

SOCIALIZING NATURALIZED PHILOSOPHY OF SCIENCE*

STEPHEN M. DOWNES^{†‡}

*Department of Philosophy
University of Utah*

I propose an approach to naturalized philosophy of science that takes the social nature of scientific practice seriously. I criticize several prominent naturalistic approaches for adopting “cognitive individualism”, which limits the study of science to an examination of the internal psychological mechanisms of scientists. I argue that this limits the explanatory capacity of these approaches. I then propose a three-level model of the social nature of scientific practice, and use the model to defend the claim that scientific knowledge is socially produced.

1. Introduction. I will argue that any attempt to naturalize the philosophy of science must take into account the social nature of scientific practice. I provide an account of the social nature of scientific practice that gives credence to the claim that scientific knowledge is socially produced. Most recent approaches to naturalizing the philosophy of science have downplayed the social nature of scientific practice, concentrating, instead, on explaining science using results from cognitive science. I will demonstrate that these approaches are flawed, first because they adopt what I call cognitive individualism, and second because scientific practice is irreducibly social. Cognitive individualism is the thesis that a sufficient explanation for all cognitive activity will be provided by an account of autonomous individual cognitive agents.¹ Cognitive individualism is related to methodological individualism as cognitive individualists take individuals as their unit of analysis in their explanatory theories, but cognitive individualism is a narrower position more akin to what philosophers of social science call psychologism. On this view what is cognitive is circumscribed by what can be described in terms of the internal psychological mechanisms of such autonomous cognitive agents. In contrast to this view I propose a model of the social nature of scientific knowledge that isolates several ways in which the term “social” can be understood, all of which must be included in an adequate model of scientific practice.

*Received January 1992; revised June 1992.

[†]This paper was completed during a post-doctoral fellowship at Northwestern University. Arthur Fine, Steve Fuller, David Hull and an anonymous referee provided helpful comments on earlier drafts.

[‡]Send reprint requests to the author, Department of Philosophy, 338 Orson Spencer Hall, The University of Utah, Salt Lake City, UT 84112, USA.

¹In conversation, Maurice Finocchiaro drew my attention to this point.

Philosophy of Science, 60 (1993) pp. 452–468
Copyright © 1993 by the Philosophy of Science Association.

My position draws on work in the sociology of science, but has more in common with “constructivist” views such as those presented in Latour and Woolgar (1986), Latour (1987), and Collins (1985) than the interest theory of the Edinburgh School most familiar to philosophers of science (e.g., Bloor 1976, and Barnes and Bloor 1982).

On my view all of scientific practice contributes to the production of scientific knowledge. Scientific practice includes carrying out experiments, extending existing theories into new domains, postulating and testing new theories or hypotheses, replicating experiments, peer reviewing of papers, and using papers published by members of a field. When I refer to scientific cognition I mean the activity that leads to the production of scientific knowledge. Scientific knowledge I generally take to refer to the certified claims that appear in finished research papers. Scientific cognition has often been taken to be a process that goes on inside the heads of individual scientists, and this is the view of many working now in naturalized philosophy of science. Defining scientific cognition as the activity that leads to scientific knowledge widens the scope of naturalistic philosophy of science, drawing attention to factors other than the psychological capacities of individual scientists. Once we take account of the varied activities that go into the production of scientific knowledge it is easier to make sense of the claim that scientific knowledge is a social product.

My argument proceeds in several stages. First, I give a definition of cognitive individualism and argue that the position arises from naturalists’ attempts to take over the explanatory agenda of traditional philosophy of science. Second, I argue that cognitive-individualist attempts to explain hypothesis evaluation in science fail. Third, I examine some attempts by cognitive individualists to deal with social aspects of science. Finally, I develop and defend a model of the social nature of science generated from three distinct interpretations of the term “social”. I conclude that with my model we can make sense of the claim that scientific knowledge is socially produced.

2. Cognitive Individualism in Naturalized Philosophy of Science. Much recent naturalized philosophy of science can be characterized as cognitive science of science. Cognitive scientists of science apply techniques from the cognitive sciences to issues previously examined by the philosophy of science, for example, theory change, theory evaluation, the nature of scientific theories, and scientific discovery. In general, cognitive-science-of-science researchers set their work apart from more traditional approaches to philosophy of science such as logical empiricism. They are naturalists who hold that empirical results from current sciences, particularly the cognitive sciences, should inform and constrain

philosophical theories. Some further suggest that philosophy should be abandoned and replaced by a particular science, so some suggest that cognitive science of science should replace philosophy of science (e.g., Giere 1988, chap. 1). Recent works by Churchland (e.g., 1990), Giere (e.g., 1988), Simon (e.g., Langley et al. 1987), and Thagard (e.g., 1988) exemplify this approach.

