Statistical arguments are a type of inductive argument, having a structure that is quite apart from the psychological activity that goes on when reasoning about one. It makes perfect conceptual sense to investigate the structural properties of those arguments, and statistical default logic captures two of their distinctive features: the defeasible acceptance that conditions of a statistical model are applicable, and the management of error probabilities for sequences of inference steps. The 'inference-implication' fallacy applies to psychologism and thinking that epistemic principles are theorems of some system of logic. The fallacy is forced by particular applications of logic, not by applied logic as such. While it is a mistake to regard the study of inductive logics as the study of human inference, it is likewise mistaken to regard the study of inductive logics to be predicated on this very error.

[8] For more on statistical default logic, see (Wheeler 2004).

Studying epistemic relations involves two components. The first is a purely formal approach that studies the logical and computational properties of formal systems which one might consider applying to a particular problem. We illustrated one example of how the mechanism of default justification can capture judgments of normality in statistical reasoning. The second component concerns judgments of fit between a formal representation and the problem domain. Traditional epistemic internalism views the methods of intuitive introspection to be a sufficient methodology for judging the fitness of an epistemic theory. We think that contemporary systems of applied logic have outgrown this approach.

4. Naturalizing internalism

One reason that internalist epistemologists have been reluctant to embrace scientific methods is that cognitive science is understood to mean cognitive psychology, and internalists like Feldman don't think that cognitive psychology provides relevant information to epistemology. Feldman thinks that epistemology should examine evaluative questions and that detailed empirical results from the cognitive sciences concerning how we actually think and reason are not essential for making progress in addressing these normative questions.

But internalist epistemologies do propose theories involving epistemic relations and descriptions of (ideal) epistemic agents, and these theoretical entities are open to formal study and empirical evaluation.

Consider then

(*c*) Some distinctively internalist relations *should* be studied scientifically.

The primary example of an internalist relation currently undergoing thorough scientific study is the relation of *coherence*. Constraint programming software, such as PrSAT [9] is being used to develop and *test* hypotheses about the relationship between (anti-)correlation and confirmation,[10] and accumulated theoretical and *experimental* results from machine learning algorithms are now informing our understanding of the differences between synchronic coherence and diachronic coherence (Wheeler 2007). The results here are largely negative: we are learning

[9] Developed by Jason Alexander and Branden Fitelson, found at `http://fitelson.org/PrSAT/`.

[10] See in particular (Shogenji 1999), (Fitelson 2003), (Bovens and Hartmann 2003), (Olsson 2005), and (Meijs and Douven 2007).

more about what coherentist justification cannot be, than what it is. But they are results and insights that are hard to imagine coming about, particularly at so quick a pace, without the aid of computational tools, empirical facts about machine learning algorithms, and the experimental techniques that underpin the application and development of these tools.

Furthermore, one of us (Pereira) has been involved in developing a dynamic logic programming language (EVOLP) that facilitates self-evolution and updating (Alferes, *et. al.* 2002).[11] The EVOLP language extends Logic Programming (LP) by permitting rules to indicate assertive conclusions that themselves can take the form of program rules. Assertions of this kind, whenever they appear in a program, can be employed to generate an updated version of that program. This process may then be iterated on the basis of that new, updated program. When a program semantics affords several program models, *branching* evolution occurs, allowing several candidate evolution sequences to be constructed.

What is interesting about this technology from the point of view of this essay is that the EVOLP framework provides a powerful testbed for *dynamic* epistemic models, insofar as an epistemological hypothesis can be encoded within EVOLP, which may reveal the behavior of the (possibly nested) epistemic relation. This is a critical tool, since it is very often the case that the properties that make for a successful *static* model of a feature do not carry over to the development of a dynamic model. EVOLP, and tools like it, can be very useful for exploring dynamic epistemic relations, which is a topic that traditional, internalist epistemology has had difficulties advancing. We suggest that the main reason for this difficulty is methodological; it is very hard to reason *a priori* about dynamic procedures, and to test proposals with intuitive examples. Applying computational tools to dynamic epistemic processes is analogous to the application of computational methods to other dynamic processes, and there are empirical and analytical methods for verification and judging performance.

