

THE 'PERIODIC TABLE' OF EXPERTISES

[COPYRIGHT HARRY COLLINS AND ROB EVANS]

This is a draft that, all being well, will become two chapters of a book entitled *Expertise*. Here we present the periodic table of expertises and drafts of the two chapters which explain it.

IF YOU REFER TO THIS PIECE IT PLEASE CITE AS

**COLLINS HARRY, AND EVANS ROBERT, (2004) 'PERIODIC TABLE OF EXPERTISES'
WWW.CF.AC.UK/SOCSI/EXPERT**

PLEASE SEND COMMENTS ON THE DRAFT TO

CollinsHM@cf.ac.uk and/or EvansRJ1@cf.ac.uk

UBIQUITOUS EXPERTISES

DISPOSITIONS

Interactive Ability
Reflective Ability

SPECIALIST

UBIQUITOUS TACIT Knowledge

SPECIALIST TACIT Knowledge

EXPERTISES

Beer-mat
Knowledge

Popular
Understanding

Primary Source
Knowledge

Interactional
Expertise

Contributory
Expertise

META-

EXTERNAL

INTERNAL

EXPERTISES

Ubiquitous
Discrimination

Local
Discrimination

Technical
Connoisseurship

Downward
Discrimination

Referred
Expertise

META- CRITERIA

Credentials

Experience

Track-Record

PERIODIC TABLE OF EXPERTISES

DRAFT CHAPTER 1

THE LADDER OF EXPERTISE

UBIQUITOUS EXPERTISES

Ordinary people are talented and skillful almost beyond comprehension. We can say 'almost beyond comprehension' with confidence because a lot of very clever people have tried to encapsulate the talents of ordinary people in computer programs entirely failing to realise how hard a task it would be. What we will call 'ubiquitous expertises' include all the endlessly indescribable skills it takes to live in a human society; these were once thought of as 'mere commonsense'.¹ For any specific society, its 'form-of-life,' or 'culture,' provides, and is enabled by, the content of the ubiquitous expertises of its members. Fluency in the natural language of the society is just one example of an ubiquitous expertise.² Others that we will refer to include

¹ We refer to *ubiquitous* expertises because the antonym of esoteric is 'exoteric,' which the *Chamber's Dictionary* defines as, 'intelligible to the uninitiated; popular or commonplace.' This definition renders the word 'exoteric' inappropriate when paired with expertise because by definition you cannot have expertise which is intelligible to the uninitiated: expertise is available only to the initiated or experienced; that is why we use ubiquitous rather the exoteric as the antonym to esoteric. Our distrust of the term 'lay expertise' also has its roots in the definition. The *Chambers Dictionary* defines a layman as follows: one of the laity; a non-professional person; someone who is not an expert. But all lay-persons possess ubiquitous expertises.

The Aristotelian term 'phronesis,' being some sort of combination of prudence and wisdom -- a practical wisdom in a moral setting -- captures part of the notion of ubiquitous expertise but not that part represented by language speaking and the like. Perhaps the idea of the 'soul' -- that which separates us from the animals and gives us language and moral judgement also captures the idea.

² Huge resources have been expended on attempts to give machines fluency in natural languages, each new supposed solution to the problem being greeted with enormous hype supporting marketing campaigns to promote close to useless products. One explanation of the failure is found in Collins 1990, where it is demonstrated that a device based on Searle's 'Chinese Room,' which encodes all the formal knowledge there can be about English-Chinese translation, would fail a test in which it was required to 'repair' badly

moral sensibility, and political discrimination. These are abilities that people acquire as they learn to navigate their way through life.³ In earlier chapters of the book and in Collins and Evans, 2002 (see link on this website), we describe the 'Problem of Legitimacy' and the Problem of Extension.' New technologies will not be politically acceptable unless 'the public' is allowed to feel part of the decision-making process. This is the Problem of Legitimacy.' It gives rise, however, to the question: 'How far should technical expertise be taken to extend among the public?' -- the Problem of Extension. In the case of ubiquitous expertises the Problem of Extension ceases to have any practical significance.

Thus, when we say that claims about the abilities of ordinary people, such as are made in the 'Politics of GM Food' report, are wrong, we do not mean that ordinary people do not have expertise, we mean only that the ubiquitous expertise of ordinary people should not be confused with the expertise of technical specialists.⁴ What we are arguing is that we must preserve a logical space for expertises that are not the property of the general public; it is impossible for the general public to have expertise in every specialist technical domain even though they have a vast store of ubiquitous expertise.

Before passing to specialist domains it is worth noting that just because some of the things we can all do are hugely skillful it does not mean that all the things we can all do are hugely skillful and this includes things of which we have great experience. For example, one might have huge experience of lying in bed in the morning but this does not make one an expert at it (except in an amusing ironic sense). Why not? Because anyone could master it immediately without practice so nothing in the way of skill has been gained through the experience.

Now we turn to the question of how much expertise in specialist domains it is possible for ordinary people to have. To answer the question we need to think about ways of having specialist expertise. Let us then construct a rough ladder of knowledge about, or expertise in, specialist domains. No doubt the bottom of the ladder could be divided up in more or less ways and along different dimensions but we need to start with something.

At some stage all human expertise touches on tacit knowledge, that is, an understanding of rules that cannot be expressed. It is, indeed, the inexpressibility of the rules of ubiquitous expertises that make them so

formed input (sentences full of grammatical or spelling mistakes, of which there are an indefinite number of variations). Yes this is a skill which is used continually by a native speakers who are always being presented with varieties of faulty input that they have never encountered before.

³ The success of lawyers suing firms such as MacDonald's for selling over-hot coffee, and the consequent growth of warnings and safeguards surrounding every consumer good, is patronising -- treating the public as incapable of learning the rules of ordinary living through the normal processes of socialisation.

⁴ 'The Politics of GM Food: Risk, science and public trust.' Swindon: Economic and Social Research Council [ESRC Special Briefing No 5, October 1999].

hard to capture in computers. 'Tacit knowledge' enters into knowledge acquisition in two ways. Some kinds of knowledge acquisition amounts to the acquisition of additional specialist tacit knowledge; other kinds of knowledge acquisition involve the exercise of tacit knowledge in the course of the acquisition of information. In the second kind of knowledge acquisition, the tacit knowledge used is that found in ubiquitous expertises. For example, the exercise of the ubiquitous expertise associated with language can be used to acquire new information (explicit knowledge) through reading.⁵ The first rungs of a ladder of specialist expertise involve only ubiquitous tacit knowledge used in this way. The higher level rungs require immersion in the tacit knowledge of the specialist domain in question so that more tacit knowledge can be acquired. For example, mastery of even a widely distributed tacit-knowledge laden expertise such as car-driving needs practice at car-driving and internalisation of the unspoken rules of roadcraft; becoming a full-blown specialist in a scientific or technical domain requires immersion in the society of the domain specialists.⁶ This gives us an initial division of types of expertise: those, the acquisition of which involves only ubiquitous tacit knowledge and those that involve specialist tacit knowledge.

EXPERTISES INVOLVING UBIQUITOUS TACIT KNOWLEDGE

Beer mat knowledge

The 'floor' upon which any ladder of human expertise stands is, then, ubiquitous expertise. The first rung of the specialist ladder is what we will call 'beer mat knowledge.' Consider the following explanation of how a hologram works.

A hologram is like a 3 dimensional photograph -- one you can look right into. In an ordinary snapshot, the picture you see is of an object viewed from one position by a camera in normal light.

⁵ Collins and Kusch (1998) argue that there is a kind of expertise -- the ability to carry out mimeomorphic actions as opposed to polimorphic actions -- that, in essence, involve no tacit knowledge even of the ubiquitous expertise kind. These kind of actions can be mimicked by machines since to execute them is, effectively, to mimic a machine. In these kind of actions the behaviour -- the physical movements -- are of the essence and intention and understanding are irrelevant. The golf-swing and mechanical dictionary translations (with all its mistakes) are examples.

⁶ Notice that an expertise such as skillful car-driving is very widely distributed but it is not ubiquitous. Car-driving is not learned integrally with learning to live in society but needs specialist training and the specialist tacit knowledge that goes along with it even though nearly everyone in certain societies can do it to a high level of skill.

The difference with a hologram is that the object has been photographed in laser light, split to go all around the object. The result -- a truly 3 dimensional picture!

This explanation, found on a beer-mat made for the 'Babycham' company in 1985, appears to give an answer to the question 'What Is a Hologram?' It is capable, presumably, of making at least some people feel that they now know more about holograms. The words on the beer-mat are not simply nonsense nor could they be taken to be, say, a riddle or a joke. Presumably there are people now alive who have studied the beer-mat and, if asked: 'Do you know how a hologram works?', would reply: 'Yes', whereas immediately before they had read the beer-mat they would have answered: 'No', to the same question. So what increment in expertise does someone have in consequence of perusing the beer-mat?