I argue that this approach in naturalized philosophy of science is inadequate. One important reason for this inadequacy is the adoption of what I call cognitive individualism. A cognitive-individualist account seeks an explanation of a cognitive phenomenon purely in terms of individuals' psychological processes, understood as internal mechanisms. For example, the choice of a particular hypothesis in a scientific field is explained entirely by appealing to the psychological mechanisms of an individual scientist in that field. In an explanation of science this approach is crucially limiting because it ignores the important social dimension of scientific cognitive activity.

The definition of cognitive individualism is both broad and vague at this point, and many explanatory approaches could be deemed cognitive individualist under such a definition. Later I will point out some of the difficulties arising specifically from the vagueness of the definition. Now I turn to the relation between cognitive-individualist explanations of science and the more traditional approach in philosophy of science.

Traditional philosophy of science gives an account of scientific theories and their explanatory capacities. I use the expression "traditional philosophy of science" to denote logical empiricist approaches such as Carnap, Reichenbach, and Hempel's, as well as recent nonnaturalistic work such as van Fraassen's (1980) and Glymour's (1981). This work involves an examination of the logical relations between elements of a theory and observations supporting that theory or predicted by it. As a result of this rational reconstructivist approach, most traditional philosophy of science is neither couched in terms of individual scientists nor social groups but in terms of theories themselves characterized as sets of sentences, or on the recent semantic approach sets of models (see, e.g., van Fraassen 1980). This result is apparent even in more naturalistic work such as Laudan's (1977, chap. 3).

Accounting for the nature of theories is also a crucial issue in naturalized philosophy of science, but here the related issue of theories' embodiment becomes crucial. To date naturalists have proposed that scientific theories are internal mental representations, and so philosophers should turn to cognitive science—whose object of inquiry is such representations—to account for their nature. Thus the embodiment issue is confronted: Scientific theories are located in individual scientists' heads. This approach is consistent with most research in the cognitive sciences

where a cognitive phenomenon is accounted for in terms of an internal psychological mechanism. For example, mechanisms have been proposed that underlie deductive logic, abductive reasoning, inductive reasoning, pattern recognition, past tense verb learning and so on. I turn to a more detailed examination of current naturalistic accounts of the nature of theories below.

The naturalistic turn of the cognitive scientists of science has some parallels with the historical orientation in the philosophy of science. Advocates of the latter argue that rational reconstructivist accounts of science were mistaken because these accounts did not fit with historical data. Historically oriented philosophers present case studies in the history of science that challenge various assumptions of rational reconstructivists' theories. For example, if a case study demonstrates that the H-D model of explanation was not operative, it is argued that this provides evidence that the model is not generally applicable as an explanation of how science proceeds. Compare this with the approach of a cognitive scientist of science who may argue that a particular philosophical theory is not generalizable because of what we know about human psychological makeup. While historically oriented philosophers claim particular historical contingencies are crucial to explaining the production of scientific knowledge, cognitive scientists of science propose specific psychological processes.

Traditional philosophers of science, through their rational reconstructions, provide an abstract or idealized account of all scientific activity, not simply an account of the ideal individual scientist. These philosophers use the relevant abstractions to reconstruct the state of a field in relation to external events, plus the relevant community standards of the field, and also some norms to govern reasoning processes. Issues such as the nature of scientific progress and the accumulation of scientific knowledge are then addressed in terms of the relevant rational reconstructions of scientific practice. Within this framework, questions such as "Does science converge on the truth?" are addressed.

It may seem a short step from logical empiricists' concern with the reasoning that supports theories to the cognitive scientist of science's account of individual psychological processing, but this step misses the crucial idealizations and abstractions that have been carried out by traditional philosophers of science. The move from the abstract entities that are theories on the traditional approach to the theories embodied in individuals' psychological processes runs into problems. In taking up where traditional philosophers of science left off, cognitive individualists end up placing properties of science as a whole within individual scientists' internal representational structure. The former's method of analyzing science builds all aspects of scientific practice into individual scientists' heads. A nat-

uralized philosophy of science must be able to address issues such as the success of science as a whole, rather than simply the relations between individual scientists' internal representations and the world. I will argue below that a correct understanding of the social dimension of science, ignored by cognitive individualists, provides a naturalistic approach more adequate for addressing such issues.

Many philosophers argue that both social and psychological factors are irrelevant to the philosophical study of science (e.g., Lakatos 1981 and Laudan 1977, along with Reichenbach and Popper). One way of setting up this argument is to make a distinction between cognitive and noncognitive factors in science. The claim is that when an account of successful science is required it should be given in terms of cognitive factors. An example would be that scientific theories can be determined as approximating the truth, and such an approximation provides the metric of the theory's success. A corresponding noncognitive account may explain the success of a theory in terms of the relation between predominant interest groups that prevailed in the wider society and those in the scientific community (see, e.g., Forman 1971). Philosophers defending the cognitive/noncognitive distinction argue that noncognitive explanations should be reserved for cases of unsuccessful or failed scientific theories only, and that success in science can only be explained by cognitive factors (see, e.g., Lakatos 1981 and Laudan 1977). This view is presented most clearly as Laudan's (1977) arationality thesis: All rational scientific practice is to be explained by cognitive factors, and anything left over can be explained by sociologists or psychologists. Ironically, on this account the cognitive science of science is concerned with noncognitive factors in science, specifically psychological processes, as under the cognitive/noncognitive distinction both psychological and sociological factors are irrelevant in explaining successful science.

