DAVID PAPINEAU

DOES THE SOCIOLOGY OF SCIENCE DISCREDIT SCIENCE?

1. THE NEW SOCIOLOGY OF SCIENCE

Why should we believe the findings of science? The traditional answer is that there are good *arguments* for those findings. Scientists start from the evidence of their senses, and then proceed by valid inferences to their conclusions. The scientific community is the rational community *par excellence*, and the findings of science are simply the conclusions forced on scientists by the dictates of reason.

This traditional view of science's epistemological authority is complemented by traditional sociology of science. Traditional sociology of science analyses scientific institutions. It looks at the norms governing such things as scientific publications, scientific education, priority disputes, fraud accusations, etc. But within traditional sociology of science there is no sociological analysis of scientific reasoning and its products themselves. The real engine of science — the processes that fix the contents of scientific theories, rather than their mere existence and subject matter — is simply taken for granted. The engine of science is simply the power of pure reason, and as such is not a fit subject for sociological analysis.

In the last decade or so a "new sociology of science" has emerged, distinguished by its readiness to investigate the processes which actually influence the *content* of scientific theories. And a central lesson of the new sociology of science seems to be that the actual behaviour of science bears little relation to the traditional rational picture. The factors influencing real scientists in their theoretical decisions do not look like good arguments at all.

We can usefully distinguish a macro-sociological and a micro-sociological school within the new sociology of science. The former school, whose stronghold is the Science Studies Unit at the University of Edinburgh, favours explanations of scientific theories in terms of their ideological significance, or their consonance with other features of the general social context. Thus, for example, the Darwinian theory of natural selection has been associated with certain social doctrines

Robert Nola (ed.), Relativism and Realism in Science, 37—57.

prevalent in Victorian England. Pasteur's success in 'disproving' the thesis of spontaneous generation has been attributed to the association of this thesis with such atheistic views as materialism and evolution, in the context of the conservative and religiously orthodox Second Empire. Or again, to take a rather different kind of case, David Bloor has suggested that the argumentative strategies adopted by mathematicians, and in particular the nature of their reactions to counter-examples, reflect the social structure of their mathematical community, and in particular the nature of its internal and external social divisions. [Cf., respectively, Barnes and Shapin (1979), Farley and Geison (1974), Bloor (1978).]

The micro-sociological branch of the new sociology of science, on the other hand, has affinities with the general ethnomethodological movement in sociology. As such it is by and large uninterested in large-scale social influences on scientific developments. Instead it studies the micro-interactions which lie behind the achievement of scientific consensus. It emphasizes the negotiations, the trial and error procedures, and the tacit understandings which precede any agreement on scientific findings. A significant factor in such negotiations is often the concern of individual scientists, or small groups, to gain credit for "establishing facts". Such establishing of facts is viewed as a matter of persuasion and argument. What kinds of experiments are conducted, what kind of controls are required, how the results are presented, where the papers are published, etc., are all explained in terms of the need of scientists to win acceptance for their "facts" from their colleagues and opponents. Often there will be significant periods of vagueness and negotiation during which the different parties will compete and manoeuvre, and when the fate of a putative fact will hang in the balance. Success in winning acceptance for facts will depend in part on the contextually variable norms of the particular scientific community, and in part on the *resources* the scientist can command: most obviously, resources in terms of money and equipment, but also, less tangibly yet more fundamentally, resources in terms of "scientific credibility", which gets built up, like capital, during a career, and which can then be re-invested in pursuit of further "facts". [Perhaps the best-known example of such micro-sociology of science is Latour and Woolgar's *Laboratory Life* (1979). See also H. Collins (1985), M. Lynch (1985).]

My question is this: does the new sociology of science discredit science? Should those of us who have hitherto been inclined to believe

the findings of science be worried by these new descriptions of scientific activity?

My answer will be that it depends. It depends on one's attitude to epistemology. Those with certain orthodox, "Cartesian" assumptions about the nature of epistemology are indeed right to be worried by the new sociology. But I shall show that there is an alternative, *naturalized* approach to epistemology which makes it possible to accommodate the new sociology without rejecting science.

One point before proceeding. As these last remarks will have suggested, I shall not be directly concerned with the factual standing of the new sociology. Perhaps some of its claims are false. No doubt many of its claims (especially in the macro-sociological mode) are speculations that demand further empirical research. But since my concern is really with the hypothetical question, "What *would* follow *if* the new sociology were true?", I am happy to take it as read that there is real factual substance in the new sociology.

2. TWO KINDS OF EPISTEMOLOGY

Epistemology is a practical business. It tells us what to do in order to ensure we have the right beliefs. More specifically, an epistemological *theory* specifies a certain kind of preferred *technique* for acquiring beliefs, and then makes the normative *recommendation* that concerned believers should acquire all their beliefs from that technique. (When an actual belief derives from the preferred technique, we say it is *justified.*)

The *Cartesian* theory of epistemology recommends that we should get our beliefs from *good arguments.* We should assent only to those beliefs that have been generated by logically valid steps from secure premises. An actual belief is justified just in case in issues from such an argument.

