12. THE CONSTITUTIVELY SOCIAL CHARACTER OF EXPERTISE

STEVE FULLER

1. A BRIEF HISTORY OF EXPERTISE

"Expert," a contraction of the participle "experienced," first appeared as a noun in French at the start of the Third Republic (about 1870). The general idea was much the same as today, namely, someone whose specialized training enables him [sic!] to speak authoritatively on some matter. However, the original context of use was quite specific. The first experts were called as witnesses in trials to detect handwriting forgeries. These people were experienced in discriminating scripts that appeared indistinguishable to the ordinary observer. Thus, the etymological root of "expertise" in "experience" was carried over as the semantically heightened way in which the expert experienced the relevant environment. In contemporary parlance, the tasks that originally required the service of an expert principally involved "pattern recognition," except that the patterns recognized by the expert were identified in terms of an implicit explanatory framework, one typically fraught with value connotations, as in the case of identifying a script as a "forgery."

When evaluating the likelihood that a script was forged, experts were not expected to publicly exhibit their reasoning. They were not casuists who weighed the relative probability that various general principles applied to the case. Rather, it was on the basis of an expert's previous experience of having successfully identified forgeries that his judgment was now trusted. This is not to say that no one could contest expert judgment, but he would have to be another expert, a colleague. If no colleague came forward to testify against an expert's judgment, then the judgment would stand. The climate of collegiality that harbored the mystique of the expert led journalists of the Third Republic to distinguish experts from the "lay" public, thereby conjuring up a clerical image redolent of the secular religion that Auguste Comte's more zealous followers had been promoting under the rubric of Positivism (Williams, 1983: 129).

Moreover, experts were contrasted not only with the lay public but also with *intellectuals.* This point is important for understanding the source of what might be called the *epistemic power* of expertise. An intellectual takes the entire world as fair game for his judgments, but at the same time he opens himself to scrutiny from all quarters. Indeed, the intellectual's natural habitat is controversy, and often he seems to spend more time on defending and attacking po-

sitions than on developing and applying them. In contrast, the expert's judg-ments are restricted to his area of training. The credibility of those judgments are measured in terms of the freedom from contravention that his colleagues accord him. The mystique of expertise is created by the impression that an ex-pert's colleagues are sufficiently scrupulous that, were it necessary, they would be able and inclined to redress any misuse or abuse of their expertise. The fact that they do not means that the expert must be doing something right.

Collegiality enables experts to exert what both Plato and Machiavelli would have recognized as an ideal form of power. In its ideal form, power thrives on a counterfactual that never needs to be realized—in less charitable terms, a persuasive bluff. If a prince's enemies believe that the prince could squash any uprising, the enemies will lie low, and the prince will seem invincible; however, if the enemies challenge the prince, and the prince defeats them only with great difficulty, then the air of princely invincibility will disappear, and the prince will need to prepare for redoubled efforts by his enemies in the future. Thus, ideal power is brought down to the level of brute force (Botwinick, 1990: 133–80). Trial attorneys continue to exploit this point whenever they try to un-dermine the very possibility of expertise in a field by pitting particular experts against one another. The lawyers do not expect a definitive judgment to emerge from the crossfire; rather, they expect to show that no such judgment can be rendered.

Philosophers have traditionally shared the attorney's desire to dissipate the power of expertise, but without in the process undermining the possibility of knowledge. Be it embodied in machine or human, philosophers have looked askance at the epistemological status of expertise. For Karl Popper, the exper-tise conferred on those trained in the special sciences serves to strategically contain critical inquiry, as the evaluation of a knowledge claim depends on the credentials of the claimant. The sociology of knowledge perennially incurs the wrath of philosophers because it seems to condone this tendency, which cul-minates in the formation of scientific guilds, or disciplines. These, in turn, di-vert inquiry away from questions of fundamental principles that go to the heart of disciplinary identity. Disciplines proliferate explanatory frameworks (jargons, if you will), while "science," in the philosophically honorific sense, unifies such frameworks. Not surprisingly, then, Popper regarded Kuhnian "normal science" as a "danger" to the advance of knowledge, if that advance is measured in terms of explanatory comprehensiveness, or the Newtonian virtue of explaining the most by the least (Popper, 1970).

Cognitive scientists display a similar disdain when they contrast the "knowl-edge-based" or "domain-specific" character of expert systems to more general purpose problem-solving machines that utilize principles that cut across do-mains. What is often called the "orthodoxy" in cognitive science holds that an adequate theory of cognition is not to be found simply by compounding a

number of distinct expert systems, just as an adequate theory of knowledge requires more than simply articulating the research conventions of all the special sciences (Haugeland, 1984). The hope, in both cases, is that whatever principles govern the special cases are not unique to them.