Cognitive-science-of-science researchers further twist this story with a new version of the distinction between cognitive and noncognitive factors. They distinguish between cognitive processes as investigated by cognitive science and social processes (see, e.g., Churchland 1990, Giere 1988, and Thagard 1988). A similar kind of explanatory hierarchy is set up as in the first version of the distinction. Some cognitive-science-of-science researchers allow that science is an activity involving many practitioners, yet claim that the important cognitive output of science will be explained by a theory of the individual scientist's psychological processes (Churchland 1990, 248; and Thagard 1989b, 80). For these researchers the extent to which science is social is simply that autonomous cognitive agents are gathered into groups. The fact that scientists are collected together in particular groups (subfields, laboratory teams, paper collaborators, and so on) is not to be taken into consideration when an account

of the cognitive output of science is desired. Nor are any properties of these groups, over and above the psychological capabilities of the individuals, taken into consideration. Cognitive individualism is exemplified in this reworking of the cognitive/noncognitive distinction. The claim is that a sufficient account of what is cognitive in science can be provided by an account of the psychological mechanisms of individual scientists.

Before turning to an examination of the social nature of science, let us look at a case that leads to problems for cognitive individualists. Consider the case of evaluating hypotheses in science. Several naturalistic philosophers of science address this issue. For example, Giere (1988, chap. 6) proposes that "satisficing" is a psychological mechanism that enables scientists to choose between rival hypotheses. Thagard (1989a) suggests that individual scientists have psychological mechanisms that allow them to evaluate hypotheses and lead them to choose the best ones. The mechanisms that each scientist possesses are simulated by a computer program called ECHO, which simulates, among other things, Darwin's selection of the evolutionary hypothesis. If each scientist possesses mechanisms for selecting the best hypotheses, it seems reasonable to ask why it is that scientists work in teams, ranging in size from two- or three-member laboratory groups to fifty strong research teams to do just this job. Perhaps the cognitive individualist would reply that individual scientists have the capacity to evaluate hypotheses by themselves, but they need not do all the work by themselves if others can assist; and more importantly, there may be time constraints on producing results, so they have to work in teams. Whatever the specific reason, working in teams is incidental to the hypothesis-evaluation process. This reply has several problems.

Consider the practical constraints on any individual attempting to do all the work evaluating a particular hypothesis that leads to a scientific discovery. First, if it is practically possible for an individual to carry out the tasks involved in hypothesis evaluation, then a sufficient empirical account of such a scientific discovery may be one that gives an account of the psychological processes of a somewhat idealized individual scientist. There is a problem with this route. Testing, or evaluating, hypotheses requires more than just psychological capabilities. For example, the manipulative skills required either to operate specialized instruments or to prepare materials for examination by such instruments are an integral part of hypothesis evaluation in many sciences. Consider the tissue slicing and preparation that precedes the examination of samples under an electron microscope. Generally labor is divided according to distribution of skills among the members of the scientific community.²

²Rouse (1987, chap. 4) influenced my formulation of this point. The underlying point is to emphasize a more thorough investigation of the role of experimental practice in science as many philosophers have done recently (e.g., Franklin 1986, Galison 1987, Gooding et al. 1990, and Hacking 1983).

Second, it may simply be practically impossible for anyone to evaluate the hypothesis by herself. Once the tasks involved are studied in enough detail, it may become clear that considerations about varying individual cognitive competence or available memory space prevent any one individual from carrying out the tasks. In this case the empirical status of the cognitive-individualist approach is threatened. Now the cognitive individualist may avail herself of the following route: turn the claim into a claim that in “ideal conditions” an individual could complete the task.

The claim that in ideal conditions an individual could carry out the required cognitive task is not a straightforwardly empirical claim. It may explain the right range of phenomena, but not under the right empirical constraints. One can claim that something could in principle be the case, given certain background assumptions, but only if the background assumptions are plausible. A useful analogy is introduced by Bach and Harnish (1982). They point out that a frictionless plane is a useful idealization in classical physics, but a surfaceless plane is not. I return to this issue about idealizations in more detail below.

Claims cast in terms of ideal conditions can be turned into normative claims about how a particular task might best be done, given, say, an infinitely capable individual processing device, but such normative talk does not address the empirical issue of how a particular task was carried out or how a particular cognitive output was produced. When developing a naturalized epistemology, empirical constraints must be at the forefront. The naturalized epistemologist faces the task of providing norms for the best cognitive performance given the empirical constraints on the agents involved. When it comes to evaluating hypotheses, scientists can deal with such constraints by working in teams and therefore dividing the cognitive labor (see Kitcher 1990). In ignoring this fact cognitive scientists of science overlook the crucially social nature of the production of scientific knowledge.