Let us investigate by analysis another such thing that one might know, the rule for the move of the bishop in chess. The rule, and we might well read it on a beermat or something similar, is 'the bishop may move diagonally only, any distance, backwards or forwards.' But it is possible to know this rule in more than one way. One might know it in the same way as an observant Jew or religious Catholic might know how to recite certain prayers in Hebrew or Latin respectively but without knowing their meaning. Thus, knowing how to 'chant' the bishop's move might enable one to score a point in a board game such as 'Trivial Pursuit' which is intended to discriminate between levels of general knowledge. Crucially, knowing the bishop's move in that way does not imply that one knows much about what it might mean. For example, you can know it in the beer-mat/Trivial Pursuit way without knowing that the term 'any distance' within the rule is to be measured in squares on the chess board and that it can never be more than eight squares nor that 'any distance' means only so long as the path is not blocked by another piece, nor that the restriction to the diagonal implies, on a chequerboard, that the bishop is restricted to squares of only one colour (nor even that there are two colours on the chequerboard). In short, knowing the bishop's move in a beer-mat-knowledge kind of way does not enable one to do anything much that one would not be able to do if one did not know it (other than scoring points in general knowledge tests). Knowing the rule for the bishop's move in the context of chequerboards and the game of chess is a rather different thing even for a novice chess-player than for a Trivial Pursuit player.⁷ The novice chess-player who knows the rule knows how to move the bishop on a chequerboard.

⁷ Knowing the rule as a more experienced chess player is yet another thing. To see the difference between a novice and a slightly experienced chess-player consider that the author who wrote this passage, who is an experienced but poor chess player, had to think hard about how to write down the formula for the bishop's move for the sake of this passage even though he had no doubt that he already 'knew' it in some non-explicit sense. A novice would have less difficulty in explicating the rule since it would not yet have become 'internalised.'

Going back to the hologram, the explanation on the beer-mat does not enable the naive reader to do anything such as make a hologram, or debate the nature of holograms, or to make a sensible decision about what should be done about holograms, or to correct anyone's mistakes about the nature of holograms, or convey any information about holograms other than the formula itself.

To give a real-life example, one of the authors of this piece 'knows' that the rest mass of the neutrino is '4eV'. Again, he knows it only as beer-mat knowledge; to him it is a discrete fact that, on its own, puts him in a position to do almost nothing that he could not already do. '40,000' would carry just as much meaning as '4' for that author of this piece when it comes to the rest masses of neutrinos.⁸

Popular understanding of science

Moving to the next rung of the ladder, much superior to beer-mat knowledge is what we will call 'popular understanding.'⁹ Popular understanding can be gained by gathering information about a scientific field from the mass media and popular books. It is the kind of understanding to which bodies such as the Royal Society's 'Committee for the Public Understanding of Science' (COPUS), directs (or at least, once directed), its efforts.

Popular understanding does involve a deeper understanding of the meaning of the information than beer-mat knowledge. For example, it may be possible to make some inferences from popular understanding of science of the kind 'antibiotics will not cure viral diseases, influenza is a viral disease, antibiotics won't cure influenza,' or 'the element is completely enclosed in my electric kettle whereas heat is wasted when I boil water on the gas stove so the electric kettle uses less energy to boil the same amount of water than the gas stove provided not too much energy is wasted converting gas to electricity in the power station.' Popular understanding of science is also transmissible from one person to another to a certain extent -- transmissible as a set of ideas rather than a set of formulae.

In the case of a long-settled science the difference between a deeper understanding of science and technology and a popular understanding is not very important in terms of public decision-making; where the science is settled, the difference between scientific knowledge as revelation and deep scientific understanding has little impact on the conclusions reached. Where the science is the subject of a dispute, however, the difference is of the essence. The last three decades of social studies of science has shown us that in disputed science, detail, tacit knowledge, and unspoken understandings of who is to be trusted among those who work in in the esoteric core of the science and who is not, are vital components of decision-making at the technical

⁸ The examples of the rest mass of the neutrino and the hologram are taken from Collins, 1990.

⁹ Matthew Harvey made the final comment that made it clear that we must include a category like this which was absent from earlier discussions.

level. Popular understanding embodies a more extreme version of the process of hiding detail and ignoring reservations that leads to decision-making in disputed science. The idea, well-established in the sociology of scientific knowledge, can be summed up in the phrase 'distance lends enchantment:' the more distant is one from the locus of creation of knowledge in social space and time the more certain will the knowledge appear to be. This is because to create certainty, the skill and fallible effort that goes into making an experiment work, or a theory acceptable, has to be hidden; if the human activity that is experimentation is seen clearly, then it is also possible to see all the things that could be wrong.¹⁰ Any redescription of events in the core of science, even when it is designed for a professional audience, is bound to simplify; when the description is for a popular audience it will simplify more brutally. But sound judgements, or at least informed judgements, in disputed science must take account of many more of these uncertainties than popular understanding allows for. For this reason, in the case of disputed science, a level of understanding equivalent to popular understanding can yield technical judgements no better than chance.

So far in this passage we have talked simply of settled and disputatious science, but disputed science comes in several forms. In a later chapter we will explore the different forms of non-consensual science in more detail and work out their implications for public decision-making.

The unsoundness of judgements based on popular understanding applies whether the conclusion is positive or negative -- whether the consumer of the simplified version believes everything they read and hear (eg they might accept that Stephen Hawking's utterances about black holes are revealed truth) or rejects the claims (eg they might be certain that everything the government says about the safety of vaccines is false). Both kinds of interpretation of evidence are strengthened and reinforced by distance and by the 'narrow bandwidth' of the media which provide popular popular understanding.¹¹

¹⁰ For 'distance lends enchantment,' see Collins 1985. For a modification see Mackenzie 1998. Ludwik Fleck, who was a sociologist of scientific knowledge before the term was invented, wrote in the 1930s that:

Characteristic of the popular presentation is the omission both of detail and especially of controversial opinions; this produces an artificial simplification. [and] ... the apodictic valuation simply to accept or reject a certain point of view. Simplified, lucid, and apodictic science -- these are the most important characteristics of exoteric knowledge. In place of the specific constraint of thought by any proof, which can be found only with great effort, a vivid picture is created through simplification and valuation. (Fleck 1935, p 112-113 -- original stress)

¹¹ Treatments that turn on the establishment of scientific knowledge as a matter of literary transformation as work passes from laboratory to the wider world (for example, Latour and Woolgar, 1979), accurately describe the certainty amplifying property of the narrowing of bandwidth. They do not, however, explain how the resulting certainty comes to be positive or negative.

One of the troubles with the old COPUS-style approach to popular understanding is that it does not distinguish between consensual science and disputed science. It tends to present even disputed science as revealed knowledge emerging from a unified community of experts. This converts any genuine effort at increasing public understanding into propaganda.¹² The obvious danger, even for those keen on propaganda, is that for each positive piece of propaganda there is a negative one which will be grasped with equally unmodulated certainty.

In sum, public understanding is big step up from beer-mat knowledge but a long way from deep understanding of scientific matters. The gap between popular understanding and deep understanding is not so important where the science is settled and consensual but it is very important where science is disputed. Not by chance, wherever there is a serious public debate involving science, the science is nearly always disputed so the enchantment brought about by distance from the research front, whether negative or positive, is very important.

Primary source knowledge

The next step after popular understanding is the kind of knowledge that comes with reading primary or quasi-primary literature. We will call it 'primary source knowledge.' Nowadays the internet is a powerful resource for this kind of material. But even the primary sources provide only a shallow or misleading appreciation of science in deeply disputed areas though this is far from obvious: reading the primary literature is so hard, and the material can be so technical, that it gives the impression that real technical mastery is being achieved.

It may that the feelings of confidence that come with a mastery of the primary literature is a factor feeding into the 'folk-wisdom view' [the view that ordinary people are wise in the ways of science and technology]. But any amateur trying to apply the knowledge gained from car-repair manuals will soon learn the bitter lesson that much less can be done as result of reading information than appears to be the case. The same applies to doctoral students in the sciences; their first experience of real research is usually a shock, however well accomplished they have become in reading the published literature and doing well-rehearsed experiments as undergraduates. But even experienced scientists tend not to understand the amount of tacit knowledge on which their abilities depend. Thus, studies of tacit knowledge transmission show, *inter alia*, that scientists will embark confidently on an experimental project having done nothing more than read the literature and only later discover the degree of joint practice and/or linguistic socialisation that is needed to

¹² The modern high priest of this approach is Lewis Wolpert, once chair of COPUS, whose book The Unnatural Nature of Science, stressed just how different a scientific grasp of matters was to a commonsense appreciation. This is a curious approach coming from one who was supposed to be developing a public understanding of matters scientific.

make a success of it (to generate the capacity to do the thing rather than talk about it).¹³ Given trainee scientists' experiences, and professional scientists' lack of reflective appreciation of their own tacit knowledge, it is no surprise that a member of the public encountering the professional journals or the internet might easily come to think that they have found a direct line to understanding.¹⁴

Actually, it can be shown that what is found in the literature, if read by someone with no contact with the core-groups of scientists who actually carry out the research in disputed areas, can give a false impression of the content of the science as well as the level of certainty. Most of the papers in the professional literature are never read so if one wants to gain something even approximating to a rough version of agreed scientific knowledge from published sources one has first to know what to read and what not to read; this requires social contact with the expert community. Reading the professional literature is a long way

¹³ Collins 1974, 1985/92, 2001.