2. SOME SENSES OF THE SOCIAL

To someone whose work is primarily in the design of expert systems, it may seem odd to juxtapose philosophical suspicions about human and machine expertise in the way I have. Perhaps this juxtaposition points to a residual positivist hankering that cognitive science and epistemology (or philosophy of science) share for the "unity of science." Without casting doubts on this diagnosis, I nevertheless want to stress some philosophically disturbing aspects of expertise that have yet to fully grip theorists in this area. As a first pass, these aspects turn on the following truism: *Expertise is a constitutively social phenomenon.* Indeed, as a card-carrying "social epistemologist" (Fuller, 1988, 1993a, 1993b), I am committed to the position that expertise can be exhaustively analyzed as a social phenomenon. But before too many eyebrows are raised, let me begin by listing—in descending order of intuitiveness—four distinct senses in which expertise is constitutively social:

1) The skills associated with an expertise are the product of specialized training. Expertise cannot be picked up casually or as the byproduct of some other form of learning.
2) Both experts and the lay public recognize that expertise is relevant only on certain occasions. No expertise carries universal applicability.
3) The disposition of expertise is dependent on the collegial patterns of the relevant experts. Protracted internecine disputes over fundamentals typically erode expertise.
4) The cognitive significance of an expertise is affected by the availability of expert training and judgment, relative to the need for the expertise. Too many experts or too little need typically devalue the expertise in question.

So far, I have concentrated on sense (3), given its salience in the historical development of the expert as a social role distinct from that of the layperson, the intellectual, and, as we have just seen, even the scientist. However, all four senses echo the twin themes of *boundedness* and *compartmentalization* as essential to the definition of expertise. The cognitive science literature offers several ways of articulating these themes: Simon's heuristics, Minsky's frames, Schank's scripts, Fodor's modules, and, most abstractly, Pylyshyn's cognitive

impenetrability. Of course, these terms do not divide up the mind's labor in quite the same way. A similar proviso would have to be attached to sociological markers of expertise, such as "indexicality" and "functional differentiation" (cf. Cicourel, 1968; Knorr-Cetina, 1981). But for our purposes, the most striking comparison may be between, so to speak, the cognitive and political impenetrability of expertise, the so-called autonomy of the professions (cf. Abbott, 1988).

The analogy I wish to draw here is not particularly difficult to grasp, but doing so may alert us to the conceptual baggage that is unwittingly imported in the images we use to characterize expertise. The profession that has most jealously guarded its autonomy—scientists—has often struck a bargain to ensure that professionally produced knowledge remains both cognitively and politically impenetrable. From the charter of the Royal Society in seventeenth-century Britain to the guild right of *Lehrfreiheit* (freedom of inquiry) that German university professors enjoyed under Bismarck, the following two conditions have been met (cf. Proctor, 1991):

(A) The state agrees not to interfere in the internal governance of the profession, on the condition that the profession does not interfere in the governance of the state.

(B) The state agrees to protect the profession from others who might want to interfere with its governance (e.g. other professional, political, or business interests), on the condition that the state is given the first opportunity to appropriate the knowledge produced by the profession, where appropriation includes the right to prevent others from subsequently appropriating the knowledge (e.g. for reasons of national security).

Most theories of the mind in cognitive science are fairly explicit in distinguishing an executive central processor or general problem-solver from domain-specific modules that function in relative autonomy from this unit. Sometimes (as in Simon, 1981; Minsky, 1986) the political imagery of governance is quite strong. However, it is not just any old image of governance, but one that is characteristic of twentieth-century thinking about the state, namely, *democratic pluralism* (cf. Held, 1987: 186–220). The pluralist portrays the state as mediating competing factions in a large democratic society by enabling the factions to flourish without letting any of them override each other or the national interest. The learning process of big democracies consists of these factions settling into interest groups and, ultimately (and ideally), professionally governed associations whose interaction is founded on recognition and respect for each other's work as essential to the business of society. Gradually, then, the state's role as mediator recedes to that of chairman of a corporate board of directors. The only point I wish to make here is that this is not the only, nor necessarily

the most desirable, image of democratic governance. It might repay the effort for cognitive scientists to examine alternative forms of governance as a source of new images for arranging the parts of the mind.