This Cartesian conception of epistemology lies behind most of the Western philosophical tradition of the last three hundred years. Even today nearly everybody takes it for granted. It was implicit in the opening remarks of this paper about the reasons for believing the findings of science. But there is an alternative way of thinking about epistemological matters. There is another epistemological theory, the naturalized theory, which gives an equally powerful account of how to get the right beliefs.

I shall explain this naturalized alternative in a moment. But first it is

worth noting how the Cartesian theory goes hand in hand with the conception of the conscious mind as transparent to itself. We are to assent only to those beliefs that follow validly from secure premises. But how are we to select those secure premises? And how, for that matter, can we ensure that the steps leading from them to the conclusion are indeed valid? By further arguments? In order to avoid the regresses threatening here we need some beliefs whose truth is ensured without argument, and some logical steps whose validity is immediately apparent. From Descartes onwards these have been provided by the conscious mind's supposed privileged awareness of its own ideas and of the relations between those ideas.

Now for the naturalized theory. On this theory, the right technique for acquiring beliefs is simply to be a *reliable belief-former*, that is, to have belief-forming processes that generally produce true beliefs. Concerned believers should try to ensure that all their beliefs come from belief-forming processes that are reliable in this sense. An actual belief is justified just in case it issues from a reliable process. [Cf. Goldman (1976), (1979).]

The difference between the naturalized and the Cartesian theory is best brought out by considering how the naturalized theory deals with perception and memory. According to the naturalized theory, perceptual beliefs and memories are justified if they come from reliable processes, and are unjustified if they come from unreliable proceses. So a good example of an unjustified belief would be little Johnny's naive belief that the pitch of the train's whistle changes as the train passed by: for the process behind this belief is unreliable, in that it characteristically generates false beliefs. Similarly, people who naively succumb to *déjà vu* experiences, and take themselves genuinely to be seeing things for a second time, will have unjustified beliefs. For, again, the processes behind their beliefs are unreliable. Of course not all the perceptual and memory processes embodied by human beings are unreliable in this way. In most cases our perceptual and memory beliefs come from reliable process and then, the naturalized theory wants to say, those beliefs are justified.

So the naturalized notion of justification doesn't necessarily involve conscious argument. On the naturalized approach, the difference between justified and unjustified perceptions and memories *can't* be due to the former having superior argumentative backing: for, after all, the distinguishing characteristic of perception and memory is that from

the perspective of consciousness they are *non*-inferential. Our perceptual and memory processes deliver beliefs into consciousness, but the processes leading up to those beliefs lie outside consciousness.

Not that the naturalized approach rejects argument altogether. For argument is itself a belief-forming process. Given existing beliefs as premises, argument is a process that moves us to new beliefs as conclusions. So, as part of the general recommendation that our belief-forming processes should be reliable, the naturalized approach recommends that our argumentative habits should be reliable too. (Though of course here the appropriate notion of reliability is conditional, rather than categorical: the conclusions should be true *if* the premises are.)

So the naturalized approach certainly allows that a belief's justification can be due to its coming from a good argument: if the premises of an argument are themselves reliably generated then the conclusion of that argument is justified just in case the argumentative habit in question is conditionally reliable for truth. But note that the naturalized attitude to argument is here quite different from the Cartesian attitude. It is not because argument is a conscious phenomenon, transparent to the mind, that it is good for generating beliefs. The point is simply that certain kinds of argumentative habits reliably generate true conclusions given true premises. Consciousness doesn't come into it. From the naturalized perspective, if a completely non-conscious being embodied reliable inferential processes, then those processes would give that being justified beliefs (or "beliefs", if you prefer).

I take it to be the overriding virtue of the naturalized approach that it doesn't need consciousness's supposed power of incorrigible self-awareness. It allows us to replace the picture of autonomous individual subjects, each locked into the mental arena of their own epistemological responsibilities, with a picture in which human thinkers are simply normal beings interacting with the rest of the natural world.

A corollary is that the naturalized approach allows us to get away from the idea that epistemology is necessarily about techniques which a single individual can use to acquire beliefs. So far in this section I have not questioned this assumption, and have concentrated on the attitude of naturalized epistemology to processes embodied in single individuals, like perception, memory and argument. But there is no need to go on thinking of epistemology in this way once we reject the idea that all epistemological techniques must operate within conscious minds. If we switch to the naturalized notion that justified beliefs require reliable

processes, processes whose operation may well lie outside consciousness, we can perfectly well think of beliefs as justified because they come from reliable *social* processes, processes which involve whole sets of people and the interactions between them, and not just from processes operating in individuals. This will be important in Sections **5** and **6** below.

3. MORE ABOUT NATURALIZED EPISTEMOLOGY

A full investigation of the strengths and weaknesses of the naturalized approach is out of the question here. But given the relative unfamiliarity of the naturalized approach, it will be worth saying a bit more about it, and in particular about its relevance to the normative question of what we *ought* to believe.