¹⁴ A familiar image is today's informed patient visiting their doctor armed with a swathe of material printed from the internet. While this kind of information gathering, especially in the context of a support or discussion group, can be valuable, it is important not to lose sight of what sociologists have shown: a great deal of training and experience is needed to evaluate such information. Sociologists of science seem to forget the lessons of their own subject rather easily.

The idea of active and informed citizens taking on established expertise is also discussed in Alan Irwin's (1995) idea of 'Citizen Science.' In advocating a Citizen Science, that is a science 'of' rather than 'for' the people, Irwin's focus is often on the recognition of the expertise that already exists among the citizenry, or about empowering local communities to acquire or recruit their expertise. As a result, Irwin is rather sceptical of the implicit agenda of many participatory exercises in which the aim is to find 'real' (that is 'uninterested') citizens, and privilege their views over those of activist communities (Irwin 2001). In other words, despite the impression of democracy, 'Citizen Science' is not about abandoning the notion of expertise, or even science. Instead, as Irwin himself writes (citing Barry Barnes 1985:110):

Our focus should not solely be on a critique of science in itself, but also on the current framework of knowledge relations which gives ascendancy to some knowledge forms over others: 'the power to determine which expert is believed is the important form of power, not the power of experts themselves' (Irwin 1995:170).

The practical challenge of Citizen Science is thus very similar to the one we have addressed in this volume -- expertise per se is not the problem. Indeed, expert knowledge is often an important part of the solution. The problem is how to identify which experts are needed and to find ways of giving them an appropriate input into decision making. To the extent that this changes institutional practices then it necessarily involves changing the power relations through which scientific expertise has been traditionally privileged.

from understanding a scientific dispute.¹⁵ The question, then, even for those who read the journals in which primary research findings are published, is whether their knowledge matches the Trivial Pursuit player's, the chess novice's, the experienced chess player's, or the chess master's understanding of the bishop's move. Our claim is that in the case scientific disputes primary source knowledge is close to the chess novice's understanding.

EXPERTISES THAT INVOLVE SPECIALIST TACIT KNOWLEDGE: CONTRIBUTORY EXPERTISE AND INTERACTIONAL EXPERTISE

The most important transformation of our ideas about expertise over the last half-century has been the move away from seeing knowledge and ability as quasi-logical or mathematical toward a more wisdom or competence-based model.¹⁶ As we have already intimated, we now see expertise more and more as something practical -- something based in what you can do, rather than what you can calculate or learn. This shift has been in part inspired by ideas coming from phenomenological philosophers such as Heidegger and Merleau-Ponty while for sociologists of science the main influence has been Wittgenstein's idea that the meaning of a concept can be understood only through its use; it is the use of a concept that establishes its meaning, rather than any kind of logical analysis or a dictionary definition.¹⁷ The Wittgensteinian frame of mind (as interpreted here), leads us to expect to find specialist knowledge located in specialists' practices rather than in books. Mastering a tacit-knowledge-laden specialism to a high level of expertise, whether it is car-driving or physics, ought, then, to be like learning a natural language -- something attained by interactive immersion in the way of life of the language speaking culture rather than by extended study of dictionaries and grammars. The first three categories of expertise, beer-mat knowledge, public understanding, and primary source knowledge, might be said hardly to enter the category of specialist expertise at all because they do not involve much in the way of mastering the tacit knowledge belonging to the subject matter of the domains; acquisition of the first three kinds of knowledge (though it depends on ubiquitous expertises), involves learning or reading rather than immersion in the culture of those who have the knowledge. 'Enculturation' is the only way to

¹⁵ Thus, there are published physics papers, making potentially momentous claims, that are known by the initiated to be of no scientific importance. (See, for example, Collins 1999; 2004b)

¹⁶ Which is one reason why Goldman's (2001) analysis of expertise is disappointing.

¹⁷ Wittgenstein 1953. Wittgenstein's writings are somewhat aphoristic and open to many interpretations. The interpretation adopted here is that of Winch 1958, and also coincides with Bloor's 1973 and 1983. The most well-known contemporary writer about expertise is Hubert Dreyfus (eg 1972, 1992) who takes his inspiration from Heidegger.

master an expertise which is deeply tacit knowledge-laden because it is only through common practice with others that the rules that cannot be written down can come to be understood.

We now take another step on the ladder of expertise. We step up to expertises that depend on the acquisition of new areas of tacit knowledge; expertises that can only be mastered through enculturation. These, as we have explained, range from some specialist expertises that are very widespread, such as car-driving, to the esoteric technical specialisms.¹⁸

What is new about our analysis is that we split this kind of expertise into two. The traditional category of ability to perform a skilled practice we call 'contributory expertise.' Contributory expertise, as its name suggests, enables those who have acquired it to contribute to the domain to which the expertise pertains: contributory experts have the ability to do things within the domain of expertise. There is little need to say much about contributory expertise since it is the most common usage of the word expertise. Much more interesting and important from the point of view of this analysis, however, is an overlooked second type of deeply tacit-knowledge-laden expertise: 'interactional expertise.' This is expertise in the language of a specialism in the absence of expertise in its practice. This may seem a strange notion given all that we have just said about the importance of practice – of doing things -- but we must look more deeply.

Why has interactional expertise been overlooked?

To simplify, as we have already intimated, within the existing academic literature analysts tend to think of knowledge as of two kinds: the formal or propositional on the one hand and the informal or tacit or the other. The formal can be expressed in rules, formulae, and facts and can be encapsulated in computer programs and the like. The informal or tacit, in so far as it is also rule-like, comes in the form of rules that cannot be explicated and are known only through their expression in action and through the fact that it is easy to see when they have been broken by those who have come to understand them by being enculturated into the form-of-life that encapsulates them.¹⁹ The perennial question that emerged with particular clarity in the debate

¹⁸ Hubert Dreyfus (see previous footnote) along with his brother Stuart (Dreyfus and Dreyfus 1986) have developed an influential fivefold developmental-stage theory of expertise which describes the way an individual comes first to learn the explicit rules of a domain which takes him or her to the level of novice but does not reach full expertise until new and unspoken rules have been 'internalised.' The Dreyfus's scheme has something in common with our classification of specialist expertises though they have no concept of interactional expertise (see below) and they take no interest in the way these different levels of expertise play out in decision-making nor in social settings in general. Later in the book we will describe the Dreyfus and Dreyfus scheme and its place in a more universal classification of expertises.

¹⁹ For example, I may not be able to say what the rules for proximity to others are in various societies but a little habituation and I will be able to accomplish them and I will also soon be made to know if I break

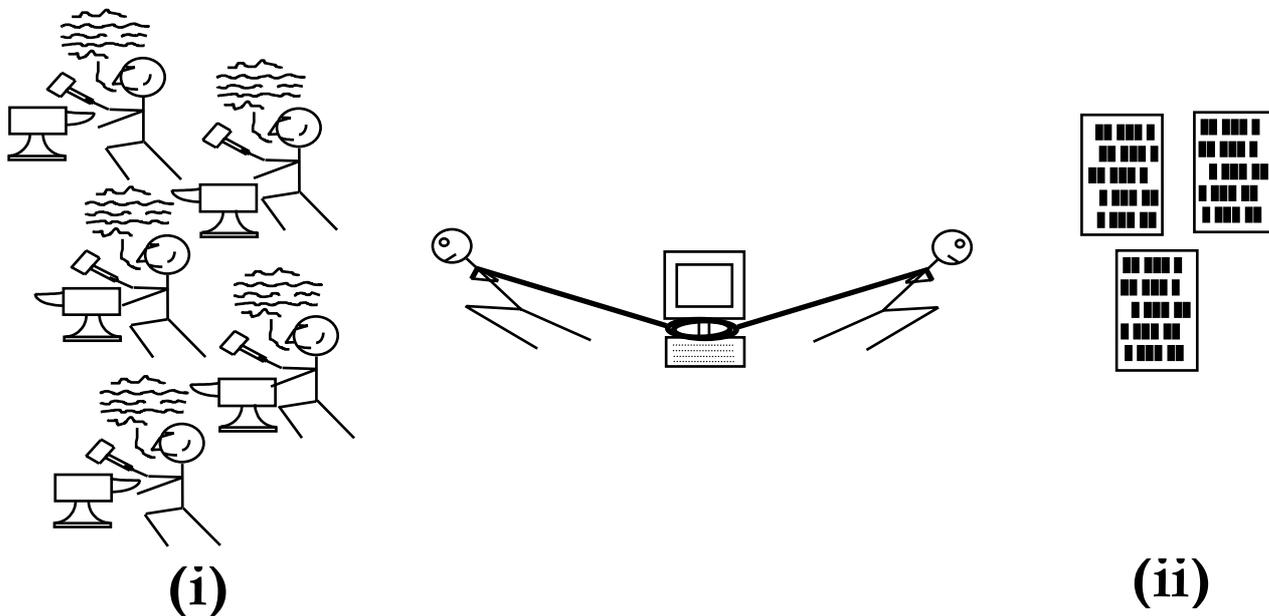
about 'artificial intelligence' is whether the informal can be reproduced by sets of formal rules if the set is sufficiently complex and ramified: 'If a computer program is complex enough can it mimic human action?' This question has tended to polarise analysts.

