

- ¹⁴See John Holland, Keith Holyoak, Richard Nisbett, and Paul Thagard, *Induction: Processes of Inference, Learning, and Discovery*, Cambridge, Mass.: MIT Press/Bradford Books, 1986; and Thagard, *Computational Philosophy of Science*, *op. cit.*. PI stands for "processes of induction" and is pronounced "pie".
- ¹⁵Kunda, "Motivated Inference", *op. cit.*, pp. 645-646.
- ¹⁶Paul Forman, "Weimar Culture, Causality and Quantum Theory, 1918-1927," *Historical Studies in the Physical Sciences*, 3, 1971-1-115.
- ¹⁷T.S. Kuhn, *The Structure of Scientific Revolutions*, second edition, Chicago: University of Chicago Press, 1970. Barry Barnes, *T. S. Kuhn and Social Science*, New York: Columbia University Press, 1982.
- ¹⁸Paul Thagard, "The Conceptual Structure of the Chemical Revolution," forthcoming in *Philosophy of Science*. Carleton Perrin, "Response to Theoretical Innovation in Science: Patterns of the Chemical Revolution". Paper read at conference on testing theories of scientific change, Blacksburg, Virginia. 1986.
- ¹⁹David Hull, Peter Tessner, and Arthur Diamond, "Planck's Principle," *Science*, 202, 1978, 717-723.
- ²⁰H. W. Menard, *Oceans of Truth*, Princeton: Princeton University Press, 1986; Gregory Nowak and Paul Thagard, "The Conceptual Structure of the Geological Revolution," unpublished manuscript, Princeton University; Paul Thagard and Gregory Nowak, "The Explanatory Coherence of Continental Drift", in *PSA 1988*, vol. 1, East Lansing, MI: Philosophy of Science Association.

TACIT KNOWLEDGE AND THE PROBLEM OF COMPUTER MODELLING COGNITIVE PROCESSES IN SCIENCE

Stephen P. Turner
University of South Florida

Once one has said that there is a "tacit dimension" of some sort to communication or to some cognitive activity, it seems like an appropriate step to make it a topic for research. The difficulty with this suggestion is that it assumes that in fact we have some sort of unproblematic mode of access to our implicit understandings, to the "tacit dimension." To be sure, sometimes portions of this "dimension" are revealed to us, as when we find that others *don't* share some procedure we follow "naturally" or habitually, or when we find ourselves in circumstances where these procedures fail, and we are forced to, and can, find replacements for them. But these are perhaps exceptional cases, revealed for us through contingent and possibly rare circumstances. And if this is so, it follows that a *project* of revealing the tacit dimension would depend on some premises that would give one some assurances that one's techniques indeed reveal the necessary parts of it. It is one thing to have a method of turning over rocks to see if there are salamanders under them; it is another to form a project of counting all the salamanders in a given area, for one doesn't know that they are all assessable by the known means.

In what follows I propose to bring out certain methodological properties of projects of modelling the tacit realm that bear on the kinds of modelling done in connection with scientific cognition by computer as well as by ethnomethodological sociologists, both of whom must make some claims about the tacit in the course of their efforts to model cognition. The same issues, I will suggest, bear on the project of a cognitive psychology of science as well.

Tacitness has had a small but persistent role in the writings of critics of computer modelling of cognition, a role smaller than that of intuition, a concept with which it is often confused. In what follows the terms will be distinguished: tacitness will refer strictly to the reproducible expectations, assumptions, presuppositions, cognitive skills and what-not that I have elsewhere called "sub-beliefs." By reproducible I mean simply socially transmitted, as the training of a clinician is, as distinct from wholly private, which intuition might be thought to be. The role of the concept of tacitness in the critical literature on computer modelling or human reasoning has been to identify a residual category of things allegedly "essential" to "real thinking" or "real discovery" that are unmodelled and perhaps unmodellable by computer programs. Weizenbaum, for example, uses Polanyi¹, Chomsky² and Chief Justice Holmes³ as a source of examples of tacitness, and, in addition to his remarks on intuition, stresses the idea that machines cannot be socialized in the same way as humans, if only because certain primal experiences cannot be shared with machines as they are with other humans.

The argument in this paper in no way depends on this thesis about ultimate or essential unmodellability, and indeed because it is in large part about methodological troubles with the concept of tacitness that extend all the way from the statisticians' concept of "belief structures" to the intellectual historians' concept of "structures of consciousness" in the historical past, it potentially undermines uses such as

Weizenbaum's, especially when these involve potentially troubled assertions about the tacit realm, such as claims about their essential character.