3. The Social Nature of Scientific Practice. In this section I develop three distinct senses of the term “social” that are relevant in the study of scientific knowledge. There are many ways of presenting the case for the social nature of scientific knowledge. For example, case studies can be presented which demonstrate that social interaction is necessary for a particular cognitive output. This has been the approach prevalent in the work of sociologists of science who have presented many case studies on many different areas of science, for example, Pickering (1984) in particle physics and Latour and Woolgar (1986) and Knorr-Cetina (1981) in biochemistry. Here I present a more general case for the social nature of scientific cognition. I first propose that scientific theories cannot be adequately understood without considering their social or public embodiment. Sec-

ond, I present some thought experiments that motivate the postulation of two further levels of the social present in scientific practice. Finally, I rebut the possible objection that only current science produces social cognition since the arrival of the research team approach after World War Two, and that previous science could be explained by a cognitive individualist account. Overall my aim in this section is to overthrow the belief that the “social” is somehow extrinsic to science in the form of “social factors” affecting science and to show that the production of scientific knowledge is an intrinsically social activity.

Before presenting my own model of the social nature of science, let me briefly examine two recent attempts to give such an account. Recently Thagard (1988) and Churchland (1990) have presented accounts of the social nature of science against the background of the working assumption that the production of scientific knowledge is best explained by appealing to the internal psychological mechanisms of individual scientists. As such these accounts have at least an implicit agenda of downplaying the social nature of scientific practice. Their attempt to marginalize so-called social factors demonstrates an inadequate understanding of the best way to analyze the term “social” with respect to scientific practice. (Of course several philosophers have attempted to develop positive accounts of the social nature of scientific practice. My approach shares features with Rouse’s 1987 and Fuller’s 1989. Several other philosophers have developed accounts of science that emphasize a positive role for social processes in producing knowledge, for example, Hull 1988b and Longino 1990.)

Thagard (1988, 186) describes the social nature of science with a model of various autonomous individual scientists sending their theories or hypotheses to a central reviewer. The reviewer in Thagard’s model has the functional role of the peer-review system in science, but the model is highly schematic. Science is social, on Thagard’s view, as a result of the fact that a group of psychologically internally equivalent scientists feed their findings to the reviewer. This account has no room for variations in the types of interactions between the scientists and the reviewer, nor does it have room to account for many of the other social interactions that go on in science. Thagard (*ibid.*, 189) does advance the claim that we could develop computer models in which the scientists modeled were not internally equivalent but behaved instead like idealized Kuhnians or Popperians, but this is hardly an attempt to model what is actually going on in scientific practice.

Churchland (1990, 248) reintroduces Laudan’s arationality thesis by claiming that social factors “corrupt” science. Churchland claims that we can distinguish between the social determination of conceptual change and the influence of the “world itself” on conceptual change, arguing that the latter can be investigated independently. He claims that if social fac-

tors determine conceptual change we have no recourse but to global skepticism. What Churchland does not establish is how one distinguishes between the effect of so-called social factors and the “world itself” on an individual cognitive system.

Thagard’s account downplays the role social interaction plays in science but at least acknowledges that it plays some role in the production of scientific knowledge. Churchland assumes that we can distinguish clearly between the influence of the world and the influence of society on conceptual change. Without a model of just exactly what the term “social” refers to, this distinction is very hard to fathom. Developing such a model is the task I turn to now.

One motivation for turning to a social rather than a cognitive-individualist approach in philosophy of science comes from considering the complexities arising from attempts to explain scientific theories. Churchland (1990), for example, claims that theories are embodied in an individual’s neural makeup. Such a view can be challenged by presenting the many features of scientific theories that are not captured by treating them as individuals’ representational structures. Theories can be embodied in textbooks; they can be crucially linked to instrumentation; students can use certain parts of a theory without ever learning other parts; theories can be general world views such as “evolutionary theory”, which change their content over time (consider the differences in evolutionary theory between 1890 and 1990; see Laudan 1977) and so on.

Cognitive scientists of science propose several different accounts of the nature of theories. Thagard (1988) claims they are schemata, Giere (1988) proposes that they are mental models, and Churchland (1990) that they are points in the weight space of a parallel distributed processing network (or partitions in weight space). These accounts all hold that the notion of a theory is captured adequately by a representational structure within an individual’s head. None of the accounts captures each particular theory’s constantly changing nature across time, or what it is about the theory that remains permanent across time, or the fact that scientists using the same theory need not have exactly the same version of that theory available to them. Examples of this latter phenomena are Dirac’s independent derivation of some of the quantum results, and Wallace’s independently developed version of evolutionary theory. An account of individual scientists’ representational structures does not do justice to these features of scientific theories if the aim of delineating the representational structure is to capture the whole theory. There are cases in the history of science where it looks as if a scientist conceived of and applied a whole scientific theory by himself, for example, Einstein’s development of special relativity. But even here, an account merely in terms of the theory as an

internal representation is inadequate. (I expand this point below in a response to the Robinson Crusoe objection.)