3. THERE MAY BE LESS TO EXPERTISE THAN MEETS THE EYE

To explore the social character of expertise implied in senses (1) and (2), consider what might be called a *behaviorist* and a *cognitivist* account of how expertise develops (cf. De Mey 1982: 216):

> *Behaviorist*: Expertise is shaped from repeated encounters with relevant environments, such that increased exposure smooths out the rough edges in the expert's practice until it stabilizes at a normatively acceptable standard, which can then be applied, "off the shelf" as it were, in subsequent encounters.
>
> *Cognitivist*: Expertise consists of a core set of skills that are elaborated in a variety of environments, most of which are unforeseen. These elaborations are stored and themselves elaborated upon in subsequent encounters, all of which serves to confer on expertise a sophistication that is evident to the observer but difficult to articulate.

Thus, whereas the behaviorist sees expertise becoming more *stereotyped* in practice, the cognitivist sees it becoming more *nuanced*. In the ongoing dispute between cognitivists and behaviorists, it is often asserted that the cognitivist won this round. But before the final verdict is delivered, I wish to offer support for the behaviorist by way of revealing the hidden social hand of expertise.

To hear cognitivists (not to mention phenomenologists) emote about the "nuanced" and "craftlike" character of expertise, one would think that an act of professional judgment was tantamount to a magic trick, one in which the audience has attended a little too closely to the magician's gestures and not enough to the circumstances under which the illusion transpires. A professional magician does not perform tricks on demand, say, by adapting his performance to play off the specific gullibilities of his audience. Of course, the magician adapts somewhat to his audience, but before he even agrees to display his expertise, the stage must be set just right, and the audience must already be in the right frame of mind to be receptive to the "magic moment." A magician who is too indiscriminate in his eagerness to please is bound to look bad very quickly. An instructive case in point is the odd magician who submits his performances to the strictures of the experimental method (Collins & Pinch, 1982: on psychokinesis).

It is worth noting that, in these scientific performances, the magician fares no worse than the expert witnesses in medicine and psychiatry whose testimony is routinely heard—and believed—in court (Faust, 1985). The reliability and validity of expert judgment in experimental settings are low all around (Arkes & Hammond, 1986). Perhaps the most celebrated historical case of the hubris of overextended expertise was the rise and fall of the Sophists in 5th Century B.C. Athens. After some fair success as teachers of rhetoric and confidants of the ruling class, the Sophists began to offer their services in competition with more established forms of knowledge in virtually every domain. According to a recent historical study (De Romilly, 1992), the Sophists were perceived as opportunistic colonizers of conventional practices who failed to cultivate the trust required for their own practices to succeed. Consequently, the Sophists were soon ridiculed by the people we now regard as the founders of classical philosophy and drama—people with a greater surface respect for tradition than the Sophists had.

Proponents of the operant conditioning paradigm in behaviorism would recognize what is going on here. Successful experts realize that the secret to their success lies in noting the sorts of situations in which clients come away most satisfied with expert judgment (and hence reward the expert appropriately) and then maximizing the opportunities for situations of that sort to arise in the future. In short, the smart expert controls the contingencies by which her behavior is reinforced (cf. Haddock & Houts, 1992). She does not easily submit to an experimental test of her expertise, over which she exerts little control. The strategy for achieving the relevant sense of control involves several tactics:

(i) pre-selecting clients before officially engaging in treatment, typically through an interview process, an important aspect of which is to gain the confidence of the prospective client;

(ii) learning to refuse certain clients, no matter how lucrative their business might be, if it appears as though they will resist treatment or, in some way, be likely to make the expert look bad in the long run;

(iii) persuading the prospective client that her avowed problem is really one that the expert has seen many times before; often this is done by first showing the client that she has not conceptualized her problem correctly; once the problem receives its proper formulation—which neatly coincides with what the expert is in the best position to treat—then treatment can begin;

(iv) obscuring any discussion of the exact method of treatment by recasting the client's problem in the jargon of the expertise, which will presumably lead the client to infer that anything that must be described in such esoteric ways must be subject to a treatment that would be equally difficult to convey.