The relevance of problems of tacitness to discussions of scientific discovery is evident from discussions of the more familiar "search procedures" of statistics, such as the kinds of curve-fitting done using least squares. The difficulties are dramatized in the following suggestion made by the statistician David Freedman.

Suppose a few of us were transported back in time to the year 1600, and were invited by the Emperor Rudolph II to set up an Imperial Department of Statistics in the court at Prague. Despairing of those circular orbits, Kepler enrolls in our department. We teach him the general linear model, least squares, dummy variables, everything. He goes back to work, fits the best circular orbit for Mars by least squares, puts in a dummy variable for the exceptional observation—and publishes. And that's the end, right there in Prague at the beginning of the 17th century.⁴

If this sounds a bit fantastic, there is another case closer to home. In the dispute between biometers and Mendelians at the beginning of this century, the inventors of these same regression methods did in fact settle for a correlational view of heredity, at least for a time. It is said that Pearson, who lived long after the dispute, never accepted the Mendelian hypothesis. He thus fulfilled precisely the fears which earlier critics of statistics, such as Mill and Comte, had expressed, the fear that the investigations would settle for statistical results that were a conceptual dead end. Glymour has responded to Freeman's suggestion by arguing that Kepler would have been saved from a dead end in much the same way as the experimental tradition saved genetics: the separate problem-tradition of Newtonian physics, and especially Newton's derivation of elliptical orbits from his law of gravitation,⁵ would have forced these scientists to push on.

Glymour's use of the term "problem-tradition" should suffice to make a small point: the search procedures themselves do not suffice to explain why these curve-fitting efforts were not dead-ends. One need not be a believer in any sort of mystical doctrine of tacitness to claim that to model discoveries of this kind, it is not enough to describe a search procedure based on the explicit assumptions initially held by the researchers. Some account needs to be taken of the ways in which these assumptions were or could be revised. And because they are typically revised by importations from other traditions rather than from mechanical revisions of the explicit assumptions of the search procedures, the imported assumptions and assumptions that make importation possible need also to be modelled. Such assumptions are "tacit" at least in the sense that they are not a part of the modelled search procedure itself. To call the assumptions tacit does not imply that they are irretrievably tacit, or unmodellable, but simply that they have not yet been made explicit or not yet been made part of the model; this limited sense of "tacit" will suffice for the discussion here. What I will have to say concerns the problem of the identification of such assumptions or traditions, and the peculiarities of their logical status. The problem of irretrievability does not of course, vanish, because we cannot assume that the assumptions are modellable.

THE MENO ACCORDING TO POLANYI AND SIMON

Perhaps the clearest confrontation on the topic of tacitness in the AI literature is found in the discussion of Herbert Simon of Polanyi's use of the *Meno* in *The Tacit Dimension*. The *Meno* is a dialogue about whether virtue can be taught, and the point Socrates makes at the end of the dialogue is essentially that, if it could be, either the

Sophists could teach it, or virtuous men could. His audience agrees with Gorgias that the Sophists only make men clever at speaking, and Socrates disposes of the idea that virtuous men might teach virtue by showing that various virtuous men had assured that their sons were well-instructed in those things which are a matter of knowledge or skill, but failed to make them "virtuous"; virtue, he concludes, is divinely given. The slave boy in the dialogue is brought in to make a related point, which is well described by Polanyi's phrase "we know more than we can tell," but is nevertheless oddly incongruent with Polanyi's larger thesis, and oddly congruent with Simon's.

Socrates establishes a) that the boy cannot correctly answer the question ("cannot tell," in Polanyi's language), of how much larger the sides of a square with double the area of another square will be, and b) that the boy thinks he knows that if a square has twice the area the sides will also be doubled.⁷ He then leads the boy through a series of inferences, each of which the boy could "tell," at least could assent in response to Socrates' "questions" formulating those influential steps, and that he could correctly multiply and add when asked. After these queries led him to the correct solution, Socrates turns to Menon, asks whether any of the opinions the boy gave in response to the queries were not his own, and then asks whether the correct opinions were not in the boy all along. Menon assents.⁸

