

Epistemology in an Age of Cognitive Science

ROBERT N. MCCAULEY

Department of Philosophy Emory University, Atlanta, Georgia 30322, U.S.A.

ABSTRACT *Like the logical empiricists many contemporary philosophers wish to bring the determinateness of scientific judgment to epistemology. Recent efforts to naturalise epistemology (such as those of the Churchlands) seem to jeopardise the position of epistemology as a normative discipline. Putnam argues that attempts to naturalise epistemology are self-refuting.*

My goal is not to defeat the project for the naturalisation of epistemology, but rather to help clarify what it does and does not amount to. I maintain that attempts to completely eliminate the normative will be either forever incomplete or inimical to the progress of science. However, because it is the first horn of this dilemma which will prevail, these considerations do not undermine the importance of future epistemology carefully attending to the results of the relevant sciences.

1 Introduction

Many modern philosophers have aspired to bring the precision of the sciences to epistemology. The logical empiricists pursued this goal by offering an account of scientific rationality that turned on precise logical reconstructions of established scientific theories. Allegedly, in these reconstructions all non-logical terms either had some logically perspicuous relation to or were terms whose meanings were based exclusively upon observations that were supposedly neutral with respect to theoretical orientation. More recently, though, many philosophers of science have been increasingly prone to deny any strong distinction between theoretical and observational notions. Consequently, when contemporary philosophers, such as the Churchlands, pursue the goal of bringing the determinateness of scientific judgment to epistemology, they have employed a different strategy.

Ultimately, inspired by Quine (1969), but unconfined by his behaviorism, the Churchlands simply construe epistemology as that part of scientific research concerned with human cognition. In short, the way to bring the determinateness of scientific judgment to epistemological claims is simply to pursue the relevant sciences. Enthusiasm for this project has increased in the past decade in part because of the startling progress in those sciences. Progress in psychology, linguistics, cognitive anthropology, artificial intelligence (especially of the connectionist variety), and the various neurosciences has rendered increasing expanses of human cognitive activity accessible to systematic empirical investigation.

The Churchlands are especially famous for their focus on neuroscience. Paul

Churchland asserts, for example, that "making acceptable contact with neurophysiological theory is a long-term constraint on any epistemology; a scheme of representation and computation that cannot be implemented in the machinery of the human brain cannot be an adequate account of human cognitive activities." Philosophical speculation about the capacity, processing, development, and organisation of our cognitive system uninformed by the achievements in these sciences risks thoroughgoing irrelevance to what systematic knowledge we have about these matters (see also P. S. Churchland, 1980).

For example, those epistemologists who rely primarily on the traditional framework of folk psychology assume that the conceptual habits that evolved prior to systematic psychological theorising are largely adequate. Against this view, the Churchlands raise two considerations that are particularly sobering. First, an induction on the relative empirical adequacy of such folk theories is hardly encouraging. Wherever folk theories have preceded, systematic scientific research has largely displaced them—the persistence of some of these theories among the folk notwithstanding (McCloskey, 1983). Patricia Churchland remarks: "it would be astonishing if folk psychology, alone among folk theories, was essentially correct" (1986, p. 395). The Churchlands and others (see especially Stich, 1983) have also emphasised how many recent findings in the cognitive sciences seem to undermine traditional presumptions about human cognition enshrined in folk psychology.

Paul Churchland has been especially forthcoming about a further consequence of naturalising epistemology along these lines. He points out that "the claim that the enterprise of epistemology should be conducted along the lines of any other natural science renders problematic the status of what we would call *normative* epistemology" (1979, p. 124). Traditionally, the task of epistemology has involved formulating general principles for the *assessment* of beliefs. The central vocabulary of epistemology has been normative throughout. Epistemologists theorise about the most *honored* of our beliefs, namely, those which merit the title 'knowledge'.