The most obvious result of these four tactics is to shape the prospective client's behavior so that her problem falls into one of the stereotyped patterns with which the expert is familiar. A less obvious but equally important result is that the expert can now exert *spin control* over how the client understands her situation after the treatment. If the client's problem is solved, the expert can confidently claim credit for the solution. No suspicions of "spontaneous remission" are likely to be raised, especially if the client has paid dearly for treatment. But perhaps more important, if the client's problem remains unsolved after the treatment, then the expert can claim, with only slightly less confidence, that other unforeseen factors intervened, including the client's own recalcitrance. Whatever happens, therefore, reinforces the expert's competence and integrity. On the basis of these considerations, I conclude that the key to understanding the distinct character of expertise may lie less in its associated skills than in the discretionary control that the expert has in deploying those skills. (Philosophers of science and cognitive scientists who are uncomfortable with brute talk of "control" may substitute this piece of genteel intellectualism: Experts have heightened metaknowledge of the *ceteris paribus* clause, or relevance conditions, for applying their expertise.) Thus, if the above pre-selection strategy fails to mold the client into shape, the expert can then tell the client that the problem lies outside her expertise, which will probably cause the client to believe that her problem remains unsolved, not because the expert was incompetent, but because the client has yet to locate the right expert, which may itself reflect the client's own failure to understand the nature of her problem. The element of trust crucial for the maintenance of expertise may be seen in the willingness with which the client holds herself responsible for an expert's inability to come to grips with her problem (cf. Gambetta, 1988). From what I have said so far, it may seem that my thinking about the social dimension of expertise has been strongly based on psychiatric encounters, which have often been subject to unflattering depictions as confidence games. However, the same observations apply, perhaps with greater import, to experts operating in the arena of public policy, especially those trained in medicine, engineering, or economics. For, the biggest single problem facing the future of democracy may be *cognitive authoritarianism*, the tendency to cede an ever larger share of the realm of participatory politics to expert rule (Fuller, 1988: 277–88). The conversion is accomplished as government officials become convinced that the public has ill-formed conceptions of its own needs, needs that are best shaped and addressed by the relevant experts. When government fails to act speedily on this conversion, and hence does not clear the political environment to enable the expert's stereotyped knowledge to take effect, the expert will often appear as a moral censor, appealing to his special knowledge—which would supposedly be efficacious if politicians secured the relevant background conditions—as a norm against which the state of society is criticized.

But can all this talk about the strategically discretionary character of expertise be applied to computerized expert systems? I do not see why not. Suppose a knowledge engineer has been asked to design an expert system that will offer advice on playing the stock market. After some time, the knowledge engineer returns with a product that she announces was constructed from in-depth interviews with four of the best stock analysts on Wall Street, making sure that respected spokespersons for all the relevant market perspectives—fundamentalists, chartists, insiders, and traders—were canvassed (cf. Smith, 1981). The effect of this pedigree on the client will be similar to that of the diploma and license that hang on the wall of the human expert's office and invariably engage the client's peripheral vision during a consultation. If the knowledge engineer designs the interface with protocols that make interaction between client and expert appear stilted, then the client will probably interpret that to mean that the expert is concerned with getting to the heart of the client's problem without dragging in superfluous information. Likewise, if the expert seems to give the client advice that causes her to lose money in the market, then the client may wonder whether a human expert could have really done any better, or that perhaps she did not input all the information that was relevant for the expert system to provide better advice. Moreover, the client's inclination to assume responsibility for the bad advice increases with the amount of money that she had to originally spend to purchase the system. That an AI pioneer, Joseph Weizenbaum (1976), should appeal to *moral,* rather than technical, grounds for restricting the use of expert systems reflects the propensity of clients to invest the same level of trust in computers as in human beings.

4. Global Constructivism and the Political Economy of Expertise

The senses (1)–(3) in which expertise is constitutively social are consistent with a *constructivist* sociological orientation (cf. Knorr-Cetina, 1981). Constructivists typically minimize the attribution of intrinsic properties to cognitive agents. Instead, they unpack such so-called intrinsic properties into relational ones, in which the relata are two mutually interpretive agents who jointly negotiate who will be credited with which properties. As we have seen, an adept expert can shift the burden of responsibility onto the client for the unpleasant consequences of following expert advice. But for the constructivist, there is no "fact of the matter" about whether the expert's incompetence or the client's recalcitrance is to blame, until the transaction has actually taken place. In discussions of cognitive science that acknowledge that society is more than a metaphor for the mind, constructivism is often presented as *the* sociological

perspective. Empirically speaking, the dominance of constructivism cannot be denied, but I will ultimately appeal to sense (4) of expertise's sociality in order to introduce a different, and more comprehensive, sociological perspective.