Polanyi thinks the *Meno* shows conclusively that if all knowledge is explicit, i.e., capable of being clearly stated, then we cannot know a problem or look for its solution. And the *Meno* also shows, therefore, that if problems nevertheless exist, and discoveries can be made by solving them, we can know things, and important things, that we cannot tell.⁹

How exactly the *Meno* does this is not made clear by Polanyi. Presumably he has in mind that the boy had a kind of precognition of the solution that he could employ when coaxed, but not formulate in response to Socrates' direct query. He claims that "the kind of tacit knowledge that solves the paradox of the *Meno* consists in the intimation of something hidden, which we may yet discover," and suggests that there are parallels in the history of science. The Copernicans "must have meant to affirm" a "kind of foreknowledge" when the "passionately maintained during the hundred and forty years before Newton proved the point, that the heliocentric theory was not merely a convenient way of computing the paths of the planets, but was really true."¹⁰ The scientist's "valid knowledge of a problem," sense of "approaching its solution," and "valid anticipation of the yet indeterminate implications of the discovery," (meaning the sense that the discovery will lead to additional but "as yet undisclosed, perhaps as yet unthinkable, consequences") are thus accounted for by tacit knowing.

Bradie presented a counterexample directed not at this broader set of claims, but at Polanyi's formulation of the paradox that "to search for the solution of a problem is an absurdity; for either you know what you are looking for, and then there is no problem; or you do not know what you are looking for, and then you cannot expect to find anything."¹¹ One horn of this dilemma, that if you know what you are looking for, there is no problem, is false, Bradie suggests. He gives the example of a mathematician who is trying to refute Goldbach's conjecture: he knows what he is looking for, namely an even number which is not the sum of two primes, but he still has the problem of finding one.¹² One suspects that Polanyi would not accept this formulation, and indeed it seems like a pun on "problem." But it is a serious pun, as Simon's appropriation of the example makes evident, for, to the extent that "prob-

lems" are those cases where "you know what you are looking for," it may be possible to represent the problem solving process formally.¹³

Simon's expansion of this point runs as follows. It is possible to have an effective procedure for testing, and an effective procedure for generating candidates, we will have a "problem," i.e. an unsolved problem, where we nevertheless "know what we are looking for" without actually having it. In the case of Goldbach's conjecture, we can set up the following procedures: generate even numbers, generate numbers named by their prime factors, and make judgements of equality. The problem then can be defined as follows: "find a number k generated by the first procedure that does not belong to the numbers generated by the second procedure."¹⁴ Thus the example fits the "general scheme for defining problem solutions prior to finding them"¹⁵ as elaborated in such papers as "The Logic of Heuristic Decision-making," in which Simon also discusses the *Meno* paradox. In this paper he defines "solution," following John McCarthy, as "the object that satisfies the following test...."¹⁶

WHAT DID THE SLAVE BOY KNOW?

Did the slave boy possess a procedure for testing? Presumably he did not, or he would have been able to see that his first answer to Socrates was wrong. Socrates shows him that he was wrong by making him go through certain inferential steps which lead to an inconsistency with his original answer, and he recognizes the inconsistency. We might describe this in computational terms by saying that the knowledge the boy possessed consisted of the subroutines that generated each inferential step, plus subroutines that matched the outcomes and recorded a failure to match. On the basis of this description of this first part of the dialogue with the slave boy, one might question whether Socrates has in fact merely "queried" the slave boy: the slave boy apparently could not have performed the routine producing the result on his own; Socrates, by putting the pre-existing subroutines together into a routine measuring the area of squares in fact taught him something, namely the higher-level routine.

It is perhaps less clear that the boy possessed a procedure for matching answers in order to test candidate solutions; after all, he didn't recognize his original answer was wrong until Socrates showed that, using the boy's matching and calculating subroutines, one comes to a different numerical result. It is the inconsistency that the boy recognizes as an error. So one might say that Socrates represented the geometrical problem in a way that the boy's pre-existing "numerical inconsistency recognition subroutine" could be invoked, which it was not by Socrates' original query. Thus Socrates has taught at least one thing to the boy in this part of the dialogue, namely a higher level routine for calculating area, and perhaps another, the technique for the representation of the problem of area as a problem of adding the squares within a square, which Socrates supplies and the boy instantly accepts, and which is necessary for producing the inconsistency the boy recognizes.