Many philosophers find it difficult to see how the naturalized theory can possibly be an answer to this practical question. If we try to follow the naturalized theory's normative recommendation that we ought to acquire all our beliefs from reliable processes, we seem to face an insuperable dilemma. The demand for naturalized justification can be read either "externally" or "internally". If we think of it entirely "externally", with justification depending on the reliability of processes some of which lie outside the domain of consciousness, then how can such "justification" be of any significance to a concerned believer? Such a believer won't necessarily know whether or not the processes involved are reliable. So how can the fact of their reliability or unreliability possibly influence the concerned believer in deciding what to believe?

Alternatively, we can read the naturalized approach "internally", as recommending that for any given belief we need to have a further belief to the effect that the original belief was reliably produced. But then we face a particularly vicious version of the familiar regress. (What about the further belief about reliability? Do we need to believe that *that* was reliably produced? Etc.)

But this dilemma begs the question. It seems insuperable only because the grip of the Cartesian approach is so strong. It takes it for granted that what the concerned believer needs is arguments. Thus: if the relevant facts are unknown, they cannot enter into the concerned believer's *arguments*; if we do try to get them into the concerned believer's *arguments*, there is an obvious regress.

The question, however, is precisely whether concerned believers

should seek out arguments, or whether they should simply aim to be reliable belief-formers. Naturalized epistemology makes the latter recommendation. But this should not be understood as a recommendation as to how better to conduct arguments. That is the very point at issue.

It might still be quite unclear how the naturalized theory can have any normative import. But once we free ourselves from the idea that argument is the only possible epistemological technique, nothing could be simpler. You want to be a reliable belief-former? Well then, you'd better do what has to be done to bring this about. You'd better investigate what belief-forming processes you already embody, and you'd better consider what alternative such habits you might adopt. And in particular you'd better investigate which of those habits, actual and possible, are reliable for generating truths. And having done all that, you should take steps to rid yourself of any bad, unreliable habits you already have, and take steps to instil any good, reliable ones that are open to you.

In effect, the idea is to think of yourself as a system for generating true beliefs. You want to be as reliable a system as possible. So you consider ways of redesigning the system, and you implement those that promise an improvement. [Goldman makes some suggestions along these lines in his (1978), (1980), and (1985). But it seems to me that he fails fully to free himself from Cartesian presuppositions: see in particular p. 30 of his (1980), and p. 40 of his (1985).]

Further questions obviously remain. Maybe we can sensibly *aim* to be reliable believers by reflecting on our belief-forming processes, and practising those we judge reliable. But will this strategy necessarily *succeed* in making us reliable belief-formers? In particular, isn't there an obvious difficulty? Our judgements as to which belief-forming processes are reliable will inevitably depend on our existing beliefs about the world and how we fit into it. So won't the naturalized recommendation at best get us into the state of our belief-forming processes *seeming* reliable (in the light of our existing beliefs), not the desired state of their being reliable?

Here the story gets complicated. In favour of the naturalized approach, note that it is by no means automatic that any given belief-forming process will vindicate itself as reliable. Somebody who habitually succumbs to *déjà vu* is quite capable of investigating this disposition of mind, as a psychologist might, and of concluding that in general *déjà vu*

"memories" do not have the facts that would make them true among their antecedents. More generally, the beliefs that *issue* from a given belief-forming process (such as: I've been here before) are not as a rule the beliefs which will *justify* that process as reliable (for this you would need something like: in general, whenever somebody has a *déjà vu* feeling, this is a result of that person having previously experienced the relevant scene).

So the naturalized epistemologist can certainly argue that the naturalized recommendation has some bite: it won't automatically leave everybody with whichever belief-forming processes they might naively have happened to start with. But still, even if the naturalized strategy leads you somewhere, this won't necessarily mean that it will lead you to reliable methods (as opposed to reliable-seeming methods). Might not different communities, with different intellectual starting-points, each apply the naturalized strategy, and yet get led off in different directions *ad infinitum*? Advocates of naturalized epistemologist will want to deny that this is in fact possible. They will urge that the world itself will act as a rudder steering different communities towards the truth (provided, of course, that those communities play their part by taking pains to ensure that their belief-forming dispositions are as well-adjusted to the world as possible).

Clearly there are questions of argumentative onus here. This last thought scarcely *proves* that the naturalized recommendation will lead to truth. But naturalized epistemologists will deny that they have any obligation to show that their method will *inevitably* be successful. Their claim is simply that it will, as it happens, succeed, and they will point out that there is no obvious reason to expect the world to frustrate such success. They will say that the desire for proof we feel here is itself a hangover from the Cartesian way of thinking: if epistemological justification did depend on Cartesian arguments transparent to the mind, then it would indeed follow that if you were justified you could prove that you were; but if justification depends only on being well-adjusted to the world, then there is no reason why being able to prove that you are justified should be a necessary condition of being justified.

Naturalized epistemology faces other problems. In giving examples of the evaluation of belief-forming methods I have concentrated on non-inferential processes like perception and memory. In such cases it is easy enough to see how reliability-evaluations are supposed to work, even if you have doubts about their further significance. But what about

inferential methods like deduction and induction? In particular, what about the kind of theoretical inferences that move scientists from their experimental findings to their theoretical conclusions? Here matters are far less clear, and showing that our inferential habits are well-designed to lead us from truths to truths is a substantial research programme for the naturalized epistemologist. (Note, however, that it's not the impossible programme of showing somebody who has no such habits that they ought to adopt some; rather it's the task of explaining to ourselves, in the light of our beliefs — including our inferred beliefs — about the general structure of the world, why it is that our inferential practices work.)