To put this another way, language, whether natural language of the language pertaining to a specialist domain, has been treated in one of two exclusive ways in the debate.

(i) Full immersion in an entire 'form-of-life' would be needed to master a language (this is the 'informal' view).

(ii) mastering the language pertaining to a domain comprises no more than the acquisition of propositional knowledge -- a set of formal rules and facts gained through reading and instruction (this is the formal' view).

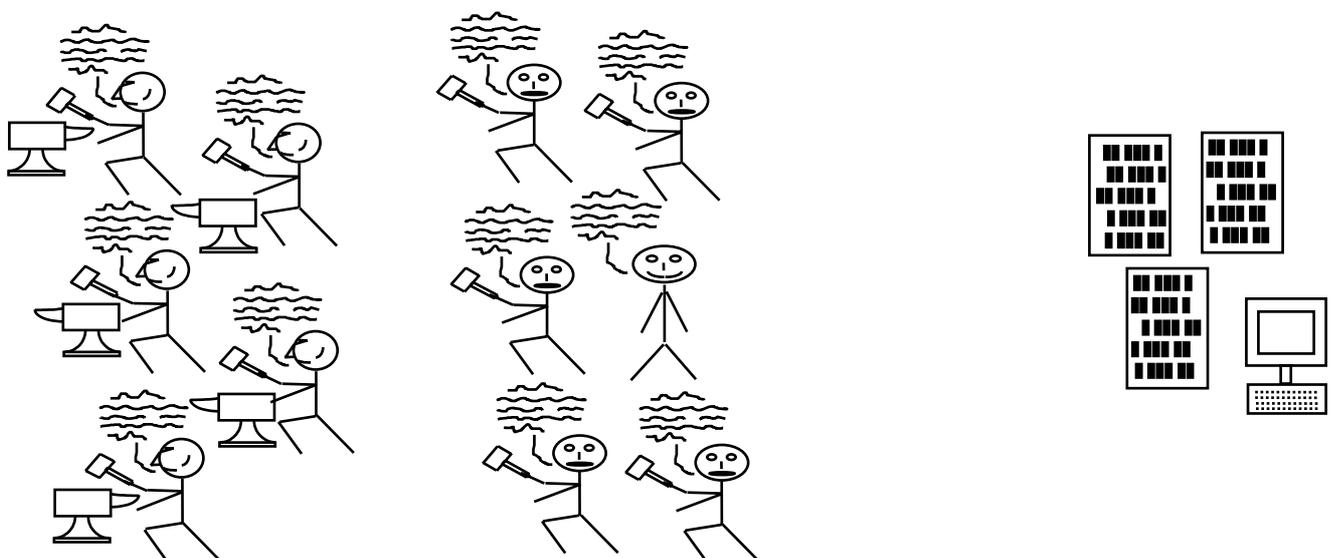
The second of these possibilities -- that of the formalists,' has, we believe, been shown to be wrong by both theoretical analysis and, as intimated above, by the long-running experiment with intelligent computers: the failure of artificial intelligence has shown (ii) to be wrong and that the language of a domain, like any spoken language, consists of more than propositional knowledge. The problem this has left us with is that any attempt to claim that language can be mastered outside of a full-blown practical immersion in a form of life is thought, by informalists, to amount to the claim that (ii) is true since the only possibilities 'on the table,' as it were, are (i) or (ii). This excludes any kind of immersion within a domain that is short of full-blown practical immersion. The awkwardness of the notion of interactional expertise comes from the fact that it stands between (i) and (ii).



them, and can easily recognise if someone in my own society breaks them (say, by standing too close to me).

The cartoon illustrates the point. On the left we have a group engaged in practical activity and discourse of a form of life. On the right we have a description of their activities in propositional form -- in books, journal articles, and the like. In the middle we have a computer, the subject of an intellectual tug-of-war between the 'formalists,' who believe it can be elaborated to the point where it will fit indistinguishably into the left hand group, and the 'informalists,' who believe it will never progress beyond, essentially, a set of propositions or symbols whose natural home is with the books on the right. The existence of tug-of-war means that the efforts and the gaze of those who think about these matters are directed exclusively one way or the other; no one is paying attention to the space in the middle between the skilled group and the books. Interactional expertise is found in this middle ground.

Interactional expertise is, however, nearer to (i) than to (ii). Interactional expertise is definitely not a set of propositions. Interactional expertise is mastery of the language of a domain and mastery of any language, naturally occurring or specialist, requires enculturation within a linguistic community. Interactional expertise is an 'embodied' expertise and cannot be expressed in propositional terms. The computer, no current of foreseeable model of which can be immersed within a language community in a way that will allow for it to become enculturated, will have to be dragged to the right. On the other hand, the idea of interactional expertise still does not amount to (i) -- full-blown immersion in a form life. The idea of interactional expertise implies that complete fluency in the language of a specialist domain can be acquired in the absence of full-blown physical immersion in the form-of-life. Note that a strong claim is being made: the level of fluency in the language of a domain that can be attained by someone who is an interactional expert only is in principle indistinguishable from that which can be attained by a full-blown contributory expert. In a later chapter, and in Collins, 2004a, (see link on this website) we investigate interactional expertise, and this strong claim in particular, in a more technical/philosophical way. For now let us provide another cartoon to make the position being argued crystal clear.



This cartoon shows the computer in its proper place while interactional expertise occupies the left centre of the middle ground. That the experts are engaged in interactional expertise is indicated by taking away their anvils so their hammers merely identify them as contributory experts engaged, for the time being, in the discourse of their form of life rather than the practical activity. In the middle of those with the hammers is a stick figure, who might be a sociologist, grinning happily having mastered the interactional expertise while never having worked with hammer and anvil. As can be seen, the discourse of the smiling interactional-but-not-contributory expert is indistinguishable from the discourse of those with the hammers. That is the claim, deliberately expressed as provocatively as possible, so as to give rise to a serious debate which is begun in Collins 2004a²⁰.

Origins of interactional expertise

Though it has wide application, the idea of interactional expertise emerges initially from the experience of the sociological fieldworker. Typically, sociologists who want to study areas of scientific knowledge that are new to them have to try to grasp something of the science itself. The sociologist begins with no specialist expertise -- which is a level insufficient to do sociological analysis of scientific knowledge. The sociologist is likely to move rapidly through public understanding and primary source knowledge, which are also inadequate to allow for competent social analysis of scientific knowledge. With luck, however, interactional expertise, which does allow for social analysis of scientific knowledge, will eventually be achieved.

The transition to interactional expertise is accomplished, crucially, by talking to the experts.²¹ Interactional expertise is slowly gained with more and more discussion of the science (or other technical

²⁰ See also Selinger and Mix, 2004; Collins 2004c.

²¹ In rare cases the sociologist might even progress to the level of contributory expertise. Contributory expertise can be attained only by practising the science and there is rarely an opportunity to do this if the sociologist has not undergone the full-scale training in professional institutions that are the pre-requisite to certification. It is not impossible, however, in sciences which are not too difficult as in Collins's contributions to parapsychological research (see, Pamplin and Collins, 1975). Also, some degree of contributory expertise may be attained where the science is not too far removed from the sociology as in the case of artificial intelligence. Collins would claim that his books on artificial intelligence (1990; Collins and Kusch, 1998), make contributions to the field itself as well as being sociological analyses. (Whether those contributions have been taken up is another matter.) Simon Cole (eg, 2001 and see later in the book), also claims to have made contributions to the special field he studies -- fingerprinting. Some sociologists, as we will discuss in what follows, have contributory expertise in the corresponding sciences before they start to study them as social scientists.

skill).²² Interactional expertise cannot always be attained -- the science may simply be beyond the capacity of the analyst. One of the authors, Collins, discovered this when he attempted to do research in the field of the theory of amorphous semi-conductors. After completing about 13 hours of taped interviews with scientists he had to conclude that he could not understand enough of the science to reach a sufficient level of comprehension of the scientists' world to make any sociological headway. He had to give up. One characteristic of such a failure is that each new interview or discussion would start with a long and tedious period of explanation of how the science worked which repeated, approximately, the explanation that had marked the start of every other interview or discussion. All parties were equally bored by these explanatory sessions and there the interviews went little further.