Described in this way, Socrates' version of the event as accepted by Menon is untrue, but Polanyi might find it more congenial, for it suggests a way for him to salvage his point against Simon, and perhaps even salvage the paradox. When the slave boy "recognizes as valid" he is doing something other than executing subroutines or stringing them together to perform a complex task, and something different from "being programmed" by Socrates: after all, the higher level routine is being accepted as "valid", not merely as a well-formed command. Of course, one might attempt to push this description further, and describe the boy's "accepting as valid

routine," that is, model the process by which he accepts as valid and rejects as invalid the identification of the result of certain routines with certain answers, in this case answers to certain question about the area of squares. To do this, perhaps one would model "recognition as valid" as using a successful identification of this sort as a signal to accept the routine as the determinant of area, i.e. as a step in self-programming or "learning."

But this is not quite true to the case, for the boy seems (and Polanyi has precisely this sort of thing in mind) to have a precognition of the right method, and it is the fact that Socrates' suggestions match with this tacit precognition which enables him to "recognize as valid" Socrates' method of arriving at the answer. But precognition might be modelled as well, for example in terms of a preference of previously experienced patterns. And of course this process of finding discrepancies and patching the program to match additional features of the process can go on indefinitely.

One might emulate the slave boy in innumerable other ways: by beginning with a program for areas of all kinds of surfaces and blocking execution of any but the programs relating to squares, even the rules of recognition could be constructed differently, for they depend on what the learning process starts with. If we think like Socrates, we might assume that all knowledge of means of calculating areas of surfaces is already present in the boy, and thus treat recognition as recollection (plus modifiable error, as in all mental processes), and construct a recognition subroutine such that the primary test the computer would use in "recognizing" would be to match the suggested learning with the previously blocked parts of the program, and, if they match, unblock them.

The Simon-like "learning" model described earlier would be an alternative to his; Polanyi presumably would prefer another approach, perhaps the one signalled by Socrates' preliminary question to Menon, "is the boy a Greek, and does he speak our language?" One might suppose, for example, that mathematical concepts are not primary, but based on ontogenetically prior concepts and experiences that may or may not be shared with persons raised differently—that there may be different paths to, and therefore different underlying inferential frameworks, onto which for thinking about mathematical concepts is grafted. Some may reason visually, as the slave boy could, while others might find it more natural to derive their geometric results from numerical relations and to infer by operations on numbers. Indeed, there are strong reasons to believe that different cultures dispose those raised within them to different branches of mathematics, or against mathematics.

THE MEANING OF MATCHING

The sheer diversity of the kinds of models that can emulate any given cognitive process of this sort raises a question about their status. Are these models merely analogies, which may be useful for various purposes? Or is there some fact of the matter which enables us to say one analogy is right and another wrong? Does the ultimately right model also map on to features of the physical computing mechanism of the human brain? Or are there no ultimately right models, but only purpose-relative models? What is the thing called knowledge that expert systems attempt to model? What sorts of things are the "processes of scientific discovery" that Simon and his collaborators say "may be describable and explainable simply as special classes of" problem-solving processes previously modelled on computers, and in what sense do these models constitute "an empirical account"?¹⁷ It should be evident that some sort of ontic claim is being made by writers like Simon, but it is in many ways a puzzle.

Under one interpretation, the claim is simple: the cognitive processes of a Kepler were processes in the brain, and therefore the correct model of Kepler's discovery will be the model which corresponds to those processes. Claims about tacit knowledge in relation to Kepler's discovery will be settled by having the correct model, which will either contain something corresponding to tacit knowledge or will not. The mental processes of Kepler are of course irrecoverable. But the model which is constructed will inevitably represent basic cognitive processes shared by many persons and validated by empirical psychological experimentation. Thus the sense in which these are "empirical accounts" is that the models in question are empirically tested by psychology and that they match with the basic physical features of cognition that have been empirically established by neurobiologists.

The place of this last clause in the argument appears modest, but it is important to see that it is in fact central. Turning a decision between competing models of a given cognitive process, such as the slave boy's geometrizing, into an empirical question entangles us in the difficulties to which any kind of human psychological experimentation is subject. The models are ordinarily attempts to describe what are taken to be probabilistic processes, for example of arithmetic errors. The tests of the fit of these models are themselves probabilistic. Thus there is inevitably a high degree of arbitrariness in the test of the model, as various assumptions about the form of curves, about underlying distributions, and so on must be made to apply tests of fit and curve-fitting heuristics, and the like. Because the models are themselves probabilistic, many models might be generated that fit equally or almost equally well; and because sharp quantitative discriminations between them must be based on arbitrary assumptions, those that fit almost equally well cannot be non-arbitrarily rejected. Moreover, what fits in a highly controlled experimental setting may simply not generalize to the real world cognitive processes to which they are assumed to be equivalent, and the difficulties in such simple cases as arithmetic errors are compounded in cases of great complexity, such as Kepler's reasoning processes.