In the case of EVOLP the framework allows one to evaluate competing epistemic models by performance on benchmark test cases, or by fitness in artificial environments. In this last respect, dynamic programming tools like EVOLP can help to *test* links between strategies for internal, first-person epistemic assessments of beliefs against the performance of other agents within a test environment.

Another example involves the use of probability logic. One barrier to using probability logics is that there are considerable differences in

[11] See also (Dell'Acqua and Pereira 2007), and (Pereira and Lopes 2007).

interpretation between one system and the next, but another barrier is that it is difficult to perform derivations. One may approach the problem of calculating probabilities by soft computing techniques and treat probabilistic reasoning as a type of *constraint-based* problem, but one may also use graphical methods to reveal logical structure within an information set that is not readily captured with tools like PrSAT. PROGICNET[12] is investigating the application of probabilistic networks to probabilistic logic to capture this structure, which should provide an additional tool for both the scientific inquiry into philosophy, and philosophical inquiry into scientific methods.

So why would an internalist resist the general proposal to embrace methodological naturalism? Feldman (2001b) has framed the discussion of methodological naturalism in terms of a position he calls *Cooperative Naturalism*. According to cooperative naturalism, epistemology should examine evaluative questions, but that detailed empirical results from the cognitive sciences concerning how we actually think and reason are not essential for making progress in addressing evaluative questions.

But as we've seen, the theoretical *and* empirical results from AI that we have in mind are not about how we reason, but rather are about the structure of information.

Feldman remarks that there are no less than three views concerning possible sources of information for epistemological theorizing: pure *a priori* reflection; *a priori* reflection plus common-sense empirical knowledge; and scientific epistemology, which proclaims 'the value of or (or need for)' empirical results for epistemology' (2001b). Feldman writes

> This three way classification complicates the discussion of Cooperative Naturalism. If Cooperative Naturalism is the view that empirical information is important for resolving epistemological issues, then armchair epistemologists can accept it. However, if Cooperative Naturalism is the view that detailed information from the empirical sciences is important for epistemology, then armchair epistemologists are likely not to agree.

But any theorist should be drawn to a new method if it is shown to provide a benefit over his current methods, and we believe that the computational sciences have matured to a point where they now are making contributions to philosophy. We are seeing this already in the case of coherentist justification; we see that non-monotonic logics and logic programming offer conceptual insights into the structure of inductive arguments, and a platform to encode and test rules utilizing

[12] More information about Probability Logics and Probabilistic Networks can be found at the PROGICNET homepage, `http://www.kent.ac.uk/secl/philosophy/jw/2006/progicnet.htm`

default negation; and we see in the application of logic to these problems reasons to reject Harman's, Goldman's, and Chisholm's dated views on the role that logic can play in epistemology. Quine was right about methodological naturalism in our view, but he had the wrong science in mind: computational and formal epistemology have both theoretical and (fledgling) experimental branches that are beginning to make contributions to one another.

One point to keep in mind in this discussion is a distinction between whether there is some value in appealing to a new method, and whether it is *necessary* to appeal to that method to carry out an inquiry. There are few cases where decisions over methodology are settled by nomological necessity. Usually the issue is whether you are likely to get better results using a different method. It should be clear that our recommendation is a practical one: philosophical understanding is advanced primarily by advancing methods, and we think the continuing methodological advances in the computational sciences are ripe for exploitation by epistemologists, particularly internalists.

Before closing, let's consider another objection to methodological naturalism. Feldman, addressing Hilary Kornblith's endorsement of methodological naturalism, writes:

> Hilary Kornblith has suggested that philosophizing in the way epistemologists often do about knowledge is something like philosophizing about aluminum. The only serious questions about aluminum, he thinks, are scientific questions. (Kornblith 1999) It is difficult to see, however, exactly why we should think that knowledge is relevantly like aluminum. For one thing, what we seek in the case of aluminum is an understanding of its physical constitution. We want to know what it is made of, how it interacts with other materials and why, and what we can use it for. Our analysis of knowledge does not call for an account of its physical constitution. It's doubtful that there is any such thing.... But knowledge isn't a substance like aluminum or a process like cell division. So, analogies such as these don't provide reasons to seek naturalistic analyses of knowledge (Feldman 2001a).