The prospect of jeopardising normative epistemology, though, does not seem to worry Churchland in the least, for, as indicated already, he holds that "the common-sense P-theory [person theory], whose categories are taken as a given by almost all approaches to normative epistemology, is a theory whose basic integrity is very much open to doubt... it is still an open question whether the fundamental parameters of rationality are to be found at the categorial level it comprehends" (1979, p. 123; see also P. M. Churchland, 1981). Although Patricia Churchland eschews making direct predictions about these normative matters, she discusses the same questions in the section of her book on *eliminative materialism*. On her view as well the elimination of the categories of folk psychology would, presumably, certify the elimination of the normative features of epistemology that those categories undergird. In addition, she offers an analogical argument in this section about the complete failure of similar sorts of last-ditch normative arguments to preserve either geocentrism or creationism in science (1986, pp. 398–99). By contrast, both of the Churchlands have discussed alternative, sub-personal approaches to knowledge and cognition that recent findings in cognitive neurobiology and connectionist artificial

intelligence suggest—approaches which they argue do not presume the P-theory that informs the class of (normative) ‘sentential epistemologies’ that they reject.

The progress of the cognitive sciences and their promise for contributing to the determinateness of epistemological judgment notwithstanding, one critic of naturalised epistemologies remains unimpressed with prognostications about the elimination of normative epistemology. Hilary Putnam has argued that “the elimination of the normative is attempted mental suicide” (1983, p. 246). Although Putnam reaches this verdict in the course of criticising Quine’s program, he clearly intends to include all attempts to naturalise epistemology within its purview.

Putnam’s argument, in short, is that if we explain all of our norms away, then we have no standards remaining by which to assess competing explanatory claims—including those concerning the explanation of our norms. Without normative standards there is nothing that we can be right or wrong about, hence our arguments and assertions are all for naught. It is that prospect that motivates Putnam’s talk of “mental suicide”. Putnam holds that all attempts to naturalise epistemology question “one of our fundamental self-conceptions . . . [viz.] that we are *thinkers*” (1983, p. 246). According to Putnam, such positions are simply self-refuting.

Although Putnam is wary of aprioristic arguments about the necessary shape of epistemic concepts, the alternative picture he offers portrays epistemology as inescapably normative, nonetheless. *Explanation*, according to Putnam, is itself an epistemic notion (1983, pp. 290–98). All explanations presuppose values. Community norms always play a pivotal role in what counts as a satisfactory explanation. If explanation contains an ineluctably normative dimension, then even more clearly do such traditional epistemic notions as ‘knowledge’, ‘justification’, and ‘rationality’. Any of the naturalisers’ proposed explanations of these central epistemic notions just as surely presupposes further epistemic norms.

2 A Dilemma for Attempts at Exhaustively Naturalising Epistemology

Naturalists’ attempts to eliminate the normative outright face a dilemma. They will either prove incomplete or, if complete, then inimical to the progress of science. Nonetheless, I propose (1) that neither this conclusion nor Putnam’s should discourage advocates of naturalised epistemology in the least from pursuing their projects and (2) that even Putnam’s discussion offers important clues about why that is so. In effect, then, my goal is not to defeat the project for the naturalisation of epistemology, but rather to help clarify what it does and does not amount to.

First, if explanation *is* a fundamentally epistemic notion, as Putnam maintains, then explanatory programs for the naturalisation of epistemology themselves embody norms. Both of the episodes in the history of science (namely, the demise of geocentrism and that of creationism) that Patricia Churchland cites contributed significantly to the evolution of new standards for acceptable scientific work—for example, that, *ceteris paribus*, we should prefer, over their competitors that do not, theories that have an array of consequences that are specific enough to submit to extensive empirical tests (Kitcher, 1982). More particularly, these episodes also had an impact on what would subsequently count as an acceptable explanation, suggest-

ing, for example, that good explanations in science not appeal to the intervention of the supernatural. But, clearly, this sort of *evolution* of new scientific standards is equivalent neither to the *abolition* nor to the *elimination* of all such standards.

Churchland's analogical argument concerning the failure of attempts to defend geocentrism and creationism on the basis of their centrality to prevailing normative positions certainly *does* suggest that arguments of this sort are incapable of preserving what have proved (in the course of research) to be deficient theories. However, it does not support the contention that normative matters in general can be eliminated once and for all. What these cases show is that *as a result of theoretical progress in science our normative commitments sometimes change* (both in epistemology and in the moral sphere).