The reason those difficulties are usually ignored is the expectation, to quote Meehl quoting Skinner in conversation, that

When the laws of behavior are sufficiently worked out in mathematical detail (in the next generation following me), and when the anatomy, especially the microanatomy, and the physiology of the brain are very thoroughly understood, there will be no problem of prematurely forcing a speculative translation because it will be perfectly apparent how the brain/behavior dictionary will read.¹⁴

Presumably it would also be perfectly apparent which of the competing cognitive models, should any remain at the time our physiological and anatomical knowledge become sufficiently advanced, are to be discarded.

If we consider the problem of underdetermination apart from this expectation, matters look quite different. The case for a model's reality, or for its explanatory force, depends on the existence of match between model and process. So this expectation serves another role in making the argument plausible, that of permitting questions about the constitution of the objects, such as "processes of discovery" to be matched with the models, to be treated as non-basic: that is, as temporary or *prima facie* assumptions which one expects to replace. Again, this is an apparently small matter, with large implications. In actual cases not only are the relations between models and processes highly underdetermined—that is, the same "process" can be modelled by a variety of computer programs, each using different devices—but the gap between the *in vivo* practices of reasoning and the models of them which presently exist is usually large. In cases of expert systems consciously designed to

emulate the knowledge of the expert other than fairly routine mathematical problem-solving and other very simple cases, there is a great deal of interpretation or skill which needs to be supplied by the end-user.¹⁵ This implies that in these cases the expert system is not a complete emulation of the relevant expert knowledge.

The idea of "completeness" is itself puzzling. It lends itself to the line of argument that Polanyi employs in his own attempts to prove the existence of tacit knowledge, as well as to attempts to show that AI programs are not "really thinking." But such arguments share a premise, or discursive practice, with their opponent, namely the idea that the description of the process to which the model or expert system is being compared has some sort of ontic significance which the model or system itself acquires or fails to acquire when it is successfully or unsuccessfully matched. The project of matching assumes that the description is itself ontically valid in some sense, or at least uniquely legitimate as a description.

There is a kind of *cul-de-sac* which arises in relation to these matching arguments that needs to be briefly considered. As I have suggested, the computer modeller can attempt to emulate a particular real-world cognitive process, and the describer of the process can elaborate the description in ways which the computer modeller can then attempt to match. The dialectic between the modeller and the describer of features of the cognitive situation can go on indefinitely in this fashion. It is in connection with this feature of the modelling process in which the question of whether a computer really "recognizes a valid," really "thinks," and so forth finds a place. At each stage in the elaboration of the model the critic can say "really thinking would be when such and such," and at each stage the modeller can offer some version of such and such, only to find that it is not clear that such and such fully qualifies. Important as these questions may be for the question of man's place in the universe and for the question of the ultimate coherence of the goals of Simon-like programs of modelling human cognition, discussions of these issues do not seem to reach determinate results. Technical solutions or substitutes for such concepts as "recognition" are necessarily in part conventions or stipulations that fix the meaning of a term and consequently change it, if it is a term with an open texture such as the cognitive terms in question in discussions of scientific discovery. But this is perhaps not as decisive an issue as it might seem.

The fixing of conventions to specify one meaning of a term and allowing others to fall into disuse is one way in which language itself changes. The history of the term "pneuma" and its present form in the word "pneumatic" is illustrative. The association of pneuma with breathing and breathing with the soul has simply no part in present usage, which applies the term to suction devices and analogies with the device. This should suffice to remind us of the defective theory of language hiding behind many of these objections, namely the museum theory of meaning, and to point to a larger set of difficulties with any thesis which presumes that meanings are eternally fixed and that the laws governing the range of application of terms are there to be "found," as jurisprudential theorists say. In relation to the problem of computer modelling, the problem may be simply put: any attempt to match meanings and claim that the computer's meaning fails to match the actual meaning of some term or mental predicate itself on a construction of the "meaning" to be used as a standard. Such constructions have the same difficulty as the computer modellers' technical substitutions. Like them, the construction fix the meaning in a particular way. Hence they are always questionable when treated as a standard to which a model is matched, in much the same way that modellers' technical substitutions are questionable.