But perhaps Kornblith's analogy isn't so far off the mark, for we are interested in the structure of knowledge, and how it interacts with our evidence about the world and our practical desires about what actions to take. We've argued that the problem of understanding justification as such is analogous to the problem of understanding information as such. It is precisely the structure of justification, evidence, and rational belief, their 'constitution', and the nature of their interactions that we wish to understand. Moreover, we have recommended that epistemologists

adopt the methods being developed to understand the latter for the study of the former.

Feldman writes that 'Some topics and questions are amenable to armchair methods and some are not. It would be foolish to extend Kornblith's line of thinking to logical concepts such as validity or conjunction, to modal concepts such as necessity, or, I believe, to moral concepts such as obligation. Some concepts have a richer conceptual structure than others.' This last point is precisely right. And the very richness of that structure may well extend beyond the capabilities of unaided imagination.

Keith Devlin has remarked that our 'information age' is very much like technological ages of the past—the 'bronze age', for example—in the sense that each epoch is marked by technological breakthrough that outpaces scientific understanding. Peoples of the bronze age were masters at making things with bronze, but nobody had the faintest idea what bronze was. Likewise, we are becoming masters of manipulating information, but we know no more about information than our long-ago ancestors knew about bronze. The study of human psychology may tell us something indirectly about what information is—by telling us what types of information processing we are capable of performing well, and what types we are not so good at performing. But traditional methodological naturalists are wrong to think that cognitive psychology will lead the way in epistemology. The internalists are right to insist that there are things like 'evidence' that bear relationships that should be studied in their own right. They are wrong, however, to resist the help of the sciences in conducting this inquiry.

5. Conclusion

The point under discussion is what methodology is likely to advance our understanding of epistemic support relations. The naturalists' complaint is that the methods of armchair philosophy cannot yield insight into complex relations and processes that are comparable to those gained when using the battery of scientific methods that are available. We agree. And internalists complain that the results of cognitive psychology are not of direct relevance to (internalist) epistemologies, which is lost in most discussions of naturalizing epistemology. Here too we agree. But where both internalists and externalists err is following Chisholm's out-dated view of logic and its application. Unlike Chisholm, we think that the relations and concepts imagined within epistemology are not unique but instead are structurally similar to various structures studied by the computational sciences. And it is

there that internalist epistemologists should turn for a methodological advantage.

References

Alferes, J., A. Brogi, J. Leite, L. M. Pereira. 2002. "Evolving Logic Programs", in S. Flesca, *et. al.* (eds.) *Proceedings of the 8th European Conference on Logics in Artificial Intelligence (JELIA 2002)*, Springer LNCS 2424, 50–61.

Blackburn, P., M. Rijke and Y. Venema. 2001. "Modal Logic", Cambridge: Cambridge University Press.

Bonjour, L. and E. Sosa. 2003. *Epistemic Justification*, Oxford: Blackwell Publishing.

Bovens, L. and S. Hartmann. 2003. *Bayesian Epistemology*, Oxford: Oxford University Press.

Carruthers, P., S. Stich and M. Siegal (eds.) 2002. *The Cognitive Basis of Science*, Cambridge: Cambridge University Press.

Chisholm, R. 1966. *Theory of Knowledge*, Englewood Cliffs, NJ: Prentice-Hall.

Conee, E. and R. Feldman. 2004. *Evidentialism*, Oxford: Oxford University Press.

Dell'Acqua, P. and L. M. Pereira. 2007. "Preferential Theory Revision", forthcoming in *Journal of Applied Logic*, 5(4).

Feldman, R. 1999. "Methodological naturalism in epistemology", in John Grecco and Ernest Sosa (ed.) *The Blackwell Guide to Epistemology*, Malden, MA: Blackwell, 170–186.

Feldman, R. 2001a. "We're all naturalists now", *Symposium paper delivered at the 2001 APA Pacific meetings*, Minneapolis, May 2001. Available at `http://www.ling.rochester.edu/ feldman/papers/naturalism.html`.