In contrast, where interactional expertise is being acquired there will be a progression from 'interview' to 'discussion' to 'conversation' as more and more of the science is understood. There is no sudden 'aha moment' which marks the switch to mastery of interactional expertise but its steady acquisition can nevertheless be recognised. Above all, with interactional expertise conversation about technical matters has a normal lively tone and neither party is bored. As things develop the day may arrive when, in response to a technical query, a respondent will reply 'I had not thought about that,' and pause before providing an answer to the sociologist's technical question. When this stage is reached respondents will start to be happy to talk about physics and even give studied consideration to critical comments. Eventually respondents will become interested in what the analyst knows about the field because he or she will be able to convey the scientific thoughts and activities of others in a useful way. The sociologist who has just come from visiting scientist X may be able to tell scientist Y something of the science that X is doing or the kind of thinking that X is engaged in in respect of some common problem. Occasionally the analyst will be able to explain the scientific position of another party in a clearer way than the scientist him or herself currently understands it; this is because the analyst has heard the position explained at great length whereas the scientists may find communication with an academic rival difficult because of their different commitments and interests. By this stage, what were once 'interviews' have become 'conversations,' not markedly dissimilar to the conversations the analyst will have with social scientist colleagues and, presumably, not that different to the conversations that one scientist will have with another. (For the sociologist sitting in on the conversations that scientists are having with one another no longer seems like eavesdropping so much as participating.) In sociologist-scientist conversation of this kind both parties can speed things along by anticipating a technical point so that a longer explanation is avoided when existing mutual understanding is indicated by an interjection. In the same way the conversational partner's expression of a point may be helped, or a memory jogged, by a phrase which anticipates what is to come. At this stage jokes, irony, and leg-pullings are all recognised. At this stage

²² It might be sheep farming.

respondents will no longer be tempted to give a 'pat' or formulaic answer to a question drawn from a ready-made set of responses which represent the canonical face of science. Respondents mostly will talk to the analyst as they would talk to a colleague rather than an outsider, knowing that the standard recipe will not do. If, however, a respondent is encountered who does not know the analyst well, and is tempted to provide 'officially approved' answers to questions, the analyst will have the skill to recognise the nature of the response and discount it or probe further; a sharp technical remark by the analyst can speedily change the whole tenor of such a conversation. As things go on the analyst may even develop the confidence to take a 'devil's advocate' role in respect of some scientific controversy and argue a scientific case with which the respondent disagrees. The counter-case may be maintained well enough by the analyst to make the respondent think hard.

Where there is no developing interactional expertise, as in Collins's experience in the case of the theory of amorphous semi-conductors, the conversations never become interesting to either party, the analyst can never transmit information, take a devil's advocate position or, crucially, distinguish between 'pat' answers and real conversational interchange, nor between jokes and irony on the one hand and serious responses on the other. Worse still, though a field might be riven with controversy (as was the theory of amorphous semi-conductors), the analyst cannot understand what the protagonists are disagreeing about, nor how deep the disagreements run, nor, with any certainty, who disagrees with who! The contrast between these extremes -- no expertise and a good level of interactional expertise are very marked and quite unmistakable -- at least by the fieldworker who has experienced both.

In spite of gaining very high levels of interactional expertise, to the extent of fulfilling a useful minor role in the transmission of scientific information among the scientists, or occasionally giving a clear explanation to one party of the scientific position of another, the analyst is not going to be given a job or let loose in a scientific laboratory in the absence of contributory expertise. The analyst who has even the highest levels of interactional expertise may be able to understand scientific things, and to discuss scientific things but is still not able to do scientific things.

We can be fairly sure that another difference between interactional and contributory expertise is that contributory expertise is self-sustaining whereas interactional expertise is not. That is to say, a contributory expertise -- such as gravitational wave physics -- can be taught to new recruits and is passed on from generation to generation by apprenticeship and socialisation; someone who has the contributory expertise can pass it to someone who does not have it. It is not at all clear that the same applies to interactional expertise. It is not at all clear that interactional expertise, which is always interactional expertise in another expertise, can be passed from one person or generation to another (in the absence of contributory expertise). Interactional expertise in a specialism seems to be learned exclusively through interaction with people who have contributory expertise in that specialism, not interactional expertise in that specialism. There would be something very odd about someone with interactional expertise in, say, gravitational wave physics, trying to teach interactional expertise in gravitational wave physics to someone who never themselves spoke to

gravitational wave physicists. One would guess that the nature of the expertise would rapidly become distorted as messages are distorted when they are passed on by word of mouth through many intermediaries.²³ The point is that interactional expertise is skill in speaking a specialist language and the nature of language is a function of the whole environment, physical and social, in which it develops. Change the environment (eg remove the physical activity which is initially an integral part of the development of a language), and the language will change. But this does not mean that an individual immersed in the linguistic community cannot learn the language without being engaged physically with the physical world that gave rise to it.²⁴

THE RELATIONSHIP OF THE SPECIALIST EXPERTISES

This completes our initial five-step ladder of expertise. It starts from beer-mat knowledge, and goes to public understanding and primary-source knowledge, all of which turn on ubiquitous expertises only. Then it makes the transition to expertises involving specialist tacit knowledge, the first of these steps being interactional expertise and the second being contributory expertise.

There is a transitive relationship between the five levels of the ladder. If you possess one of the higher levels you will possess, in principle all of the lower levels but not vice-versa. There are, however, a few practical exceptions to the transitivity relationship. First, as we will discuss in the next section, a contributory expert's interactional expertise may be latent not realised. Second, contributory experts may know the journal literature only at second hand rather than have the first hand acquaintance of those whose knowledge extends only as far as the primary sources. This is important because it can make primary source knowledge alone very misleading since the peer reviewed journals often contain a model of scientific knowledge that is quite different to that of the true domain expert. For example, one of the authors of this book found during his

²³ As in the game of 'Chinese Whispers' or 'Telephone.' We will return briefly in a later chapter to consider what such a self-sustaining interactional expertise endeavour might look like.

Remarkably, it is possible to find anthropologists or ethnographers who believe that the criterion of a good ethnography is that the reader of it should be able to pass as a native (in our terms, contribute to the expertise). These analysts would think that the passing on of interactional expertise through talk would be a trivial undertaking since the passing on of even contributory expertise through writing is possible. The many studies of tacit knowledge, or attempts to embed human knowledge into computers has shown that this idea is nonsense; the analysis of this book would suggest it is outrageous nonsense.

The fading of a language in the absence of practical usage is the subject of stories such as H. G. Wells's 'The Country of the Blind.' Likewise it reminds one of the language of 'cargo-cults.'

²⁴ The crucial difference between the individual embodiment thesis and the collective embodiment thesis is discussed in Collins 2004a.

fieldwork that he was one of only about three people who actually read what was, on the face of it, a pioneering paper in gravitational wave detection; the main body of scientists had dismissed it out of hand.²⁵ We have to accept that the model of truth represented by the specialist group is superior to that of those who assess the content of the journals in the absence of social contact with specialists. Third, it may well be that specialists in general knowledge quizzes and the like could have a greater width of beer-mat knowledge than a domain specialist.

Hand-in-hand with the transitivity of the specialist expertises goes the transitivity of their pattern of distribution among the population. As we move up the scale from no specialist expertise, through beer mat knowledge, popular understanding, primary source knowledge, interactional and contributory expertises, we find ourselves looking at smaller and smaller groups of people; the expertise becomes more and more esoteric. Popular understanding is limited to the numbers who read popular science books and articles in the science magazines and broadsheet newspapers. Once we get to primary source knowledge we encounter still smaller numbers, who tend to be driven by special health needs, local circumstances, or burning political agendas -- forces which may also lead them to mix in the kinds of scientific circles where they are exposed to a deeper understanding of the issues.²⁶ Those with interactional expertise are fewer in number still since gaining interactional expertise requires crossing social boundaries and spending a long time in alien environments to which there is restricted access. Persons with interactional expertise, therefore, number only slightly more than those with contributory expertise, who, in highly technical sciences, may be limited to somewhere between a half-dozen and a few hundred.