THE TACIT

We may distinguish the matching problem, where the computer program is held to the standard of a description of a process or a definition arrived at independently and taken as valid, from a kind of test which turns the problem of matching into an empirical or naturalistic problem. The Turing test is the classic strategy for "empirically" evaluating the attempt to emulation; similar problems arise in attempts to "elicit" expert knowledge that an expert system should contain. The way in which knowledge is elicited is by a series of comparisons, in which failures to match expert judgements are "corrected" by adding new rules. These rules may be said to represent the expert's tacit knowledge, in at least this sense: many of the rules make explicit what the expert did not think to or perhaps even could not articulate as part of his or her "knowledge."

The problem of knowledge elicitation has some familiar analogues in sociology, and shares the features of a family of methods of revealing tacit aspects of an activity in which a comparison between practices that differ shows what assumptions a given practice depends on to produce the same practical results as another. The study of what Mauss called body techniques is a simple example.²⁰ We assume our way of sitting is natural. It comes as a shock to us to learn that other cultures find it *unnatural*, and "sit" quite differently. We then can see that sitting is learned, and even see how it is learned. But without the comparison we would have been hard-pressed to thematize "sitting" as a technique at all. Conceivably someone would have invented, or happened upon, a full-blown alternative, that was equally "comfortable" or "natural." But when we consider how much learning (from the earliest months, when infants are praised for sitting up) has gone into sitting, we can see that this is not likely. Nay other position would "feel wrong" to a person without the appropriate childhood preparation. So without an actual lived alternative, we would very likely never come to recognize it as a technique that was tacitly taken by us as natural.

The recognition of such a practice as a practice thus depends on our ability to form a contrast between this and some alternative practice, a contrast where the two practices are similar enough that a hypothesis about the character of the alternative practice can be formulated and supported by descriptions of the activity. Similarly, the recognition of something as tacit knowledge depends on a contrast. The contrast may be to an idealization, such as an expert system program, or to a divergent form of a practice or body of knowledge, or it may be, as in the case of Garfinkelian ethnomethodology, a contrast to a description of the activity that identifies some feature of the activity—some "orderliness" to it—that can be accounted for only by supplementing the usual description with some hypothesis about tacit sense-making practices of some sort.

In either case, it is contingent as to what is "revealed," and the attribution of an assumption, tacit knowledge, a skill, or a practice is relative to the project of comparison or the means of generating the alternative description. In the case of expert systems, for example, what counts as knowledge is relative to the body of decisions, or judgements chosen as the standard. "Completeness" has a sense here only relative to the comparison, and this is characteristic of attributions of tacit knowledge: there are no natural wholes here, only artifacts of comparison.

This is perhaps a banal point, but its implications are often misunderstood. Aaron Cicourel, for example, in a discussion of expert measurement systems, his famous argument, "that traditional social measurement must incorporate the pragmatic conditions of everyday life in order to satisfy ecological validity conditions," to "studies

of problem solving or decision-making and of beliefs, attitudes, opinions, and general knowledge of the world" by suggesting that the "on-line, adaptive potential of computer simulation models"²¹ can be used to improve social measurement. He claims that "our use of language presumes various folk models or cognitive schemata of ideas, objects, events, and action,"²² and that consequently "we must design our expert systems to reflect the ecological validity of problem-solving conditions that are presupposed by general reasoning processes and the use of standardized language and thus model the way everyday reasoning, language, and action enter into all attempts at social measurement or the creation of expert systems."²³

The actual case material which Cicourel deals with involves problems users have in responding to particular queries: the solutions are attempts to get a better grip on the techniques users employ in reading the queries, so as to get better responses. The role of discussions of "presuppositions" is thus bound to a practical end, and occur within a community with shared practices, to which the problem itself is relative. In my terms, the "practices" which are hypothesized and "proved" in some sense by the removal of the difficulties are "practices" relative to this problem. Cicourel appears to think of these practices differently, and this is reflected in the way he speaks of modelling "the way everyday reasoning enters" into "all attempts at...the creation of expert systems" as though this were a *general* modelling task, that is, as though there was a class of objects, namely folk models, that could themselves be modelled, and even as though the completion of this task would result in a scheme in which "ecological validity" could itself be accounted for. If the failing of more primitive systems was the result of a failure to include "everyday" thinking, the inclusion of these models would presumably make this failure good.