Cognitive-science-of-science researchers' accounts of the nature of scientific theories may be a legacy from traditional philosophy of science. Philosophers of science have traditionally been concerned with the *post facto* reconstruction of theories, which traditionally required no concessions to their historical development. So two competing theories are formulated in their entirety and then assessed by their ability to cope with the relevant observations. Cognitive scientists of science explicitly strive to overcome this "static" view of theories (see, e.g., Giere 1988, chap. 3). But in their new accounts of theory evaluation, the static picture is recreated in terms of internal mental representations. For example, Thagard (1989a) reconstructs Darwin's psychological process of comparing evolutionary theory with creationism. A computer program, ECHO, compares the ability of hypotheses from both creationism and evolutionary theory to cope with appropriate data. Giere (1988, chap. 8) reconstructs the process of comparing continental drift theory with its predecessor. He suggests that individual scientists use a form of decision matrix to choose between rival theories construed as models. Both accounts present pictures of an individual researcher who internalizes two theories, say *T1* and *T2*, and uses a stock of psychological mechanisms to decide which theory is superior.

I believe that a theory can persist through generations of research (despite its gradually changing nature) because of its embodiment in textbooks, research papers, instruments, computer memories and so on, all of which are the shared property of the scientific community (see Hacking 1983). On this view the problems of how one transfers one whole theory from the head of one researcher to the next dissolve. One can also attribute beliefs about a theory's truth to scientists without having to assume that they have in their heads a representation of that entire theory which they have independently assessed as best describing the world. If science is conceived as a communal project involving the communication of knowledge and the accumulation of knowledge for use by other contemporary scientists, as well as future generations of scientists, it becomes obvious that theories need some form of embodiment in addition to internal representational structures of individual scientists. This conception of scientific knowledge communication and accumulation is similar to Campbell's (1969) "fish scale model of omniscience".

Cognitive-science-of-science accounts of scientific theories share the implicit claim that one account will suffice for all scientific theories. This claim is overly ambitious. The many ways that any one theory is manifest, at the same time or over periods of time, indicate that a universal account of "the nature of scientific theories" will be illusory (see Downes

1992). We may have to be satisfied with several different accounts of theories, each having its own appropriate ranges of application, as well as a variety of objects that count as theories. One advantage of my approach is that this possibility is not ruled out.

I now address a possible objection to my alternative account of the nature of scientific theories. This objection is made from the point of view of someone capitalizing on the vagueness in my definition of cognitive individualism.³ Imagine someone named Robin Crusoe is shipwrecked and stranded on an island. She is a reasonably mature adult and prior to her sea voyage she had some training in science, physics and some botany. To while away the hours during the day time, she carries out some rudimentary experiments on plants, and at night she attempts to plot out the planet's paths. We should be willing to say that she is doing science, at least on any reasonable definition of science. The objection to my account refers to this scenario, and runs thus: Surely on the definition of cognitive individualism given above, we can give a cognitive-individualist explanation of Robin Crusoe's scientific practice. She only has access to the theories she has remembered and hence brought with her as mental representations, and so in this case we do not need any more of an account of scientific theories than one that treats them as her mental representations.

My reply involves introducing another related scenario. Robinson Crusoe is stranded on an island after a shipwreck, but he is only ten years old. He is determined that he will survive until rescued, and so over time he tests various plants and beasts to see how they taste and eats the ones he can tolerate. At night he gazes at the stars and over time begins to realize that some of them move relative to the rest. Now in this case we may be reluctant on any account to call any of Robinson's activities science. My challenge to the cognitive individualist is to explain, entirely in terms of internal mental representations, the crucial differences between the two scenarios that lead us to consider Robin to be practicing science and Robinson not.

The crucial differences between the two activities are in the relation that Robin's activity bears to a shared body of knowledge in which she was able to partake. She knows that various planets cannot be seen without a telescope and that they have elliptical orbits. She also knows how to identify several species of plants, the ones that are particularly widespread for example, and can spot resemblances between unfamiliar ones and familiar ones. The theoretical knowledge she has is not in virtue of her mental representations alone but in virtue of their relation to a com-

³William E. Morris presented this objection, and discussion with him led to the formulation of my reply.

mon stock of knowledge. Where she got these representations and what they relate to is important. What distinguishes Robin's from Robinson's activity cannot be made explicit without moving outside the cognitive-individualist account and referring to the communal nature of scientific theories and alternative embodiments for them. Thus the cognitive-individualist approach is insufficient to account for the role of scientific theories in scientific practice.

The public embodiment of scientific theories is one of the social aspects of scientific practice, which I will call level 1. I now introduce two further distinct social levels in scientific practice and provide some conceptual arguments supporting the claim that scientific cognition is a social product.