INTERACTIONAL EXPERTISE AND INTERPERSONAL AND REFLECTIVE ABILITY

Interactional expertise is very easy to confuse with other kinds of capacity that are part-and-parcel of the job of the sociologist, journalist, art critic, architect, and so forth. All these professionals need the ability to interact with other people, to talk fluently about the domain which they have chosen to study or within which they exercise their judgment, and to reflect upon their subject matter so as to articulate their findings or judgements. These are capacities not necessarily shared by those with contributory expertise in the domain and this raises a question about the transitivity of the relationship between contributory and interactional expertise. We claim that if you have contributory expertise you also have interactional expertise but if one

²⁵ See Collins, 2004b

²⁶ There are many examples in the literature in which it can be seen how the key citizen activists are driven by some combination of these motives and interests. For example, many of the case studies contained in Irwin and Wynne (1996) have this quality, as do the studies of RSI patients (Arksey 1998); AIDS treatment activists (Epstein 1995, 1996) and nuclear protestors (Welsh 2000).

does not have a ready ability to talk and reflect then one is likely to have little in the way of interaction with others in respect of the expertise. As intimated, the resolution is to say that in the absence of the other kinds of capacities the interactional expertise of the contributory expert will be latent rather than expressed.²⁷ That there must be a difference between latent interactional expertise and an absence of interactional expertise is easy to see: one could, at least in principle, reveal the interactional expertise of an inarticulate and non-reflective contributory expert by skilled and persistent probing -- from skilled interviewing you could learn something about the domain (this is what sociologists and journalists do in the case of inarticulate and unreflective respondents). In contrast, no amount of probing will extract deep information about a domain from someone with neither contributory nor interactional expertise.

We will give the labels 'interactive ability' and 'reflective ability' to the capacities that turn latent interactional expertise into expressed interactional expertise. (These are the 'dispositions' found in row 2 of the periodic table.)

Interactive ability

To repeat, possession of contributory expertise guarantees possession of at least latent interactional expertise. To realise the interactional expertise it is also necessary to possess interactive ability.

A lack of interactional expertise combined with a high level of contributory expertise is very typically exhibited by many fine artists who consider that their work must 'speak for itself.' 'If the meaning of a painting could be expressed in words there would be no point in painting,' as they might and do say. They make the point in practice by refusing to speak fluently about their work and allowing their reflective discourse to atrophy.

On the other hand, as explained, in the role of art critic, journalist, sales representative, television or radio interviewer, and interpretative sociologist, the skills of interacting smoothly and fluently with others is crucial. Without these skills the job cannot be done. In these roles, a high level of interactive ability, is part of the contributory expertise pertaining to that particular specialism (though the specialism itself may be almost entirely devoted to gaining interactional expertise in other specialisms).²⁸

²⁷ Evan Selinger tells us that you cannot have a latent expertise; a latent expertise is no expertise at all. So let us define the term 'latent' in this context as 'having the potential to develop interactional expertise very rapidly as a consequence of long practical experience in a domain.' We will return to this point in due course.

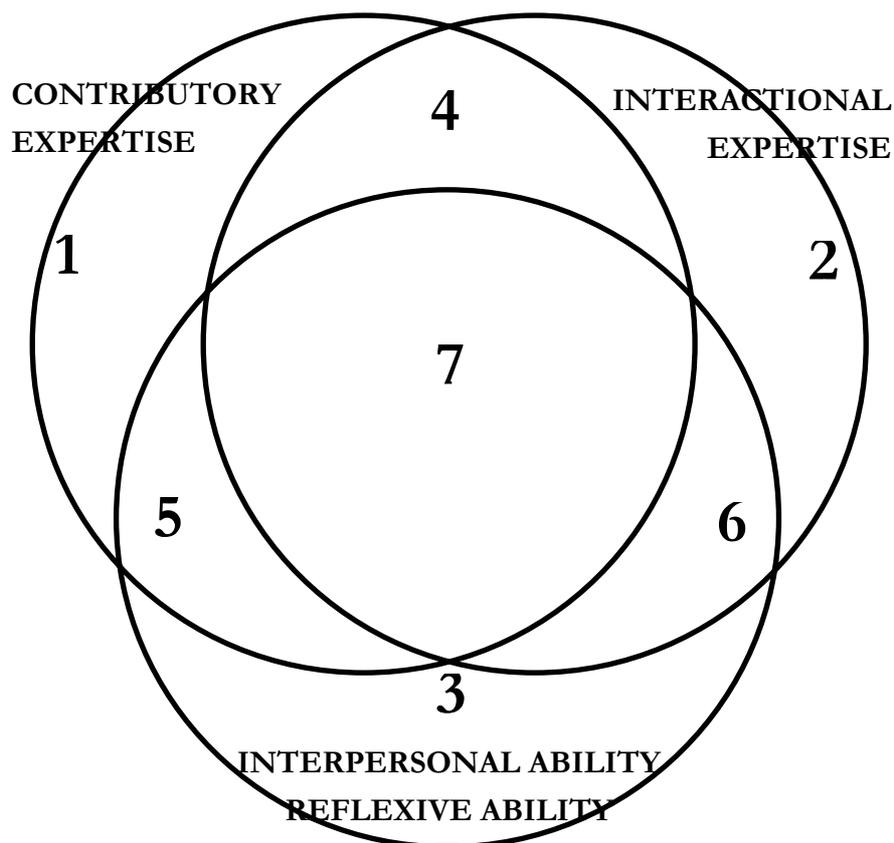
²⁸ In logic, though presumably not in practice, a person could gain interactional expertise in interactive ability. That is to say, a person could spend their time mixing with lots of people who have good social interaction skills, such as journalists, without trying to develop the ability themselves (without becoming a journalist).

An important difference between interactional expertise and interactive ability is that interactive ability, unlike interactional expertise, can be passed from generation to generation. Interactive ability is, as we will call it, a 'disposition,' like kindness, or a loving nature, or a gift for observation, rather than a specialist skill. For example, parents who have the 'gift of the gab' are likely to pass this ability on to their children. The point is that interpersonal skills are generalised abilities not an expertise in a special domain. It is because interactional expertise is expertise in something that it is unlikely that such an expertise could be passed on in the absence of continued contact with the 'something.' To repeat, one cannot imagine that interactional expertise would do other than die out if not refreshed from time-to-time by contact with those actually doing the thing -- the contributory experts. It is the contributory experts not the interactional experts who define and develop the content of the language that the interactional expert tries to master.

Reflective ability

Another generalised skill, which, with a little stretching, can also come under the heading of a disposition, is reflective ability. This, it is true, is a more professionalised and specialised ability than interactive ability because it is taught, quite self-consciously, in sociology and philosophy courses and the other critical disciplines. Like interactional ability, reflective ability is enormously useful in the building of interactional expertise. It is sometimes the only thing that marks out the difference between the analyst and the practitioner of a scientific specialism. Indeed, some scientists are proud of their lack of reflective instincts, boasting that 'philosophy is no more use to science than ornithology is to birds.' This is perfectly correct but carries the corollary that it is not birds (scientists) who one should consult to learn about flight (science). Reflective ability, like interactive ability, can exist sui-generis and be passed from generation to generation. Reflective ability is not reflective ability in something, it is just reflective ability. Reflective ability is, again, part of the contributory expertise of the social analyst of science, the art critic, and so forth.

We can assemble some of these relationships in diagrammatic form.



The large circles in the figure represent the concepts of contributory expertise, interactional expertise, and a combination of interpersonal and reflective ability respectively (combined for simplicity's sake). If, as we say, contributory and interactional expertise are related transitively, then spaces 1 and 5 are empty of people: those who possess contributory expertise will also possess interactional expertise, either latent (located in space 4), or realised (located in space 7). In practice anyone who has acquired interactional expertise without acquiring contributory expertise is likely to be a member of one of those professions that turn on interactive ability and will therefore be found in space 6 rather than space 2. Thus, space 2 is likely to be empirically empty if not quite as logically empty as spaces 1 and 5. Space 7 contains those few social analysts who are also technically competent as well as articulate scientists and technologists. Space 3 is occupied, among others, by philosophers of science and sociologists of science (as opposed to sociologists of scientific knowledge), who do not need interactional expertise for their style of work. In practice that group is likely to stress reflective ability rather than interactive ability.

PROBLEMS OF CATEGORISATION

We have embarked on a categorisation of expertise and, no doubt, many readers will already be finding fault with it. But all categorisations of expertise are going to be flawed -- for example, there will always be boundary problems. One reason is that, like any other categorisation, one has to deal with 'ideal types.' There will always be cases where one kind of expertise shades into another. Another reason is that experts often express their objections to a rival's conclusions by questioning their expertise. To give an example close to home, 'science warriors' often say that sociologists' analyses of science are flawed because they do not have enough expertise in science.²⁹ These problems must be taken seriously but not to the point of academic paralysis. Social scientists should not aspire to a greater degree of perfection than the scientists they describe. Just as in natural science, many of the flaws in social scientific work have to be ignored if distance is to be allowed to work its enchantment which it must if new knowledge is ever to be generated. This, of course, is not a way of avoiding assiduous critical scrutiny of our categories -- no doubt there are better ways to do it but we have made a start. The point is to understand the need for a table of this kind whether it is this one or another.