Admittedly, these cases show that what were once *obviously* matters of great moral significance may no longer be so conspicuously so. However again, *that* does not demonstrate that matters of normative concern in general are candidates for elimination. Furthermore, so long as contemporary creationists, in their zealous opposition to 'secular humanism', continue to advocate special creation, the conflict between their views and the theory of evolution remains one of profound moral significance for them at least. The crucial point is that what counts as a normative matter pertaining to science (either moral *or* epistemic) is not always decided exclusively on the basis of considerations *internal* to science. (The contrary position covertly resurrects the logical empiricists' assumption that the standards for a satisfactory explanation can be formally specified once and for all.)

It is the respective positions of the competitors in a dispute and those matters about which they take issue with one another that largely shape the perceived contours of normative issues. What counts as normative within science depends, at least in part, on the relative prominence of the dissenting views (both within the scientific community and within the culture at large) that are critical of the dominant methods and theoretical developments of some science. It also depends upon the kind of challenge that dissenters present to scientific standards and achievements. (Hence, it is much easier for us to understand how the origin of humanity still constitutes an issue of moral significance than it is for us to understand how the position of the earth does.) The polemical context carries the major share of the responsibility for highlighting the normative issues.

These comments are of a piece with Paul Roth's argument to the effect that "... we cannot clearly mark off what is normative from what is not" (1987, p. 32). Roth (1987, p. 43) points out that, no less than the distinction between the conceptual and the empirical, the distinction between the descriptive (or explanatory) and the normative is overwhelmingly mitigated by the sort of holism Quine proffers in semantics—a position, incidentally, to which all the parties in this discussion (Roth, the Churchlands, Putnam, and myself) generally subscribe. Holistic considerations constrain claims about justification no less than other claims within the web of belief. Roth argues that, depending upon the conceptual and theoretical frameworks at hand, different issues will seem more plainly normative than others. One consequence of this view, then, is that what counts as a normative issue, even in epistemology, is, indeed, largely a function of the polemical context.

This is not to argue that any particular assumption is intrinsically normative come-what-may, but it does imply that *all positions involve presuppositions (e.g. about methods) that function normatively* (in the same way that all positions involve conceptual presuppositions that function *within those positions* like analytic truths). (Paul Churchland emphasises this last point in his 1979, chapter 3). *Whether* normative dimensions of assumptions are manifest usually depends upon the ferocity of relevant disputes. Typically, *which* normative dimensions are manifest depends upon the problems and the purposes of the disputants.

It follows that every new explanation of norms will itself presuppose further norms which in turn require explanation. The proposed encapsulation of the normative within the framework of our explanations, which the Churchlands foresee and which Putnam deems suicidal, will surely remain forever incomplete. This is the first horn of the dilemma.

Just as each normative consideration submits to systematic explanation, each of these explanations embodies further normative assumptions. On this picture of inquiry, then, the relation of the normative and the explanatory is one of a perpetual and productive interaction where the borders between the two domains *could* shift (often only imperceptibly, perhaps) with each and every advance. Crucially, though, the persistence of both normative and explanatory dimensions in science is a prerequisite for its progress. Hence, the elimination of the normative is antagonistic to the advance of science and *that* is the second horn to which I now turn.

The pervasiveness of normative assumptions within cognitive endeavours (manifested in disputation especially) applies with undiminished force to theoretical explanation in science. If the confluence of opinion about our explanatory successes ever completely overshadowed the normative assumptions that inform scientific work, it would signal the absence of disputes. The absence of disputes would indicate the cessation of theoretical competition, and the cessation of theoretical competition would put an end to any substantial scientific change. (It is indifferent to the success of this argument whether or not we regard that change as progressive.) *When science ceases to change, science ceases.* Only on the most naive forms of realism could science ever grind to such a halt—presumably, because it had finally gotten all of its descriptions right! Only on such views is science ever in danger of wholly deteriorating into metaphysics.