This is plainly utopian, for several reasons. First, like expert systems themselves, these models would be highly underdetermined, that is, there would be a wide array of alternative models that would emulate the same results, as in the *Meno* example. So the idea that *one* of them corresponds to the real thing is highly problematic. It may be that each of us goes around with differently constructed models, each of which more or less emulates other persons, as in Quine's famous analogy between individuals' theories of the world and shrubs, each of which is differently structured internally, for examples with different arrangements of branches, but all of which have been trimmed to look the same.²⁴ In this case even the success of some attempt to emulate the general model of common sense would tell us nothing about the total structure of individuals' reasoning practices.

Second, the leap from specific problem-frames to general claims about commonsense reasoning is wholly unwarranted: Cicourel's own models, like expert systems, are also going to be subject to ecological validity limitations. Indeed, the cases are precisely parallel. The only reason Cicourel can have for believing that his overarching model would be potentially free of such limitations is that it in some sense accounts for the "ecological validity" of the expert system it improves. But its actual successes would be relative to the matching or comprehension problems of particular users of a particular expert system, which warrant no non-relative claims whatsoever.

The ideal to which he holds expert measurement systems is global identity, or something similarly ambitious. If what I have suggested about the salamander analogy is correct, this ideal is incoherent. One might say that the belief in its attainment depends on generalizing erroneously from successes in revealing *parts* of the tacit dimension to some non-relative "whole." These successes, however, depend on "techniques" which assure us nothing about the possibility of an *exhaustive* survey of the tacit dimension, and indeed there is no whole of which the results of these

analyses can be understood as parts. Worse: the salamanders are not absolute salamanders, but relative salamanders; each of which is like a hypothesis and consequently underdetermined in its description; absent some lucky accident by which these descriptions would be synthesized or a privileged class of true non-relative descriptions could be selected, this underdetermination seems largely irreducible. In any case, none of the means by which the tacit is made accessible warrant claims about the "whole" of the tacit knowledge. So "global identity" between any model and the whole of tacit knowledge or any larger thing of which tacit knowledge, taken as a whole, is a part, such as "the scientific cognitive processes of community x," is a hopeless aim.

MODELLING THE TACIT IN PRINCIPLE AND IN PRACTICE

The picture that emerges from these considerations is this: both "artificial" matching projects, such as the Langley et al model of Kepler's discovery, in which the "process" to be watched is "constructed" such that other descriptions have equal claim to ontic validity, and "natural" matching problems, such as knowledge-elicitation for an expert system, warrant claims about tacit knowledge, but these claims are very weak. First, they are highly underdetermined hypotheses: a wide range of different descriptions of the tacit practice will solve a given specific, matching problem. Second, they are relative to the particular contrast-problem solved by the hypothesis: there is no sense to claims about the complete characterization of a person's or community's tacit knowledge, because the programs of comparison are all themselves necessarily limited.

What this, to paraphrase Weber's famous remark in a different context, we come to matters of faith. In view of the highly under-determined and context-bound character of our claims about tacit knowledge and therefore the cognitive processes which employ it, a few reflections are in order about the *project* of a cognitive psychology of scientific discovery, and about the relation in this project between reasonable expectations for its future achievements and the element of faith. One aspect of the Simonian image of cognitive process that has been bracketed in the present discussion is the prospect that the brain-behavior dictionary will be apparent once the relevant physiology and psychology is at hand. When Skinner formulated this expectation, it is evident that he envisioned a moment when there would be a much simpler body of "psychological models" than the diverse, underdetermined, and contextually limited simulations I have described here. What now can be said to remain of this article of faith?

One might retain the hope that some simulacrum of this moment might yet arrive. It might take the following form: certain AI routines widely useful in various cognitive emulation projects might be shown to match particular neurophysiological processes, thus resolving at a stroke the problem produced by the existence of a large class of competing descriptions of cognitive processes by selecting out this class of procedures as ontically valid descriptions and frameworks for description. But even if we grant Simon's premise that the brain is a computer, there is precious little reason to believe that this day will arrive. The idea of opening up the black box and finding in it the mechanisms of mind is a powerful image. But Mill's image of a nomic social science was a powerful image as well: its realization was foiled by underdetermination and complexity, its early successes were later seen to be very limited in scope and significance, and practical difficulties over evidence that appeared at first to be minor gradually came to be appreciated as major. Most of the false starts in the history of social sciences, and there have been many, have had these characteristics.