Constructivism is a more heterogeneous doctrine than it may first appear. Weizenbaum and others whose opposition to the cognitive authority of either computer or human experts is mainly moral presuppose that such experts can, indeed, exercise all manner of authority, if they are not limited by convention. However, a milder species of constructivism is represented by Daniel Dennett's (e.g., 1987) instrumentalist approach to intentionality, which makes one's cognitive status dependent on another's interpretive stance. Dennett's constructivism is "asymmetrical," in that he does not grant the computer the same degree of agency as most human beings in constructing their respective identities. By contrast, a committed sense of symmetry is the hallmark of the more radical constructivists, the ethnographers who can be increasingly found on the sites where knowledge engineering occurs (cf. Greenbaum & Kyng, 1991). Here knowledge engineers, human and computer experts, clients, and other people and artifacts are portrayed as engaged in a mutually reinforcing cooperative venture. I will now address the limitations of this version of constructivism.

Drawing on the work of cultural anthropologists, especially Clifford Geertz (1983), ethnographic constructivism makes much of the "local" character of expert knowledge, which is brought out very clearly in the design of expert systems. The idiosyncrasy of locales is brought out in the knowledge engineer's interviews with both client and expert, followed by the process of adapting the expert system to the client's specific needs and abilities. These ethnographic accounts are meant to stand in striking contrast to the accounts typically given by the designers of AI systems in basic research settings, who often influence the way in which the applied researchers, the knowledge engineers, conceptualize their activities. Specifically, workers in AI tend to attribute properties to their programs (typically once embodied in a machine) that the ethnographers would prefer to regard as "boundary objects" in terms of which the identities (or cognitive capacities) of a variety of agents are negotiated. For example, the degree of satisfaction that the client gets from using an expert system has implications for how much of human expertise was successfully transferred to the computer program. While I wholeheartedly endorse this reinterpretation of expertise as a corrective to the accounts given by AI researchers, it nevertheless shares the fatal flaw of its opponents. Ironically, the flaw is the tendency to universalize from a single case. Let me explain.

It is one thing to say that all knowledge is local. That is rather boring. It is another to say that all locales are different from one another. That is more exciting. Ethnographers infer the exciting proposition from the boring one all the time. But what licenses this inference? Certainly, there have not been enough ethnographies of knowledge engineering to license the exciting propo

sition as an empirical generalization. In fact, the few ethnographies available do not depict locales so radically different from one another. My point here is not to decide the issue by philosophical argument but to observe that the ethnographic appeal to locality presupposes a *conceptual cartography* whereby one imagines that the spatio-temporal distance between locales represents a conceptual distance as well. In that sense, "the local" presupposes "the global," an image of the whole. Admittedly, sometimes this image turns out to be right; but other times it doesn't. However, that is a matter for empirical inquiry, and it is not clear that the ethnographic brand of constructivism encourages the appropriate sort of empirical inquiry. For, to learn about the global properties of knowledge engineering, one needs to discover the pattern by which expertise is *distributed* across a representative sample of locales and the *aggregated* consequences of such a distribution for the knowledge system as a whole. Here I start to speak the language of *political economy*, and to signal the quest for statistical correlations among variables that are hypothesized to be salient for understanding how expertise works.

What sorts of people are the producers and consumers of expert knowledge? Provided with a serviceable answer, we can study the distribution of expertise by focusing on a cognitively relevant locus of scarcity: A client has only so much time and money to spend consulting an expert, be it human or computer. Under what circumstances does the client feel she has gotten what she has paid for, and, when she does not feel that way, how is the blame apportioned: Who receives the lion's share of incompetence—the client or the expert? A key to understanding the distribution of expertise is to see how each side tries to convert its own sense of frustration into a perception of the other's liabilities. It would be fair to suppose that expert computers today receive far more attributions of incompetence than expert humans. A constructivist would diagnose this difference in terms of the client's lack of time, imagination, or interest in interpreting the computer as performing intelligently—perhaps because the client feels either that she has better things to do at this point, and the computer is in no position to prevent her from doing them, or that she would have to end up interpreting the computer as doing something other than she would have wanted (cf. Fuller, 1993b: 179–85). However, as people become more accustomed to dealing with expert computers, this difference in attribution is likely to disappear. But before concluding that we are projecting a future in which experts of all sorts are engaged in mutually satisfying relationships with their clients, we need to consider the aggregated consequences of people increasingly turning to computers for advice.