As the Churchlands see it, the disappearance of the normative is a consequence of the elimination of folk psychology in favor of some future, successful, neuroscientific account of cognitive functioning. This accomplishment, on their view, will undercut the categorial framework upon which normative judgments depend. For the Churchlands our self-conception as thinkers is theoretical through and through (and, hence, subject to the sorts of criticisms that we can bring against any theory in science). Furthermore, they point out that sometimes in the history of science a theory proves so incompatible with other reigning theories and with a superior competitor in particular that all traces of it disappear from the scientific landscape once that competitor becomes entrenched. On the basis (1) of the recent successes within the neurosciences, (2) of the perennial and persistent flaws of folk psycho-

logy, and (3) of the (allegedly¹) stark incommensurability of the two, the Churchlands argue that we have many good reasons to anticipate just this sort of elimination of folk psychology. Since the folk psychology of propositional attitudes provides the categorial superstructure of normative epistemology, they have also suggested that such traditional approaches to epistemology will share the same fate.

I have argued elsewhere that the Churchlands' specific eliminativist argument fails (McCauley, 1986a, especially pp. 188–98). The problem, in short, is that they employ a model of intertheoretic relations in science that is too coarse-grained. The elimination of theories in science of the sort that the Churchlands foresee reliably occurs *within* a particular analytical level in science, *not between* analytical levels. The abrupt elimination of theories in science occurs when some newly ascendant theory (the oxygen theory of combustion) is overwhelmingly incompatible with its immediate predecessor (the phlogiston theory of combustion) *within the same science*. When theoretical succession in a science is less tumultuous, the elimination of the remaining traces of (much) earlier predecessors is less obvious, perhaps, but just as inevitable. Hence, it was only with the appearance of Newtonian (and not Galilean) mechanics that the vestiges of the late medieval, Neo-Aristotelian mechanical theories were ultimately purged.

The crucial philosophical point is that the elimination of theories in science is an *intra*level not an *inter*level phenomenon. Perhaps, the most obvious evidence for this general philosophical claim is sociological. Once a science has enjoyed momentum sufficient for its practitioners to identify themselves as such, the abrupt elimination of an entrenched theory within such a science *always* comes as a result of a revolutionary challenge that a new theoretical competitor presents *at the same analytical level*, i.e. within the same science. It was not theoretical developments in chemistry or physics that eliminated phrenology from neuroscience, but rather alternative *neuroscientific* positions that proved more responsible empirically (Bechtel, 1982 and 1988).

Returning to the issue at hand, then, psychology (of any cognitivist sort) and neuroscience seem to operate at levels of analysis that are sufficiently removed from one another that the straightforward elimination of (folk) psychology by developments in neuroscience is extremely unlikely. This is, in part, a function of just the sort of radical discontinuities between folk psychology and neuroscience that the Churchlands have documented. Precisely, because their differences are substantial on so many different points, it seems improbable that the neurosciences and high level cognitive theorising (folk psychological or not) would ever evolve in such a way that they would both be incorporated within a single analytical level. So, since the elimination of theories in science is an *intra*level process, theoretical advances in

¹ It is worth noting that when the Churchlands have discussed neuroscientific theories of cognitive functioning in detail, some of the time their accounts seem to *support* various sub-systems of our intuitions that are, presumably, informed at least in part by folk psychology. Most of the cases in question concern the qualia (and relationships between them) that we associate with perceptual experiences—such as the relative proximity of various perceived colors. See, for example, P. M. Churchland, 1986, especially pp. 299–305.

neuroscience are not likely to eliminate psychological theories occurring at a recognisably higher level.

Some clarifying comments are in order. First, the model of intertheoretic relations implicit in this alternative account does not preclude the elimination of folk psychology, it simply predicts that if folk psychology is eliminated, it will not be from the neuroscientific level of analysis that we should expect its likely successor. In all probability the neurosciences will share some responsibility in the crime, but it is much more likely to be some descendant of contemporary cognitive psychology that will do the actual dirty work. Consequently, this is *not* to deny Patricia Churchland's prediction that future psychology and neuroscience are likely to substantially co-evolve, but, rather, to emphasise that the language of co-evolution is a welcome corrective to the Churchlands' talk of eliminative materialism, especially since they still employ both (see P. S. Churchland, 1986, chapter 9, especially).