That AI modelling of cognitive processes understood as a program of explanation and description should be so marked by each of these traits, even at this early stage in the development of the project, inspires not faith but skepticism.

Notes

1. Joseph Weizenbaum, *Computer Power and Human Reason: From Judgment to Calculation*. New York: W. H. Freeman, 1976. 71, 124-5.
2. *Ibid.* 137.
3. *Ibid.* 226.
4. David Freeman, "Statistics and the Scientific Method," in William M. Mason and Stephen Fienberg (eds.) *Cohort Analysis in Social Research: Beyond the Identification Problem*. New York: Springer-Verlag, 1985. 343-366.
5. Clark Glymour, "Social Science and Social Physics," *Behavioral Science*, 28. 1983. 129-30.
6. Meno. 99.
7. *Ibid.* 81-2.
8. *Ibid.* 85.
9. Michael Polanyi, *The Tacit Dimension*. Garden City, N.Y.: Doubleday, 1966. 22.
10. *Ibid.* 23-4.
11. *Ibid.* 22.
12. Michael Bradie, "Polanyi on the Meno Paradox," *Philosophy of Science*, 41. 1974. 203.
13. One may also question whether Bradie adequately formulates the issue, for the "problem" of refuting Goldbach's conjecture in this fashion, i.e. by finding a number, is apparently insoluble. A short retort by Polanyi might be to deny that insoluble problems are problems, or simply to restrict the paradox to soluble problems, or to deny that the problem of refuting a conjecture is properly speaking a "problem," though indeed it may involve searching.
14. Herbert Simon, *Models of Discovery*, Dordrecht: Reidel, 1971. 340.
15. *Ibid.*
16. *Ibid.* 161.

17. Pat Langley, H. A. Simon, Gary L. Bradshaw, and Jan Zytkow, *Scientific Discovery: Computational Explorations of the Creative Process*, Cambridge, Mass.: MIT Press, 1987. Preface.
18. Paul E. Meehl, "What Social Scientists Don't Understand," in Donald W. Fiske and Richard A. Shweder (eds.), *Metatheory In Social Science*. Chicago: University of Chicago Press, 1986. 319.
19. M. Collins, R. H. Green, R. C. Draper, "Where's the Expertise?: Expert Systems as a Medium of Knowledge Transfer," in Martin Merry (ed.), *Expert Systems 85*, Cambridge: Cambridge University Press, 1985. 323-334.
20. Marcel Mauss, "La Technique du Corps,"
21. Aaron Cicourel, "Social Measurement as the Creation of Expert Systems," in Donald W. Fiske and Richard Shweder (eds.), *op. cit.* 1986. Note 18. 261.
22. Ibid. 261-262.
23. Ibid. 268.
24. W. V. O. Quine, *Word and Object*,.

Part Three

The Modularity of Scientific Cognition

SOCIOLOGY OF THE SCIENCES

A YEARBOOK

Managing Editor:

R. D. Whitley

Manchester Business School, University of Manchester

Editorial Board:

G. Böhme, *Technische Hochschule, Darmstadt*

N. Elias, *Amsterdam*

Y. Ezrahi, *The Hebrew University of Jerusalem*

L. Graham, *Massachusetts Institute of Technology*

T. Lenoir, *University of Pennsylvania*

E. Mendelsohn, *Harvard University*

H. Nowotny, *Institut für Wissenschaftstheorie und Wissenschaftsforschung, Vienna*

Claire Salomon-Bayet, *Paris*

R. Schwartz-Cowan, *State University of New York at Stony Brook*

T. Shinn, *Groupe d'Etude des Méthodes de l'Analyse Sociologique, Paris*

P. Weingart, *University of Bielefeld*

VOLUME XIII – 1989

THE COGNITIVE TURN

Sociological and Psychological
Perspectives on Science

Edited by

STEVE FULLER

Virginia Polytechnic Institute, U.S.A.

MARC DE MEY

Rijksuniversiteit Gent, Belgium

TERRY SHINN

Centre National de la Recherche Scientifique, France

and

STEVE WOOLGAR

Brunel University, U.K.



KLUWER ACADEMIC PUBLISHERS

DORDRECHT / BOSTON / LONDON