This is not the place to try to convince you of the superiority of the naturalized approach. But I hope I have done enough to show that it is a serious competitor to Cartesianism. Of course there are problems facing naturalism. But that is scarcely a reason for continuing to view all epistemological issues through Cartesian spectacles. After all, three hundred years of philosophical toil notwithstanding, there are plenty of familiar problems facing Cartesianism too. We shall be reminded of some of these in the next section. [For further defence of the naturalized approach to epistemology, see Papineau (1987).]

4. CARTESIANISM AND THE NEW SOCIOLOGY

In this Section I want to explain why Cartesians will take the new sociology of science to discredit science. The answer to this question might seem obvious. Don't Cartesians take the worth of science to depend on the supposed purity of the scientific mind, on the supposed fact that the scientific consciousness is swayed by reasoned argument alone? And doesn't the new sociology of science shows this supposed fact is not a fact at all?

However, there is a possible response open to Cartesians who want to defend science against the new sociology. Maybe, they can allow, scientists are on occasion influenced by ideological prejudice, or by personal ambition, or by other unscientific motives. But, they can then insist, it by no means follows that the views those scientists espouse are to be rejected. For there may still *be* good arguments for those views, in the sense that those views may still be ones that an ideally rational person acquainted with the evidence *would* uphold. That is, Cartesian defenders of science can appeal to the distinction between the norma-

tive question of which views ought to be accepted, and the genetic question of what actually led certain particular people to adopt those views.

An immediate difficulty facing this response is that it gives no grounds for a *general* faith in scientific findings. Maybe specific scientific theories can be independently vindicated by appeal to abstract standards of reason. But if the causes which actually operate in persuading scientists to accept theories are quite independent of the abstract reasons that can so normatively vindicate those theories, won't it be a freak if the findings of science are generally belief-worthy?

At this point a Cartesian needs to appeal to some kind of "invisible hand", some kind of unintended social mechanism which will ensure that in general the personal motives of scientists will always cancel out and leave us with belief-worthy theories. Now, perhaps a reasonable case can be made for such an invisible hand: perhaps there is a story about technological success and funding support, say, which will show that only the good theories survive in science. Indeed I myself shall argue for a somewhat different kind of invisible hand in Section **6** below. But let me leave this question up in the air at this point. For I now want to show that, whatever we think about invisible hands, there is a rather deeper reason why Cartesians have to resist the new sociology if they are to defend science.

I have been speaking of 'good arguments' and the 'dictates of reason'. But what are these dictates of reason? What authenticates the abstract normative principles with which the Cartesian hopes to vindicate scientific theories?

When we see how this question has been answered in recent philosophy of science, we will see why the new sociology of science raises a quite particular and intense difficulty for the Cartesian friends of science. But first we need a bit of general stage-setting. Note first that there is a strong tendency for Cartesians to reject *realism*. By realism I mean the view that success in judgement consists in correspondence to an independent reality. Cartesians who are realists face an uphill task. They have to show from first principles that the standards of reason which (a) recommend themselves to conscious human minds are (b) guaranteed to produce beliefs that correspond to an independent reality. It is not at all clear how to show this. So the natural move is to reject (b) and embrace *anti-realism*.

By anti-realism I mean the position that makes *reason* prior to *truth*.

Where the realist starts off with the notion of truth as correspondence (and so incurs an obligation to show that reasons are a good guide to truth), the anti-realist takes reason to stand on its own feet. The anti-realist takes success in judgement to consist directly in beliefs being backed by reasons, in their having the right provenance, rather than in terms of their correspondence to reality. And so for the anti-realist there is no question of showing that reasons are a good guide to truth. 'Truth' and 'reality', if they are used at all, are simply epithets attached to the picture of the world that reason leads us to. 'Truth' and 'reality' become by-products of rationality, rather than its aim.

Of course, not all Cartesians have been anti-realists. Descartes wasn't, for one. Nor was Locke. But Descartes and Locke could appeal to a God-given natural light to link up reason and realist truth. This had ceased to be a philosophical option by the middle of the eighteenth century. And since then the Cartesian tradition has unquestionably been overwhelmingly anti-realist.

So I have it that Cartesian epistemology leads to anti-realism. What now of the question raised above, "How can Cartesians authenticate the standards of reason?" It is all very well being told that reasons should be introspectively available, that they should automatically recommend themselves to consciousness. But can consciousness be sure that it has hold of the right reasons? Mightn't we be mistaken in assuming that inductive support, say, or simplicity, were good reasons for believing some scientific claim?