This alternative account of intertheoretic relations in science, then, leaves open the possibility of eliminating folk psychology, but *only as an explanatory theory within psychological science* (McCauley, 1988). Nothing follows about the ability of something perhaps quite like our current folk psychological framework to continue to serve (even) the explanatory needs of social scientists—for example, along lines Dennett (1981, p. 50) suggests for employing what he calls "pure intentional system theory" in characterising an ideally rational consumer. Nor must the elimination of folk psychology as an explanatory theory in psychology necessarily preclude its providing the conceptual framework for our most overtly normative projects.

Finally, not only is an elimination of folk psychology along the lines the Churchlands predict unlikely, it is undesirable. Maintaining the integrity of separate analytical levels in science is important to scientific progress. If cognitive psychology were to collapse into cognitive neurobiology or even to be driven by it in the way that the Churchlands sometimes seem to expect (e.g. P. S. Churchland, 1986, p. 294), we would sacrifice important sources of theoretical and problem solving inspiration (currently supplied by psychological research) as well as numerous funds of evidential support at the altar of a unified, but impoverished, science. Too often, the co-evolution the Churchlands seem to anticipate is one where all of the selection pressures seem to be exerted exclusively from the bottom up. Retaining some distance (and independence) between the sciences operating at different analytical levels has its advantages too.

3 Naturalised Epistemology Nonetheless!

Neither these arguments nor Putnam's should deter philosophers' interests in the bearing of the cognitive and the neuro-sciences on epistemology. Nor should those who champion either folk psychology, or ordinary language, or the preeminence of intuition find much comfort here. None of the considerations raised in the previous section undermine Quine's fundamental insight that once we are bereft of appeals to First Philosophy, it is the standards of those pursuits that contribute most obviously and most successfully to human flourishing generally and to our knowledge about

the world in particular that offer the clearest and most defensible standards of rational justification. I assume that it is non-controversial that science deserves a prominent, if not the most prominent, position among the candidate activities (see Roth, 1987, p. 34). Standards that are *simultaneously external to, superior to, and independent of* those of science simply do not exist. The standards to which scientific communities aspire and the results at which they arrive jointly constitute compelling grounds for regarding scientific endeavor as the paradigm case of rational inquiry.

This preference for science concerns the determinateness, not the quantity, of its successors. It looks to the (relative) determinateness of both its findings and its methods. The results and the methods of our knowledge-seeking activities fall along a continuum on this count. The problem, though, is that the farther away from physical science we move on this continuum, the less determinate truths and methods are and the less clear we are both about what to make of them and about what to do with them. Although scientific justification is not the only kind, science has the distinct advantage of yielding relatively determinate results in accord with relatively determinate methods.

Since science can never definitively justify itself, it does not and it cannot monopolise the truth. Nonetheless, even Putnam's position inevitably must turn back to science. Truth according to Putnam "...is an idealization of rational acceptability" (1983, p. 200). It is difficult to see how to cash such formulations in without devoting substantial attention to *the standards of science generally and the findings of the cognitive sciences in particular*. Although not all truths are truths of science, the truths of science stand much nearer to the core of our concept of truth than do our less determinate, veracious claims (McCauley, 1986b). No doubt, this assertion looks a bit question begging, since whatever force it possesses seems to be psychological only. But once all versions of foundationalism have been found wanting, the accounts of rationality whose begging of questions is least offensive are precisely those which attend most conscientiously to *our best accounts of the relevant facts*, and the relevant facts here are largely psychological. Progress within the cognitive sciences concerning the acquisition, processing, and utilisation of information seems, *prima facie* at least, to add to this body of relevant facts². (Certainly, this does not imply that anything is finally settled. We can always dig deeper. For any inquiry, another is always possible which examines the underlying assumptions of the first. Reliably, though, these more basic inquiries take on an air of circularity themselves, and they sacrifice more and more precision, i.e. determinateness, in the bargain. The returns diminish exceedingly fast).

The normative *is* in some sense ineliminable. Contrary to Putnam, though, the persistence of the normative only precludes the possibility of ever *exhaustively* naturalising epistemology. Putnam does not show either that our current epistemic norms nor that the categorial framework in which they are posed must persist, but

² So, for example, on the notion of radial categories that informs the comments on truth that immediately precede and that appear in McCauley, 1986; see Lakoff, 1987.

only that normative discourse of some sort will. The persistence of normative discussions does not establish any principled autonomy for epistemology. Just as surely as all debates about the facts presuppose norms, all normative discussions make assumptions about the facts.