Consider the following kind of interaction: If a particular output could be produced by a group of people, and yet it could not be analyzed in terms of the contributions of the individuals in a group (say by adding each of their contributions together), we could say that the relevant product was a social product. The following example captures this notion of social interaction, which I will call level 2. Marching bands often perform routines that leave their members distributed in such a way on the football field as to write the name of their team. The "output" of the band, the name of the team read by the fans in the stands, is not something that can be analyzed in terms of the "outputs" of individual band members. Imagine a case in which the band wrote an obscene word on the field. We could interpret this only as a case of group defiance, or mischief. Perhaps extreme coercion on behalf of a small subgroup might produce the word, but the relevant output still cannot be analyzed in terms of the contributions of individual band members.

To sharpen this example imagine what idealization, or idealizations, would enable an individual band member to produce the desired effect. The point of this extension of the thought experiment is to highlight the difficulty of appealing to idealizations. Arguments claiming that in ideal conditions an individual could produce a certain effect generally do not hold for cases of social interaction that conform to this model. How would an individual band member go about producing the desired output? The kind of ideal band member we would have to conjure up would certainly have to defy many physical laws, such as being in many places at one time. In cases like this Bach and Harnish's (1982) analogy is telling. Recall that they argued that certain idealizations were informative and useful in explanations, such as frictionless planes, but others, such as surfaceless planes, were not. If an idealization invokes inconceivable notions, such as a surfaceless plane or an individual who can be in hundreds of places at the same time, it is not useful. In this kind of case it is more productive to seek an explanation that does not rely on such idealizations,

specifically an explanation that acknowledged that social interaction is a crucial explanatory factor.

A more general naturalistic position underlies this specific point. Philosophers who invoke brains in vats or science fiction provide empirically impossible explanations for phenomena. Similarly, some attempt to account for counterexamples to a theory with progressive idealizations of the original explanatory theory. Naturalistic philosophers should be satisfied with having their theories account for empirical facts, and idealizing their theoretical entities within the bounds of what is empirically plausible. Wilkes (1988) and David Hull (1988a) make a related point to this one.

In cases that conform to the marching band example, I propose that the output results from particular kinds of social interaction. In science this kind of interaction occurs in complex laboratory work where no one researcher is entirely responsible for the ultimate cognitive output. This kind of social interaction in science is discussed in some detail by Rouse (1987, 109), who uses Fleck's (1979) account of the discovery of the Wassermann reaction to focus the point.⁴ The relevant cognitive output in such cases is a validated hypothesis or, in more concrete terms, a research paper. This is level 2 in my three-level model.

Let us now turn to the third level. Consider the question "Are there any lone scientists?" The following thought experiments are based on this question. Imagine two women identically dressed doing identical tasks in identical rooms. They are performing operations with various test tubes full of liquid; they pour the liquid into a machine, then write down figures displayed on a screen at the side of the machine. One woman is involved in an important scientific experiment and the other is rehearsing a role in a play, which depicts a post-nuclear-holocaust society, who performs the actions of experimental scientists as a religious ritual that invokes a pre-holocaust past. Neither woman defines her own context by her activities. Each woman's involvement in a greater picture, the first in the biochemistry community and the second in the theater group, provides the context for these activities.⁵

⁴There is a similarity between the social interaction that leads to scientific cognition and the social interaction that Putnam (1975, 271) claimed contributed to meaning. Putnam's notion of the "division of linguistic labor" (*ibid.*) influenced my thinking on the division of cognitive labor in science.

⁵A related point can be made by imagining how one could stop the women from being scientists. As we can establish that they are scientists only from their context, we can prevent them from being scientists only by manipulating their context. These types of considerations cause problems not only for cognitive individualists, but also for naive laboratory anthropologists/ethnographers who have no theoretical considerations to fall back on concerning larger issues in science such as validation of discoveries and the honorific nature of facts. Latour and Woolgar (1986) and Latour (1987) do not fall into this category as they devote a large amount of space to the validation of scientific claims once they leave the laboratory.

Consider another scenario in which two identical computer programs are running on identical machines connected to two larger identical machines in two identical rooms; both produce a printout of numbers and sentence fragments in English. One set of computer printouts is used by an artist to form a continuous frieze down the entire spiral length of the Guggenheim Museum, and the other is taken by a group of scientists who interpret its numbers and sentence fragments and submit them in the form of a paper to *Science*, claiming that the computer has made a discovery. One cannot distinguish between the activities of the two computers simply by appealing to facts about their output and their physical makeup. The second computer is interpreted as doing science because it has a certain place within the scientific community.

A final example further illustrates the level 3 sense of the social. At a coronation the crowning of a queen and her role makes sense only against the background of the beliefs of the nation in which she is crowned.

These thought experiments introduce a level of the social that helps focus both the scale of investigations of science and its inescapably social nature. The wider context is crucial in making sense of an activity as a contribution to science. In each of the first two examples above, facts about the wider context of the specific activities provide answers to most of our questions about the significance of each of the two identical types of activity and enable us to distinguish them from one another. The coronation example invokes a notion of social interaction that is shared with the thought experiments. Scientific discoveries are analogous to the coronation case. Often the acceptance of a discovery and its place within a body of communally shared knowledge distinguishes it from similar episodes that are not accepted as scientific discoveries (see Brannigan 1981).