At this point the standard anti-realist strategy is to reject the problem. Anti-realists will argue that doubts as to whether intuitively good reasons really are good reasons only arise if we succumb to the realist misconception that the job of reason is to produce beliefs that correspond to reality. Perhaps, the anti-realist will concede, there are no higher meta-arguments which *prove* that inductive support, or simplicity, are guaranteed to produce true beliefs. But why demand such meta-justifications? Given anti-realism, there is no need to ask whether human standards of rationality are well-suited to generating true beliefs: a belief's being true just *is* a matter of its being generated by human rationality.

But there remain obvious worries about the 'human standards of rationality' being invoked here. We seem to be presupposing that standards of rationality are universal. But don't different people have different standards of rationality? Some anti-realists take a short line

here too. They simply argue that, despite any appearances to the contrary, alternative rationalities are not possible: the very idea of a community with radically different standards for evaluating beliefs makes no sense. Thus, for instance, some interpreters of the later Wittgenstein take the moral of his stories about people who measure piles of wood by cross-section rather than volume, etc., to be that such apparent alternatives are incoherent, that these alternative possibilities of thought are not real possibilities after all. [Cf. Lear (1982).] And the later Putnam has argued similarly that we have no alternative but to deem people who diverge from our basic standards of rationality to be crazy, to be incapable of human thought at all [Putnam (1981), Ch. 6].

However this blunt denial of any possibility of alternative rationalities will seem too quick to anybody working in the philosophy of science, to anybody concerned specifically with the rationality of scientific theory choice. For it is manifestly *possible* for people to suppose that modern astrology, or creationism, are supported by reason. And it would be tendentious at best to insist that such people must be crazy.

A natural ploy at this point is for the anti-realist to shrink the sample of humans whose thought is taken to manifest rationality. The anti-realist can, so to speak, reduce the data base against which hypotheses about the right rationality are to be tested. Not all human thinkers are *per se* rational. Rather it is the mature, serious thinkers, those with the time and resources and inclination to investigate nature critically, to whom we should look to identify the real standards of rationality. In short, it is the intellectual activity of *scientists* which serves to define the right reasons for belief.

And this, in effect, is the view of scientific rationality adopted by many contemporary philosophers of science: proposed methodologies, proposed sets of standards for scientific theory choice, are to be evaluated by comparing them with the actual intellectual practice of past scientists. The way to identify the right methodology for doing science is to think of proposed philosophies of science as themselves *theories*, which can then be tested against the data provided by the history of science. The right methodology is the one that best fits the history of science. [See in particular Lakatos (1978); Laudan (1977).]

Taken on its own this approach can seem puzzling. How can we hope to get normative conclusions about the right way to do science from factual premises about the way certain past people have reasoned

on past occasions? But this only remains a puzzle if you think of human intellectual practice as answering to some higher standard (such as, for instance, the need to generate conclusions that correspond to an independent reality). If there isn't anything more to the 'right' standards of rationality than simply the standards that come naturally to mature human thinkers, then how else should we identify those standards except by familiarizing ourselves with the habits of thought of the central figures in our scientific tradition? This way of adjudicating between philosophies of science may seem odd when taken out of context. But it makes perfect sense when seen against the background of Cartesian anti-realism.

So for Cartesian anti-realists, especially for those working in recent philosophy of science, rationality is by definition the way that scientists think. But now the intense threat posed to Cartesians by the new sociology of science should be clear. If you distil rationality from the history of scientific practice, then you are no longer in any position to argue that although scientists are often moved by ulterior motives, scientific theories themselves are in the abstract rationally supportable. For the things that move scientists are *per se* what's rational. And so, if scientists are no better than the rest of us, if their thoughts too are characteristically swayed by ideological prejudice and social context and self-interest, then it seems to follow that there is no such thing as rationality after all, beyond a degree zero a-rationality which leaves all beliefs, however ill-founded, on a par. For those who aim to found rationality on scientific practice, the new sociology of science discredits not just the motives of scientists, but the very idea of rationality itself.

Let me sum up the reflections of this section. I do not want to suggest that all Cartesians necessarily have the rationale I have outlined for resisting the new sociology. Just as there are some Cartesians who are not anti-realists (Descartes, Locke), so there are some anti-realists (Wittgenstein, Putnam) who do not attempt to distil scientific rationality out of the history of science. But, even so, there are a significant number of contemporary philosophers who have been led to attempt this, and I hope I have made it clear why these philosophers at least cannot allow space to the new sociology. [In this connection, it is worth observing that those philosophers whose history-based notion of rationality is most explicit, and whose anti-realism is correspondingly closest to the surface, are also those who are most insistent in resisting the inroads of sociology. Both Lakatos (1978) and Laudan (1977) are quite

adamant that sociology has no place whatsoever in the analysis of proper science.]

5. NATURALISM AND THE NEW SOCIOLOGY

I now want to consider what a naturalized epistemologist ought to say about the new sociology of science.

Contrast the situation of naturalism with that of Cartesianism. In the last section I argued that there were two routes by which the new sociology threatened Cartesian defenders of science. In the first instance, it raised immediate questions as to whether scientific findings were generally the result of good arguments. And then, in response to the thought that an invisible hand might ensure that in the abstract scientific findings were always rationally support*able*, even if the scientists had not so supported them, the new sociology threatened the very idea of rational support itself.