²⁹ For typical work by 'science warriors' see Social Studies of Science Vol 29, No 2 and Dawkins, 1999; Gross and Levitt 1994; Gross, Levitt and Lewis 1996; Koertge, 2000; Wolpert 1992

DRAFT CHAPTER 2

META-EXPERTISES: EXPERTISES FOR JUDGING OTHER EXPERTISES

In the last chapter we set out a number of levels of expertise. Decision-making is about judging expertise -- judging which person is the better expert or which conclusion represents the most expert input. We are going to have to say something about the resources used by an expert for judging experts and expertise. But in the course of this exercise we are not going to provide formulae for actually making such judgements. Such judgements can only be made at the time by the experts in question. Studies of expertise are about setting some boundaries on how expertise, including the judging of expertise, is and should be exercised, it is not actually doing things or making judgements. The normative work that we are trying to do is to define the broad class of experts and expertise, *inter alia*, defining the broad class of non-expertise and non-experts (in particular the general public when it comes to technical expertise). To repeat, this is not the same as actually making technical decisions or even helping the class of experts in their thinking except in so far as excluding certain classes of judge from a decision contributes to a judgement.

Now we can turn to expertises used to judge other expertises. There are two kinds: external expertises, which do not turn on acquisition of the expertise itself, and internal expertises, which do involve an intimate acquaintance with the substance of the expertise being judged.

EXTERNAL JUDGEMENTS

Ubiquitous discrimination

We can make a start on defining the class of those who are in a position to judge experts and expertises by noting that in respect of some kinds of judgement the boundaries do extend to the very edge of the general public. Just as there is ubiquitous expertise (described at the beginning of the last chapter), there is ubiquitous meta-expertise, which we will call 'ubiquitous discrimination.' In such a case, once more the Problem of Extension disappears. Ubiquitous discrimination, like other ubiquitous expertises, is acquired as part-and-parcel of living in our society in the way natural language fluency is acquired.

For example, those with little scientific knowledge can sometimes make what amounts to a technical judgement on the basis of their deeper social understanding. The judgement turns on whether the author of a scientific claim has the appropriate scientific demeanour and/or the appropriate location within the social networks of scientists and/or sufficient distance from the interests served by the claim.

Ubiquitous discrimination is what we have all been learning since we could speak and is just a particular application of our judgements about, friends, acquaintances, neighbours, relations, politicians,

salespersons, and strangers, applied to science and scientists. For example, most members of a society, just by being members of that society, are able to discriminate between what counts as science and what counts as non-science. This is the ubiquitous judgement on which we rely when we dismiss certain fringe beliefs such as astrology from the list of contributors to the scientific element in technical decision-making. Most members of our society have sufficient judgement to know that the standards, and the social and cognitive networks of astrologers do not overlap with the standards and social and cognitive networks of scientists.³⁰

Another illustration of this point emerges from the dispute over the Moon landings. There is a view that the Moon landings were faked by the Americans, the events being filmed somewhere in a desert in the Western USA.³¹ The groups who believe this cite various anomalies in the films, such as the quality of the shadows, the way the flag flapped, or some such. Technically we are in no position to judge whether the films were real; certainly it would be quite possible with the film technology of the time to have faked the whole thing. If we turn to those with more than the average level of technical expertise we find that even there fakery of this sort is not easy to rule out. For example, it is reported that some American astronauts and other technical people disputed the filmed evidence of the first Russian 'space walk' (EVA) in 1965. As the astronaut David Scott reports his reaction at the time:

I was on my feet pacing by this point. 'If this EVA is real, they're not only ahead but pretty far ahead, at that. What proof do we have that this guy really went outside?'

The first grainy photographs released to the world press of Alexei Leonov floating in space sparked a heated debate in the West. Some claimed the photos were faked. They simply would not accept that the Russians had chalked up another first. (Scott and Leonov, 2004, p124) [Scott, David and Leonov, Alexei, (2004) Two Sides of the Moon: Our Story of the Cold War Space Race. London: Simon and Schuster.]

Given astronauts' ability to doubt the Russian space walk there is nothing technical to stop ordinary people doubting the story of the Moon landings.

³⁰ Poor social judgements are the problem with those who believe in, say, newspaper astrology as a scientific theory. They are making a social mistake -- they do not know the locations in our society in which trustworthy expertise in respect of the influence of the stars and planets on our lives is to be found. In a later part of the book we will develop these ideas further using Wittgenstein's notion of 'family resemblance' as the basis of our theory of what should count as a fringe science and what sciences should count as continuous with Western science. We will compare our treatment of astrology and the like with that of Turner (2001).

³¹ We are grateful to Susan Carter for helping to develop this line in respect of the Moon landings.

What stops us doubting their validity is, once more, ubiquitous discrimination which turns on social judgement. It is beyond the bounds of sociological credibility, even ordinary people's sociological credibility, that the thousands of people involved in the Moon missions could all have been organised to lie so constantly and consistently; we know that if there were any possible credence to the story, the Russians, deeply involved in the cold war as they were, would have exploited the doubts -- yet they did not. Our social discrimination allows us to be sure about the Moon landings even while the technical discrimination of even those who were fairly close to the events, does not.

To see how this kind of discrimination works in a more difficult case, consider cold fusion. Most reasonably literate members of this society, who know nothing of cold fusion beyond what they have seen on the news or read in the newspapers, 'know' that cold-fusion has been tried and found wanting. Though there was a time when cold fusion was contiguous with science as we know it, it is now understood that its cognitive and social networks no longer overlap with those of legitimate scientific society. This knowledge has nothing to do with scientific competence. On the contrary, it is vital to ignore scientific credentials and even track-records of success if a sound judgement is to be made.³² Thus Martin Fleischman, the co-founder of the cold fusion field, has an enviable track record for success in the sciences, is immensely well-qualified, and has both interactional and contributory expertise in cold fusion yet still believes in the effect, whereas the consensus is that he is wrong. What people in Western societies have in common is what they have heard about cold fusion in the broadcast media (popular understanding). Their consensual view, in so far as they have one, emerges from making of social judgements about who ought to be agreed with, not scientific judgements about what ought to be believed. To expect the citizen to be sufficiently educated in science as to be able to make scientific judgements in disputed esoteric matters like this is the completely crazy aim of the old COPUS-like organizations. (Because 'distance lends enchantment' they will, however, find many who will enthusiastically endorse their aims -- ordinary citizens who think that having read a popular book or two they now understand cold-fusion better than a Fellow of the Royal Society with a lifetime of relevant practical experience, such as Martin Fleischman.³³) The crucial judgement, however, is to 'know' when the mainstream community of scientists has reached a level of social consensus that, for all practical purposes, cannot be gainsaid in spite of the determined opposition of a group of experienced scientists who know far more about the science than the person making the judgement. This ability is gained through membership of what The Guardian newspaper calls the 'chattering classes.' Note that this is not the sort of judgement we would expect even an immaculately qualified scientist from 'another planet' where the Guardian does not circulate, to be able to make. A scientist from another planet, reading published papers for and against cold fusion, would

³² cf Goldman's theory of success track records

³³ Which is, as we will see, the flaw in the 'deficit model' of public understanding of science.

have a difficult decision to make; the scientifically ignorant citizen of this planet, in contrast, has an easy decision to make.³⁴

A rather different kind of ubiquitous discrimination, based on demeanour and assessment of the balance of interests, was what led the general public to distrust Sir Walter Marshall's account of the British Nuclear Fuels train-crash test of a nuclear fuel flask which we is discussed on Collins 1988. Lord Marshall just did not come across well and his position seemed to be too closely tied to the commercial interests of a particular industry. Occasionally, it may also be possible to spot inconsistencies in a position, the kind of inconsistencies that political interviewers are good at drawing out.³⁵

Local discrimination

There is a quite different version of discrimination that pertains to specialist groups. At this point we need to introduce one of the most influential studies of the nature of expertise, Brian Wynne's examination of events on the Cumbrian fells following the Chernobyl nuclear meltdown in 19XX. (Wynne's work has been important in both establishing the idea that technical expertise can be found beyond the normally recognised qualified groups and in giving rise to many of the confusions that reinforce the folk wisdom view.) Wynne looked at the interaction between UK Ministry of Agriculture Food and Fisheries (MAFF) scientists and the Cumbrian sheep farmers after radioactive fallout contaminated their pastures.³⁶ As we will discuss at greater length below, Wynne argued that the expertise of the sheep farmers in respect of sheep ecology should not have been ignored by the scientists but he also found that in addition to this contributory expertise the farmers possessed what we will call 'local discrimination.' Soon after the Second World War, the Windscale-Sellafield nuclear processing plant was built on the Cumbrian fells so the farmers had long experience of the nuclear industry's pronouncements concerning radioactive contamination; they knew that these pronouncements could not be taken at face value. An outsider, with less experience of discussions of radioactive contamination in this particular social and geographical location, would not have judged them with such finesse. Thus the sheep farmers were able to discount statements by the nuclear industry's spokespersons much as the general public discounted Lord Marshall's assurances but in this case it was a result of local experience rather than a more generalised discriminatory ability developed over a lifetime of social and political education.