As we have seen, the history of expertise teaches that the expert is not a universalizable social role. There are no "experts" in areas that are regarded as commonsense or part of general education or easy to acquire without specialized training. Consequently, knowledge engineers are in the curious position

of potentially destroying expertise as they diligently codify it and make it available to more people in user-friendly packages. While this consequence is bound to elude any on-site description of the knowledge engineer's work, it is nevertheless felt by professional associations that believe that knowledge engineers are indirectly deskilling their members. For, even as the human expert retains discretionary control over when, where, and how she uses her expertise, she may be losing discretionary control at the meta-level, namely, over who—or what—else counts as an expert in her field. Librarians have so far been most vocal in their concerns (cf. Pfaffenberger, 1990), but attempts by many doctors and lawyers to limit the scope of the interviews they give to knowledge engineers reflect similar worries.

At first glance, it may seem that the proliferation of expert systems is the ideal vehicle for democratizing expertise, as it would seem to put expert knowledge within the reach of more people. Just because the knowledge engineer can extract elements of expertise from her interviews with experts, it does not follow that the expertise remains intact once it is programmed into a computer for a client. After all, if expertise is indeed constitutively social, then altering the context in which expert knowledge is deployed should alter the character of the knowledge itself. Such change may be witnessed in the course of designing the interface that enables the client to interact with the expert system. Here the tendency has been to "go ergonomic" by designing interfaces that require the client to change his ordinary patterns of thought and behavior as little as possible (Downes, 1987). Less charitably put, the ergonomic approach reinforces the client's cognitive biases and thereby minimizes the learning experience that he might derive from engaging with the expertise as a form of knowledge. A potentially "dialectical" exchange is thus rendered merely "instrumental" (cf. Adorno & Horkheimer, 1972). The result is a spurious sense of autonomy, whereby the client's powers appear to be extended only because the environment in which he acts has been changed to his advantage (Fuller, 1986).

Thus, while the expert humans may lose some of their power as their clients increasingly rely on computerized systems, the clients themselves may not, in turn, become epistemically empowered. *Experts are deskilled without clients being reskilled.* Where has the original power of expertise gone? That power would seem to have dissipated somewhere in the knowledge engineering process, specifically when expertise was converted into a *tool* that exerted few of its own demands on the user (cf. Fields, 1987).

The utopian vision of democratized expertise is foiled by the simple fact that expertise, and perhaps knowledge more generally, is what economists call a *positional good* (Hirsch, 1977). A positional good is one whose value is directly tied to others not having it. Economists have generally refused to count knowledge as a positional good, preferring instead the classical philosophical position that knowledge is an "ethereal good," one whose value is not determined

by the laws of supply and demand (Fuller, 1992). However, professional associations realize all too well the positional character of expertise.

The existence of positional goods is the dark secret of the welfare state. According to welfare economics, capitalism can avert a Marxist revolution because lingering inequalities of wealth will be resolved once a level of productivity is reached that enables everyone to be supported at a minimally acceptable standard. At most, stabilizing this situation will require a modest redistribution of income through progressive taxation. Overlooked in this scenario is that, as more goods are made more generally available, the perceived value of the goods may decline as they no longer serve to discriminate people in socially relevant ways (Bourdieu, 1984). Knowledge-intensive goods display such positionality effects. Higher education is perhaps the most obvious case in point: As it becomes easier for people to complete college, more postgraduate degrees are needed to acquire the same credentials. Should it become impossible either to stop the production of degree-holders or to set up additional barriers in the credentialing process, higher education will then no longer be seen as imparting an especially valued form of knowledge. Instead, it will take the place of bare literacy and the high school diploma as the minimum threshold for entry into the job market. Can anything be done reverse such positionality effects, or is the value of knowledge-intensive goods doomed to continual deflation?

The political scientist Yaron Ezrahi (1990) has argued that the scientific enterprise has come to consume so many resources and to produce so many questionable consequences that we may be reaching the end of the period (which began with the Enlightenment) when knowledge is presumed to be a public good. Ezrahi envisages that scientific forms of knowledge will gradually acquire the social character of artforms: their support will be privatized and their products customized to client tastes, which are presumed *not* to be universalizable. The expansion of intellectual property law to cover more instances of "basic research" suggests that Ezrahi's prognosis is already taking shape (Fuller, 1991). Given their own interest in customizing expertise for user demand, knowledge engineers clearly contribute to this overall trend toward privatization. Indeed, human experts may soon find the need to seek legal protection for their expertise, if only to earn royalties from the expert systems designed on the basis of it. In that way, knowledge engineers would not benefit too much from the "Japan Effect" of learning how to manufacture expertise more efficiently than the original experts themselves (cf. Weil & Snapper, 1990). This form of legal protection would, in turn, require a new category of intellectual property beyond the usual three of patent, copyright, and trademark.