With level 3 I attempt to capture what some consider to be the boundary between society in general and the society of scientists. I agree with critics of Mertonian sociology of science and later critics of Edinburgh School sociology of science who have pointed out that this boundary is not distinct (see, e.g., Mulkay 1976, Woolgar 1988, and Latour 1987). The fact that scientific work has a point is relative not only to a particular community of scientists but to society as a whole. This claim that level 3 plays an important part in the production of scientific knowledge does not reduce to the more familiar claim that wider societal interests determine the content of scientific theory (see Forman 1971, Bloor 1976, and Barnes and Bloor 1982).

By distinguishing the three levels I aim to show that the term "social" does not simply denote "that which is outside science". My aim in introducing the thought experiments is suggestive rather than an attempt to demonstrate the necessity of the three social levels. I have attempted to use the thought experiments to draw out three senses of the term "social"

that are run together in many of the sociological and historical case studies. My claim is that all three levels of the social are intrinsic to scientific practice and, hence, contribute to the production of all scientific knowledge.

Let us consider the objection that we are tempted to a social account merely because science has recently become such a large-scale affair. Of course, the objection runs, science is *prima facie* a social activity in the present day because of the advent of large research teams. What is objectionable is the claim that scientific knowledge in general is a social product. For example, would not the work of a seventeenth-century scientist such as Boyle be better accounted for by cognitive individualists?

Shapin and Schaffer (1985) show that Boyle required technicians to set up and operate his famous airpump experiments. This alone does not support the claim that scientific knowledge in Boyle's time was socially produced. Shapin and Schaffer's crucial point is that these technicians were not able to report the results of experiments. Experimental results were established by the invitation of a selection of appropriate witnesses, members of Boyle's own social peer group, to observe the experiments. This procedure was required to corroborate the experimental results. Shapin and Schaffer's add that Hobbes believed that there could be no science in the state of nature, as without a properly ordered society there could be no communication.

In my model, witnesses are necessary in the production of the results, level 2, and are also necessary in certifying these results in the wider society, level 3. The role as nobles and gentlemen is bestowed upon the witnesses by society, and so they are suitably qualified to present these results to that society. The claim is not simply that science is usually practiced by groups of people who, given the right conditions, may be dispensable but, rather, that scientific knowledge, the final validated outcome of scientific practice, is a result of certain crucial social interactions.

4. Conclusion. I have highlighted several aspects of the social nature of science. First, I pointed to the public embodiment of theories, which was necessary for their transmission across communities and across time; I call this level 1. Notice that part of this aspect of the social nature of scientific theories need not refer to people themselves, as theories are in part the books, instruments, and computer programs in which they are embodied. The notion of social at work here relies on the idea that such embodied theories are public or communal property (see Latour and Woolgar 1986, Latour 1987, Woolgar 1988, Rouse 1987, Ackermann 1985, Hacking 1983). Second, I used a thought experiment to illustrate how social interaction can lead directly to cognitive products; this is level

2. Third, I developed some thought experiments to motivate the idea that the wider society is crucial for the validation of scientific discoveries, providing a point to the scientific enterprise, and making sense of the relationship between the scientific community and society at large. This latter is level 3. Finally, I used the three-level model of the social to demonstrate that seventeenth-century science provides us with an example of the social production of scientific knowledge, thus blocking the objection that science is only *prima facie* social in its current big-science guise. Once the notion of social is analyzed in this way it is easier to understand claims that scientific knowledge is socially produced. All scientific practice involves at least the three levels of the social I have isolated.

Two sets of implications can be drawn from my arguments. The first set is descriptive and addresses the descriptive adequacy of the cognitive-science-of-science type of approach in naturalized philosophy of science. The main thrust here is that we should not be content with any descriptive account of science that is couched purely in terms of individual scientists' putative psychological processes. The second set is normative and connects with the difficult problem of providing any naturalistic yet normative account. The issue here is centered around the maxim "ought implies can" and is that if we attempt to derive norms for particular activities from descriptive accounts, the norms must appropriately apply to those activities as they actually occur. My general conclusions are that cognitive science of science cannot provide adequate descriptive accounts of science if it retains cognitive individualism, and second, that a naturalistic philosophy of science with a normative dimension must have a descriptive basis that takes into account the distinctive social nature of scientific practice.