Feldman, R. 2001b. "Naturalized Epistemology", appearing in *Stanford Encyclopedia of Philosophy*, `http://plato.stanford.edu/entries/epistemology-naturalized/`, July 22, 2001 version.

Fitelson, B. 2003. "A Probabilistic Theory of Coherence", *Analysis* 63: 194–99.

Foley, R. 1987. *The Theory of Epistemic Rationality*, Cambridge, Mass: Harvard University Press.

Goldman, A. 1979. "What is justified belief?", in George Pappas (ed.) *Justification and Knowledge*, Dordreche: Reidel, pp. 1-23.

Goldman, A. 1986. *Epistemology and Cognition*, Cambridge: Harvard University Press.

Harman, G. 1986. *Change in View*, Cambridge, MA: MIT Press.

Harman, G. 2002. "Internal Critique: A Logic is not a Theory of Reasoning and a Theory of Reasoning is not a Logic", appearing in *Studies in Logic and Practical Reasoning, Vol. 1.* Gabbay, D. *et. al.* (eds.). London: Elsevier Science, pp. 171-86

Harman, G. and S. R. Kulkarni. 2006. "The Problem of Induction", *Philosophy and Phenomenological Research*, 72: 559-575.

Hempel, C. 1965. "Inductive Inconsistencies," in *Aspects of Scientific Explanation*, New York: Free Press.

Kim, J. 1988. "What is 'Naturalized Epistemology'?", *Philosophical Perspectives 2*, J. Tomberlin [ed.]. Atascadero, CA: Ridgeview Publishing, 381-406.

Kornblith, H. 1999. "In Defense of a Naturalized Epistemology", in *The Blackwell Guide to Epistemology*, John Greco and Ernest Sosa (eds.), Malden, Ma: Blackwell, pp. 158-169.

Kornblith, H.(ed.). 2001. *Epistemology: Internalism and Externalism*, Oxford: Blackwell.

Kyburg, H. E., Jr. and C. M. Teng. 1999. "Statistical Inference as Default Logic", *International Journal of Pattern Recognition and Artificial Intelligence* 13(2) : 267-283.

Kyburg, H. E., Jr. and C. M. Teng and G. Wheeler. 2007. "Conditionals and Consequences", *Journal of Applied Logic* 5(4), in press.

Mayo, D. 1996. *Error and the Growth of Knowledge*, Chicago: University of Chicago Press.

Douven, I. and W. Meijs. 2007. "Measuring Coherence", *Synthese*, in press.

Olsson, E. 2005. *Against Coherentism: Truth, Probability and Justification*, Oxford: Oxford University Press.

Pereira, L. M. and G. Lopes. 2007. "Prospective Logic Programming with ACORDA", *Workshop on Empirically Successful Computerized Reasoning*, 2006 Federated Logic Conference Seattle, USA.

Quine, W. V. 1969. *Ontological Relativity and Other Essays*, New York: Columbia University Press.

Shogenji, T. 1999. "Is coherence truth conducive?", *Analysis* 59:38–45.

Wheeler, G. and Damásio, C. 2004. "An Implementation of Statistical Default Logic", in J. Alferes and J. Leite (eds.) *Logics in Artificial Intelligence: JELIA 2004*, LNAI Series No. 3229, Berlin: Springer-Verlag, 121-133.

Wheeler, G. and L. M. Pereira. 2004. "Epistemology and Artificial Intelligence", *Journal of Applied Logic* 2(4): 469-493.

Wheeler, G. 2004. "A Resource Bounded Default Logic", in J. Delgrande and T. Schaub (eds.) *Proceedings of the 10th International Workshop of Non-Monotonic Reasoning (NMR 2004)*, Whistler, British Columbia, Canada, 416-422.

Wheeler, G. 2005. "On the Structure of Rational Acceptance", *Synthese*, 144(2): 287–304.

Wheeler, G. "Applied Logic without Psychologism", forthcoming in *Studia Logica*.

Wheeler, G. 2007. "Synchronic and Diachronic Coherence", unpublished manuscript.