But even if expertise were to become entirely market-driven, the skills surrounding the expertise would still attract human practitioners for the same reasons as art continues to do. The skills would be detached from the fame,

fortune, or power that had been previously tied to them. Most of the perverse consequences of positional goods rest on such coupling (Crouch, 1983). For example, higher education is populated by a few people who are interested in the education process itself and many more who view it as a credentialing process, the surest route to a job. Decoupling those two groups would presumably help restore the integrity of higher education.

Knowledge engineers have a crucial role to play in the future disposition of expertise and knowledge more generally. Customized expert systems will hasten the demise of expertise and turn Ezrahi's image into a reality. However, the proliferation of such systems may also limit the client's potential for cognitive growth. A page from the history of manufacturing may prove instructive in resolving this dilemma.

Once the demand for manufactured products grew to a critical level, customization yielded to mass production (Beniger, 1986: 291–343). This transition was accompanied by the design of quality control standards for the mass produced goods. In the process of defining the minimum level of acceptability for a particular good, manufacturers effectively forced potential customers to adapt their behavior to the set dimensions of the good. Typically, these adaptations were dictated by the manufacturer's desire to cut costs, but knowledge engineers could collectively set guidelines for the design of expert systems, the successful use of which required clients to expand their cognitive repertoire. The sort of behavioral changes I envision here may be quite subtle. For example, an online library search system may discourage disciplinary provincialism by requiring the client to initiate searches by using protocols that are tied less to the jargon of specific disciplines and more to the exact topic or problem that client wishes to tackle. The system's database would, in turn, draw on the literatures of several disciplines so as not simply to confirm the course of inquiry that the client would be naturally inclined to follow (Cronin & Davenport, 1988: 316–27).

5. But Is Expertise Really Knowledge?

It remains an open question whether the epistemic power of science is tied more to its sheer practice or to its status as a positional good. Those keen on retaining the ethereal quality of knowledge may no doubt want to make a strong distinction between "genuine knowledge" and "mere expertise" (cf. Ford & Agnew, 1992). They will object to my apparent conflation of these two concepts. Whereas expertise may ultimately reduce to matters of status and trust, the objector may argue, the test of knowledge is precisely that it does not lose its force as its availability increases. In response, let me grant the objector's distinction as the basis for an empirical hypothesis. If there are indeed types of

"information" or "skills" whose power does not diminish with their increased availability, then I will gladly call them "genuine knowledge." However, my guess is that there are limits to the optimal distribution of these cognitive products. Consider the following homely observations:

(a) Everyone in a town may know the location of a particular store and the time it opens for business. However, if many of these people decide to act on that knowledge at roughly the same time to purchase the same goods, then a larger percentage of them will probably return home empty-handed than if fewer of them knew about the store in the first place. *Here knowledge lacked efficacy because the knowledgeable got in each other's way.* (A more realistic version of this situation is one in which everyone decides to take the same expert's advice on which stock to purchase.)

(b) It is often assumed that information freely exchanged among a large network of peers breeds the sort of critical inquiry that is necessary for genuine epistemic progress: the larger and freer the network, the more critical the inquiry. Unfortunately, this assumption presumes, contrary to fact, that inquirers have an inexhaustible ability and inclination to attend to each other's work. Yet, by the time the network of inquiry attains the dimensions of "Big Science," inquirers become more concerned with finding allies than opponents, and hence are likely to simply ignore work that cannot be immediately used for one's own purposes (cf. Fuller, 1994). *Here knowledge lacked efficacy because more of it was available than could be assimilated.*

Knowledge was undermined in (a) because too many people possessed the same information, whereas in (b) it was because each person possessed too much information. In neither case did these skewed distributions actually convert a truth into a falsehood, but from a pragmatic standpoint, they might as well have. In other words, attention to the socially distributed character of knowledge may help explain the intuitions that have traditionally led philosophers to posit a conception of knowledge that "transcends" the constitutively social character of expertise.

REFERENCES

Abbott, A. (1988). *The System of Professions.* Chicago: Chicago.

Adorno, T. & Horkheimer, M. (1972). *The Dialectic of Enlightenment.* New York: Continuum.

Arkes, H. & Hammond, K., Eds. (1986). *Judgment and Decision Making.* Cambridge: Cambridge.

Beniger, J. (1986). *The Control Revolution.* Cambridge: Harvard.