REFERENCES

- Ackermann, R. J. (1985), *Data, Instruments and Theory*. Princeton: Princeton University Press.
- Bach, K. and Harnish, R. M. (1982). "Katz as Katz Can", *Journal of Philosophy* 79: 168–171.
- Barnes, B. and Bloor, D. (1982), "Relativism, Rationalism and the Sociology of Knowledge", in M. Hollis and S. Lukes (eds.), *Rationality and Relativism*. Cambridge, MA: MIT Press, pp. 21–47.
- Bloor, D. (1976), *Knowledge and Social Imagery*. London: Routledge Kegan Paul.
- Brannigan, A. (1981), *The Social Basis of Scientific Discoveries*. Cambridge, England: Cambridge University Press.
- Campbell, D. (1969), "Ethnocentrism of Disciplines and the Fish Scale Model of Omniscience", in C. Sherif and M. Sherif (eds.), *Interdisciplinary Relationships in the Social Sciences*. Chicago: Aldine Press, pp. 328–348.
- Churchland, P. M. (1990), *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. Cambridge, MA: MIT Press.
- Collins, H. M. (1985), *Changing Order*. London: Sage.
- Downes, S. M. (1992), "The Importance of Models in Theorizing: A Deflationary Se-

- mantic Approach", in M. Forbes and D. Hull (eds.), *Proceedings of the Philosophy of Science Association*, vol. 1. East Lansing: Philosophy of Science Association. pp. 142–153.
- Fleck, L. (1979), *The Genesis and Development of a Scientific Fact*. Chicago: University of Chicago Press.
- Forman, P. (1971), "Weimar Culture, Causality, and Quantum Theory, 1917–1927: Adaptation by German Physicists and Mathematicians to a Hostile Intellectual Environment", in R. McCormach (ed.), *Historical Studies in the Physical Sciences*, vol. 3. Philadelphia: University of Pennsylvania Press, pp. 1–115.
- Franklin, A. (1986), *The Neglect of Experiment*. Cambridge, England: Cambridge University Press.
- Fuller, S. W. (1989), *Philosophy of Science and its Discontents*. Boulder: Westview Press.
- Galison, P. (1987), *How Experiments End*. Chicago: University of Chicago Press.
- Giere, R. (1988), *Explaining Science: A Cognitive Approach*. Chicago: University of Chicago Press.
- Glymour, C. (1981), *Theory and Evidence*. Princeton: Princeton University Press.
- Gooding, D.; Pinch, T. and Schaffer, S. (eds.), (1990), *The Uses of Experiment*. Cambridge, England: Cambridge University Press.
- Hacking, I. (1983), *Representing and Intervening*. Cambridge, England: Cambridge University Press.
- Hull, D. (1988a), "A Function for Actual Examples in Philosophy of Science", in M. Ruse (ed.), *What the Philosophy of Biology Is*. Dordrecht: Kluwer, pp. 309–321.
- . (1988b), *Science as Process*. Chicago: University of Chicago Press.
- Kitcher, P. (1990), "The Division of Cognitive Labor", *Journal of Philosophy* 87: 5–22.
- Knorr-Cetina, K. (1981), *The Manufacture of Knowledge*. Oxford: Pergamon Press.
- Lakatos, I. (1981), "History of Science and its Rational Reconstructions", in I. Hacking (ed.), *Scientific Revolutions*. Oxford: Oxford University Press, pp. 107–127.
- Langley, P.; Bradshaw, G. L.; Simon, H. A. and Zyrgow, J. M. (1987), *Scientific Discovery*. Cambridge, MA: MIT Press.
- Latour, B. (1987), *Science in Action*. Cambridge, MA: Harvard University Press.
- Latour, B. and Woolgar, S. (1986), *Laboratory Life*. 2d ed. Princeton: Princeton University Press.
- Laudan, L. (1977), *Progress and its Problems*. Berkeley and Los Angeles: University of California Press.
- Longino, H. (1990), *Science as Social Knowledge*. Princeton: Princeton University Press.
- Mulkay, M. (1976), "Norms and Ideology", *Social Science Information* 15: 637–656.
- Pickering, A. (1984), *Constructing Quarks: A Sociological History of Particle Physics*. Chicago: University of Chicago Press.
- Putnam, H. (1975), *Mind, Language, and Reality*. Cambridge, England: Cambridge University Press.
- Rouse, J. (1987), *Knowledge and Power*. Ithaca: Cornell University Press.
- Shapin, S. and Schaffer, S. (1985), *Leviathan and the Air-pump*. Princeton: Princeton University Press.
- Thagard, P. R. (1988), *A Computational Philosophy of Science*. Cambridge, MA: MIT Press.
- . (1989a), "Explanatory Coherence", *Behavioral and Brain Sciences* 12: 435–502.
- . (1989b), "Scientific Cognition: Hot or Cold?" in S. Fuller, M. DeMey, T. Shinn, and S. Woolgar (eds.), *The Cognitive Turn: Sociological and Psychological Perspectives on Science*. Dordrecht: Kluwer, pp. 71–82.
- van Fraassen, B. (1980), *The Scientific Image*. Oxford: Oxford University Press.
- Wilkes, K. (1988), *Real People*. Oxford: Oxford University Press.
- Woolgar, S. (1988), *Science: The Very Idea*. London: Tavistock.