Neither of these discrediting implications go through if we adopt the naturalized perspective. Take the second threat first. This does not arise for naturalism, for there is nothing to draw naturalists to the anti-realist idea that rationality is nothing but those ways of thought natural to humans, and so, *a fortiori*, there is nothing to draw the naturalist to the idea of distilling standards of rationality from the history of science. For naturalists, a belief is rational just in case it is produced by a belief-forming process that is reliable for truth. Since naturalism here rejects the whole idea that justification demands arguments whose cogency is transparent to consciousness, naturalists by-pass the problem, faced by Cartesians, of showing from first principles that conscious human rationality is well-suited to generating beliefs that correspond to reality. And so there is no motive whatsoever for naturalists to become anti-realists. To put it at its simplest, naturalists build the idea of producing beliefs that correspond to reality into the very idea of rationality. And so, whatever other difficulties this approach might raise, they are not going to be difficulties which force naturalists to a truth-independent notion of rationality.

Consider now the first, more immediate, threat raised by the new sociology: it shows that in general scientific findings are not the results of good arguments. But why should this worry a naturalized friend of science? Once more, the naturalized epistemologist does not think of the belief-worthiness of science as depending on there being good

arguments for the findings of science: the requirement is only that those findings should generally be produced by reliable methods. And it is by no means immediately obvious that the new sociology succeeds in casting doubt on *this* requirement.

This last observation is the central point of this paper, and I want to develop it in some detail. But first let me make one thing clear. I am not here producing arguments to show, nor am I presupposing, that the findings of science are in fact generally true. I do, as it happens, believe this. And I also, as it happens, believe that the naturalized perspective allows us to see, in general terms, how it is that we are entitled to this belief. But that is not my present project (though the remarks introducing naturalized epistemology in Sections **2** and **3** above give some indication of how this larger project might be carried out). All I want to establish here is that a certain *negative* argument *against* believing in science does not carry the weight it seems to: that is, I simply want to show that, given naturalized epistemology, the new sociology of science should not disturb any faith in science you might independently happen to have.

My central claim, then, is that the new sociology of science does nothing to show that scientific practice is not generally reliable for generating true theories. It may well show that scientists are often swayed by prejudice, ambition and other ulterior motives. It may well show that the internal mental motivations of scientists are no different from those of the general public. But it by no means follows that the overall structure of scientific practice is not reliable for truth.

Let me elaborate this point in connection with that central plank of the new sociology, the "symmetry thesis" that true beliefs should receive just the same kind of explanation as false ones. [Cf. Bloor (1976), Ch. 1.] In a sense the naturalized friends of science can happily admit this thesis. They can accept that all scientific beliefs are produced by natural processes, and that there is no need for any special category of conscious reasons to play a special part in the generation of scientific beliefs. But they do need to insist on one point. The natural processes generating scientific beliefs can not, if they are to be reliable for truth, include *only* social factors, if this is understood to exclude the influence of the truth conditions of those beliefs, the facts that the beliefs in question are supposed to be about. Thus, to take an analogous case, while the reliability of our long-term memory does not demand that we derive our memories from some kind of Cartesian cogitation, it clearly

does require that amongst the causal inputs to the processes generating our memories should be the facts those memories are about.

So there is a minimal demand to be made by the naturalized friend of science: the explanation of creditable beliefs needs to differ from those of beliefs in general at least to the extent of allowing that amongst the causes of those beliefs are the truth conditions of those beliefs. But this now is why the new sociology fails to discredit science: it gives us no reason to suppose that this minimal demand is not satisfied. That the relevant natural facts play *no* part in the causation of scientific theories is an extremely strong claim, which is by no means established by the empirical findings of the new sociologists.

While they are not always explicit on the matter, many new sociologists seem to *think* that their work discredits the findings of science. But I would say that, if they do, this is because they are taking Cartesianism for granted. For a Cartesian, as we have seen, the fact that social influences play some part — a large part — in influencing the mind of the scientist, is indeed enough to discredit science. But something much stronger — that social factors play the *only* part — is needed to persuade the naturalized epistemologist against science.

Let me consider the two branches of the new sociology in turn. According to the macro-sociologists of science, scientists believe theories that fit their ideological position, or are otherwise consonant with the large-scale social context. If it were in general true that such social influences were *sufficient* for the acceptance of scientific theories, then we (I now abandon any further pretence of impartiality and identify myself with the naturalized epistemologist) should indeed be suspicious of science, for any conformity between scientific theories and the facts they were about would be purely happenstantial. But it does not matter that much if such social factors are *necessary* for the acceptance of theories. Thus, to take a familiar example, we could allow that a climate of high competitive capitalism was necessary for the acceptance of natural selection theory (and I mean the *acceptance*, not just the 'discovery', the thinking up of the theory — we can allow that natural selection would not have been accepted even if it had been proposed, as it arguably was by Buffon, in the *ancien regime*) and still hold that the social factors weren't *sufficient* for acceptance: that, in short, the truth of the theory was necessary as well. Provided the relevant natural facts are necessary for acceptance along with the social ones, there is no reason to doubt that scientific practice as a whole is reliable for truth.