Botwinick, A. (1990). *Skepticism and Political Participation.* Philadelphia: Temple.

Bourdieu, P. (1984). *Distinction.* Cambridge: Harvard.

Cicourel, A. (1968). *Cognitive Sociology.* Harmondsworth: Penguin.

Collins, H. & Pinch, T. (1982). *Frames of Meaning.* London: Routledge & Kegan Paul.

Cronin, B. & Davenport, E. (1988). *Post-Professionalism: Transforming the Information Heartland.* London: Taylor Graham.

Crouch, C. (1983). Market failure. In A. Ellis & K. Kumar (Eds.), *Dilemmas of Liberal Democracies.* London: Tavistock.

De Mey, M. (1982). *The Cognitive Paradigm.* Dordrecht: Reidel.

De Romilly, J. (1992). *The Great Sophists in Periclean Athens.* Oxford: Clarendon.

Dennett, D. (1987). *The Intentional Stance.* Cambridge: MIT.

Downes, S. (1987). A philosophical ethnography of human-computer interaction research. *Social Epistemology,* 1, 27–36.

Ezrahi, Y. (1990). *The Descent of Icarus.* Cambridge: Harvard.

Faust, D. (1985). *The Limits of Scientific Reasoning.* Minneapolis: Minnesota.

Fields, C. (1987). The computer as tool. *Social Epistemology,* 1, 5–26.

Ford, K. & Agnew, N. (1992). Expertise: socially situated, personally constructed, and "reality" relevant. Paper presented at the AAAI Spring Symposium on the Cognitive Aspects of Knowledge Acquisition. Stanford.

Fuller, S. (1986). User-friendliness: friend or foe? *Logos,* 7, 93–98.

Fuller, S. (1988). *Social Epistemology.* Bloomington: Indiana.

Fuller, S. (1991). Studying the proprietary grounds of knowledge. *Journal of Social Behavior and Personality,* 6 (6), 105–28.

Fuller, S. (1992). Knowledge as product and property. In N. Stehr & R. Ericson (Eds.), *The Culture and Power of Knowledge* (pp. 157–90). Berlin: Walter de Gruyter.

Fuller, S. (1993a). *Philosophy of Science and Its Discontents.* 2nd edn. (Orig. 1989) New York: Guilford.

Fuller, S. (1993b). *Philosophy, Rhetoric, and the End of Knowledge: The Coming of Science & Technology Studies.* Madison: University of Wisconsin Press.

Fuller, S. (1994). The social psychology of scientific knowledge: another strong programme. In W. Shadish & S. Fuller (Eds.), *Social Psychology of Science.* New York: Guilford: 162–80.

Gambetta, D., Ed. (1988). *Trust.* Oxford: Blackwell.

Geertz, C. (1983). *Local Knowledge.* New York: Basic.

Greenbaum, J. & Kyng, M., Eds. (1991). *Design at Work.* Hillsdale NJ: Lawrence Erlbaum.

Haddock, K. & Houts, A. (1992). Answers to philosophical and sociological uses of psychologism in science studies. In R. Giere (Ed.), *Cognitive Models of Science* (pp. 367–99). Minneapolis: Minnesota.

Haugeland, J. (1984). *Artificial Intelligence: The Very Idea.* Cambridge: MIT Press.

Held, D. (1987). *Models of Democracy.* Oxford: Polity.

Hirsch, F. (1977). *Social Limits to Growth.* London: Routledge & Kegan Paul.

Knorr-Cetina, K. (1981). *The Manufacture of Knowledge.* Oxford: Pergamon.

Minsky, M. (1986). *The Society of Mind.* New York: Simon & Schuster.

Pfaffenberger, B. (1990). *Democratizing Information.* Boston: G. K. Hall.

Popper, K. (1970). Normal science and its dangers. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the Growth of Knowledge* (pp. 51–58). Cambridge: Cambridge.

Proctor, R. (1991). *Value-Neutral Science?* Cambridge: Harvard.

Simon, H. (1981). *The Sciences of the Artificial.* Cambridge: MIT.

Smith, C. (1981). *The Mind of the Market*. Totowa NJ: Rowman & Littlefield.

Weil, V. & Snapper, J., Eds. (1990). *Owning Scientific and Technical Information*. New Brunswick NJ: Rutgers.

Weizenbaum, J. (1976). *Computer Power and Human Reason*. San Francisco. W. H. Freeman.

Williams, R. (1983). *Keywords*. Oxford: Oxford.