The crucial point here is that this persistence of the normative has no bearing whatsoever on the urgent need of contemporary epistemologists to carefully attend, henceforth, to developments in the cognitive sciences and to take as their primary occupation the task of making sense of those findings. (In fact, most disputes about science's norms take care of themselves in the course of scientific practices.) From the fact that epistemology will never be completely co-extensive with science it follows neither that it cannot nor that it should not become *much more* responsive on these counts.

My differences with the Churchlands, then, concern (1) the inevitability of normative matters resurfacing even after attempts to explain them away and (2) the range of future scientific work that is likely to contribute to our understanding of cognition. My comments about intertheoretic relations amount to suggesting that the Churchlands' (commendable) predilection for determinateness in epistemology must be tempered by other values. These differences, however, do not concern our common convictions concerning the relative desirability of the day to day developments within the cognitive sciences informing all subsequent work in epistemology.

It is not that the final theories are already in place; rather, it is that many of the characteristic topics of epistemology and philosophical psychology can now be approached by the methods of empirical science. The point is not that the work has all been done, but rather that science is now eminently capable of directly contributing.

Acknowledgement

I am indebted to Mark Johnson and Paul Roth for their helpful comments on an earlier version of this paper and to Bill Bechtel both for discussions of the issues with which this paper is concerned and for encouragement.

References

- BECHTEL, W. (1982) Two common errors in explaining biological and psychological phenomena, *Philosophy of Science*, 49, pp. 549-74.
- BECHTEL, W. (1988) *Philosophy of Science: An Overview for Cognitive Science* (Hillsdale, NJ, Lawrence Erlbaum Associates).
- CHURCHLAND, P.M. (1979) *Scientific Realism and the Plasticity of Mind* (Cambridge University Press).
- CHURCHLAND, P.M. (1981) Eliminative materialism and the propositional attitudes, *Journal of Philosophy*, 78, pp. 67-90.
- CHURCHLAND, P.M. (1986) Some reductive strategies in cognitive neurobiology, *Mind*, 95, pp. 279-309.
- CHURCHLAND, P.M. (1988) On the nature of theories: A neurocomputational perspective, in: C.W. SAVAGE (Ed.) *Minnesota Studies in the Philosophy of Science, Volume XIV* (in press) (Minneapolis, University of Minnesota Press).

- CHURCHLAND, P.S. (1980) A perspective on mind-brain research, *Journal of Philosophy*, 77, pp. 185-207.
- CHURCHLAND, P.S. (1986) *Neurophilosophy* (Cambridge, MA, MIT Press).
- DENNETT, D.C. (1981) Three kinds of intentional psychology, in: R. HEALEY (Ed.) *Reduction, Time and Reality* (Cambridge University Press).
- KITCHER, P.S. (1982) *Abusing Science* (Cambridge, MA, MIT Press).
- LAKOFF, G. (1987) *Women, Fire, and Dangerous Things* (University of Chicago Press).
- MCCAULEY, R. (1986a) Intertheoretic relations and the future of psychology, *Philosophy of Science*, 53, pp. 179-199.
- MCCAULEY, R. (1986b) Truth, epistemic ideals, and the psychology of categorization, in: A. FINE & P. MACHAMER (Eds) *PSA—1986 Volume 1*, pp. 198-207 (East Lansing, MI, Philosophy of Science Association).
- MCCAULEY, R. (1988) Psychology in mid-stream, *Behaviorism*, 16 (in press).
- MCCLOSKEY, M. (1983) Intuitive physics, *Scientific American*, 248, pp. 122-130.
- PUTNAM, H. (1983) *Realism and Reason* (New York, Cambridge University Press).
- ROTH, P. (1987) *Meaning and Method in the Social Sciences* (Ithaca, Cornell University Press).
- QUINE, W. (1969) Epistemology naturalized, *Ontological Relativity and Other Essays*, pp. 69-90 (New York, Columbia University Press).
- STICH, S. (1983) *From Folk Psychology to Cognitive Science: the case against belief* (Cambridge, MA, MIT Press).