(We should of course allow that sometimes, as in the Lysenko episode, the acceptance of theories is forced by social facts alone. But that such things can on occasion happen, in special circumstances, does not show that *scientific* institutions will in general allow such purely social determination of theories.)

Let me now turn to the micro-sociologists of science. It is interesting to note that from a naturalized perspective many of the descriptions micro-sociologists give of scientific practice are extremely reassuring about the reliability of science. No doubt it is true that scientists are often primarily motivated to attach their names to facts, and to build up the scientific credibility that will enable them to do this. But consider how they go about doing this. A prime concern, if they are to persuade others, is to ensure that their opponents won't be able to pick holes in their published claims. So scientists take care that their experiments are repeatable. They try to design experiments so that the move from the observed results to the desired theoretical conclusions depends on as few disputable auxiliary hypotheses as possible. They try to devise a variety of experiments, with a suitable range of controls, so as to leave their opponents no room for manoevre. And so forth.

To the naturalized epistemologist this is likely to seem just what is needed to ensure that scientific acceptance is reliable for truth. (And even the periods of "negotiation", when the fates of putative new facts hang in the balance, can be interpreted as periods of healthy agnosticism, during which further analysis and experimentation is needed to decide whether the new claims should indeed be entered in the archives.)

6. THE GAME OF SCIENCE

Consider this argument. "There is no compelling transcendent logic which forces scientists to their conclusions. That certain facts get written into the archives depends on socially contingent processes involving power, negotiation, and argumentative work. So scientific findings are merely social constructions, with no further epistemological authority."

At one level the appropriate response is obvious. It is only within the Cartesian perspective that justification requires arguments whose logic is compelling to consciousness (or at least to the consciousness of mature, scientific thinkers). Naturalism, by contrast, allows that justified

beliefs can come from whatever processes you like, however contingent or arbitrary they might seem from the point of view of conscious logic, as long as those processes are reliable for truth.

But the underdetermination of scientific practice by simple logic does point to a more substantial difficulty. If sound scientific habits of thought are not logically inescapable, is there not, then, a pressing need to explain why scientists do adopt their peculiar habits of thought? And surely at this level the only possible explanatory factors will be social ones.

And if so, won't this in itself discredit science? Put it like this. Maybe, say, the particular community of Californian endocrinologists studied by Latour and Woolgar do indeed embody habits of thought that are reliable for truth. But if they are sociologically caused to have those habits of thought, isn't there a sense in which they are just *lucky*? For all *they* have contributed to the enterprise, for all they know and care about, they could as well have been caused to adopt *un*reliable habits of thought.

And, more generally, if sound scientific practice is just a matter of luck, what is to ensure that scientists in general, as opposed to such particular communities as Californian endocrinologists, will have reliable habits of thought? In Section **2** I touched on the possibility of Cartesians evading their initial difficulties by positing an "invisible hand". But now I seem to be committed to an invisible hand myself, and indeed one that is required to do rather more work: for where the Cartesians just wanted a mechanism to select out rationally supportable *beliefs*, I seem to need a mechanism which will ensure that the belief-forming *methods* of scientists in general are reliable producers of truths.

It might seem that naturalized epistemology can once more dig in its heels and deny the difficulty. What does it matter *why* scientists have reliable habits of thought, provided that they *do* have them? Maybe the Californian endocrinologists *are* just lucky. Maybe it *is* just a matter of luck that scientists in general think in ways that produce truths. But why bemoan such luck, if it indeed obtains?

There is indeed a sense in which the fact of reliability is more important than its explanation. But, still, it would be a puzzling freak if scientists in general adopted reliable habits of thought, and yet there was no general explanation of why they did so. And, in fact, if that were the case, any claim to the effect that scientists are in fact generally

reliable would lose all plausibility: piecemeal demonstrations that particular scientific communities at particular times happened to have reliable processes would not prove anything, if we did not think there was some general explanation for this.

But fortunately there is an explanation. Namely, that scientists are their own naturalized epistemologists. Scientists are themselves perfectly capable of reflecting on what is required for the reliable production of truths. They are aware that unrepeatable experiments, or inferences based on *ad hoc* hypotheses without independent support, are unreliable guides to the facts. More significantly, scientists working in any particular speciality will have detailed ideas about the alternative hypotheses that need to be taken into account, the kinds of experimental controls that are required, the experimental techniques that are called for, etc., if truths about their specific subject matter are to be got at. And they will use this expertise to dismiss and discredit claims to knowledge that are based on unreliable methods.

In effect I am here appealing to a point made back in Sections **3** and **4**: it is perfectly possible for human beings to reflect on whether the belief-forming practices they engage in are reliable for producing truths, and to reconstruct themselves and their communities in the light of such reflections. As we saw then, it is an open question, which raises delicate issues of argumentative onus, as to whether such self-adjustment in pursuit of reliability will succeed in producing beliefs that actually are reliable (as opposed to ones which merely seem reliable). But we do not have to resolve this issue here. For the present objection is the specific one that it would at best be an accident if scientific practice was generally reliable (and that therefore we can not suppose that it is). And it is sufficient answer to this objection to point out that, whether or not the self-adjusting attempts of scientists to be reliable are *guaranteed* to succeed, it will scarcely be an *accident* if reliability is in fact the upshot of such efforts.

This last suggestion might seem to be going back on my factual concessions to the new sociology. Aren't I now denying that scientists are swayed by ulterior motives, and instead presenting them as impartial seekers after truth? But the conflict here is only apparent. Take the analogy of a sport like tennis, or, perhaps better, a gambling game like poker. Very few tennis players or poker players are primarily motivated by the desire to conform to the rules of their game. They are after money, or glory, or self-esteem, or perhaps just the satisfaction of

winning. And they will use such resources as they have at their command to achieve their ends: their psychological or physical advantage, or their superior skill, or simply their bigger bank-roll. Indeed many will be prepared to break or bend the rules, if they can get away with it, in order to succeed. But, still, there is a clear sense in which the rules are essential to an understanding of their activities. Whatever their underlying motives and means, they are still *playing tennis*, or *playing poker*, and as such still accept that (surreptitious infractions aside) certain rules govern their behaviour.

Similarly, it seems to me, with science. Individual scientists will have all kinds of extraneous ambitions (such as money, or fame, or prestige) and will use whatever resources they command (such as influential friends, or an established reputation, or somebody else's insights) to achieve them. But, still, the activity they are engaged in is *science*, in the sence of uncovering truths about nature, and in consequence scientists will accept that the very nature of their enterprise means that they are supposed to conduct themselves in ways that will reliably discover truths.

It might seem that the difficulty raised in this section emerges again a level up: isn't it just a matter of luck, susceptible of nothing but sociological explanation, that modern Western society should contain an institution devoted to developing reliable ways of uncovering truths? I agree this calls for sociological explanation: many societies do lack such institutions, and muddle along instead with uncritical acceptance of received doctrines, and it is by no means clear what accounts for the emergence and persistence of critical science in modern Western society. But I do not think that *this* discredits the findings of science in the slightest. On the contrary, I think we ought to celebrate our good fortune in having institutions for finding out truths.

One final caveat. It is no part of my purpose to engender dumb respect for modern scientific institutions, to assure you that the men in white coats are invariably finely-tuned truth-detectors, whose every dictum should be swallowed whole. Of course there are times where the ideology of a whole community, or the power of some group within it, gets some proposition accepted as fact despite its lack of appropriate pedigree. Of course the pressures for theories to be reliably generated do not always hold complete sway, and sometimes untruths will get written into the archives. But the appropriate response to this danger is

not to belittle science as a social construction, but simply to urge and constrain scientists to be more reliable in their findings.

University of Cambridge

REFERENCES

Barnes, B. and Shapin, S. (1979) 'Darwin and Social Darwinism: Purity and History', in Barnes, B. and Shapin, S. (eds.), *Natural Order: Historical Studies of Scientific Culture*, London, Sage.

Bloor, D. (1976) *Knowledge and Social Imagery*, London, Routledge and Kegan Paul.

Bloor, D. (1978) 'Polyhedra and the Abominations of Leviticus', *British Journal for the History of Science* **11**, 245—72.

Collins, H. (1985) *Changing Order*, London, Sage.

Farley, J. and Geison, G. (1974) 'Science, Politics, and Spontaneous Generation in Nineteenth-Century France: The Pasteur-Pouchet Debate', *Bulletin of the History of Medicine* **48**, 161—98.

Goldman, A. (1976) 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* **73**, 771—91.

Goldman, A. (1978) 'Epistemics: The Regulative Theory of Cognition', *Journal of Philosophy* **75**, 509—23.

Goldman, A. (1979) 'What is Justified Belief', in Pappas, G. (ed.), *Justification and Knowledge*, Dordrecht, D. Reidel.

Goldman, A. (1980) 'The Internalist Conception of Justification', in French, P., Uehling, T. and Wettstein, H. (eds.), *Studies in Epistemology: Midwest Studies in Philosophy, Volume* **5**, Minneapolis, University of Minnesota Press, 27—51.

Goldman, A. (1985) 'The Relation Between Epistemology and Psychology', *Synthese* **64**, 29—68.

Lakatos, I. (1978) 'History of Science and its Rational Reconstructions', in Worrall, J. and Currie, G. (eds.), *The Methodology of Scientific Research Programmes: Philosophical Papers, Volume* **1**, Cambridge, Cambridge University Press.

Latour, B. and Woolgar, S. (1979) *Laboratory Life*, London, Sage.

Laudan, L. (1977) *Progress and Its Problems*, Berkeley, University of California Press.

Lear, J. (1982) 'Leaving the World Alone', *Journal of Philosophy* **79**, 382—403.

Lynch, M. *Art and Artifact in Laboratory Science*, London, Routledge and Kegan Paul.

Papineau, D. (1987) *Reality and Representation*, Oxford, Basil Blackwell.

Putnam, H. (1981) *Reason, Truth and History*, Cambridge, Cambridge University Press.