³⁴ See Collins, 1999 for a similar argument in respect of the rejection of claims about the existence of gravitational waves.

³⁵ See also Goldman 2001.

³⁶ Wynne 1989, 1996a, 1996b

Another study by Wynne reinforces both the point and the confusion between local and ubiquitous discrimination. Wynne describes the experience of apprentices working in the radioactive materials industry. He suggests that the apprentices felt they had no need to contribute to their own safety by trying to understand the science of radioactivity because they were 'intuitively competent sociologists' and 'vigilant and active seekers of knowledge ... tacitly and intuitively, positioning themselves, using their knowledge of their social relationships and institutions.'³⁷ Wynne argues that the apprentices used their social understanding as a basis for trust in their employers. In a later article, referring to the same group, he says that these apprentices' 'technical ignorance was a function of social intelligence.'³⁸

There are two ways of looking at Wynne's discussion of the apprentices. It could be an example of local discrimination. In that case the apprentices would be seen as using their hard-won specialist competence in understanding the trustworthiness of their particular employers and their own place within the social networks of trust operating in that particular workplace to assess the safety of their procedures in which they were involved. On the other hand, the apprentices might be doing no more than any of us do when we trust the institutions which surround us. For example, when we put money in the bank we do not say that we have no need to understand economics because we are 'intuitively competent sociologists' and 'vigilant and active seekers of knowledge ... tacitly and intuitively, positioning ourselves, using our knowledge of their social relationships and institutions ...' Again, the language seems quite unnecessarily folksy and romantic. Would we say we are using our social understanding as a basis of trust in the bankers and that our economic ignorance was a function of social intelligence? Well, yes, but it is a passive kind of knowledge -- something that we learn in the same way as we learn a language. It is the kind of 'knowledge' that is the basis of all trust in society. It will differ from society to society and it will occasionally be breached (as when there is a 'run on the banks'). It would be, nevertheless, a ubiquitous expertise and, as such, would not really solve any local and technical problem. What should be clear is that for discrimination to be counted as some kind of esoteric technical expertise it must be local discrimination not ubiquitous discrimination. In the case of the apprentices, if their discrimination was of the ubiquitous kind then no amount of sugaring the pill makes them other than uninformed in respect of their own safety; if it was local discrimination -- they had learned who to trust and who not to trust -- then that is a different matter.

Local discrimination, like ubiquitous discrimination is an expertise for judging other expertise that is external to the expertise being judged because it does not depend on the understanding of the expertise being judged but upon an understanding of the experts. We now move to internal judgement of expertise.

³⁷ Wynne 1992 p 39.

³⁸ Wynne 1993 p 328.

INTERNAL JUDGEMENTS OF EXPERTISE AND THEIR PROBLEMS

What other ways of judging between experts are there? The standard method of choice is by reference to the qualifications of the expert.³⁹ This we now know to be inadequate because it is possible to have expertise, and that includes specialist expertise, in the absence of qualifications. Probably the most well-known study to have established this point is Wynne's study on the Cumbrian sheep farmers mentioned above. Wynne found that the sheep farmers knew a great deal about the ecology of sheep, the prevailing winds, and the behaviour of rainwater on the fells that was relevant to the discussion of how the sheep should be treated so as to minimise the impact of the radioactivity. Unfortunately, the MAFF scientists failed to take the farmers' advice seriously because they were not qualified scientists.

A still better documented case of unqualified person gaining expertise is Stephen Epstein's study of AIDS cure activists in the 1980s.⁴⁰ In 1985 the new drug, 'AZT,' was about to be subjected to double-blind randomised control trials but AIDS sufferers were concerned that too many of them would die before the drug was approved. They therefore began a campaign for the introduction of speedier testing regimes, the relaxation of test protocols, and earlier release of potentially beneficial treatments. At first the suggestions of the activists -- a group whose members' dress codes and presentation-of-self was as far from the world of medical orthodoxy as it was possible to be -- were resisted. Robert Gallo the co-discover of HIV is reported as being initially hostile to the AIDS activists, saying to members of one of their pressure groups -- 'the AIDS Coalition to Unleash Power' or ACT UP:

I don't care if you call it ACT UP, ACT OUT or ACT DOWN, you definitely don't have a scientific understanding of things. (Epstein 116)

The activists, however, undertook an arduous course of self-education and learned the language of medical discourse. They added to a developing understanding of microbiology and statistics their experience of how AIDS sufferers would actually respond to the demands placed upon them by the protocols of randomised control trials. They knew that these demands were unrealistic where death was in constant prospect: the groups regularly smuggled untested cures from Mexico, continued to take other drugs which were banned which had the potential to confound test statistics, and even shared placebos and trial drugs between experimental groups.

³⁹ But bear in mind that even the most well-trained and accredited professional sometimes turn out to be incompetent.

⁴⁰ Epstein, 1996

Eventually, the activists gained so much interactional expertise that, allied with their experience, they were able to make real contributions to the science which were warmly embraced by the scientists. Gallo was to come to say of one of their leaders that he was 'one of the most impressive persons I've ever met in my life, bar none, in any field...I'm not the only one around here who's said we could use him in the labs.' Gallo is also said to have described some activists' scientific knowledge as of and 'unbelievably high' standard. He said 'It's frightening sometimes how much they know and how smart some of them are...' (Epstein, 338). The AIDS activists, though unqualified in any field that bordered on medical science, eventually trained themselves to a point beyond that of the Cumbrian sheep farmers, a point at which the science community took them very seriously indeed.

Any criterion of expertise has to allow groups such as the Cumbrian sheep farmers or the AIDS activists to be included in the category of expert and that is why the criterion of formal qualification or accreditation is too exclusive. We suggest that a more important criterion than qualifications is experience. If there is to be a general criterion of expertise, experience is the leading candidate. The criterion of experience would include the Cumbrian sheep farmers, the AIDS activists, and the like and we will discover more of its advantages in due course.

Note that there is no problem about judging expertise at the lower levels of the ladder. A general knowledge quiz such as Trivial Pursuit can discriminate adequately between levels of beer mat knowledge, while higher level quizzes or examinations can discriminate proficiency in popular understanding or primary source knowledge. Indeed, much of our education system is dedicated to discriminating precisely at these levels and hence the perennial complaints of employers that graduates come to them unfitted for the work place where a different kind of expertise -- the level of expertise associated with doing -- is required. It is at these higher levels of expertise that the problems of judging expertise, both practical and conceptual, arise. We start with practical problems of the kind which are thrown up by those who pass themselves off as experts whereas according to most criteria they are not. We start, then, with hoaxers, frauds, and confidence tricksters.

Hoaxers, frauds and confidence tricksters

There are bogus doctors, bogus lawyers, bogus nurses, and bogus paramedics, bogus gas and electricity meter readers, and bogus traffic police. There has been at least one bogus Oxford don, at least one bogus army general -- who turned out to be a woman posing as a man -- a bogus catholic priest, and bogus CIA men. It seems, then, that people are ready to attribute almost unlimited expertise and authority to almost anyone given the right conditions.

The conundrum has been marvellously exploited in fiction. To pick one example among many, Jerzy Kosinski's book, Being There, subsequently made into a film of the same name, shows an educationally sub-normal, sexually neuter, but well-dressed gardener, the eponymous 'Chauncey Gardner,' rise from

unemployment to become President of the United States complete with a reputation for innovatory sexual adventure.⁴¹ It happens because a band of hangers-on re-interpret his minimal vocabulary as the profound and gnostic discourse of a sage and rake.

Since we are concerned with judgement of expertises, however, we need to ask which kinds of role are more or less difficult to fake -- which kind of expertises is it easy to make a mistake about, which more easy to judge, and why. Thus, in the case of the 'expertise' of, say, lying in bed in the morning there is no expertise to fake so anyone can say they are an expert in it without fear of contradiction: there are no confidence tricksters when it comes to lying in bed. There are, on the other hand, few or no confidence tricksters when it comes to solo violin-playing, at least, no confidence tricks in which a performance with an orchestra playing a well-know piece is part of the scam. In between are all the interesting cases, some more easy to fake than others.

A famous case is that of the trivially simple computer program, ELIZA, which was easily able to fake the expertise of a Rogerian psychotherapist.⁴² More recently, and uncomfortably 'close to home' for many of the readers of this book, there is the case of the Sokal hoax. Alan Sokal submitted a manuscript for publication in the journal Social Text using the stylistic cliches of the semiotic turn in the cultural analysis of science; the journal published the paper only to have it revealed as a hoax, Sokal proclaiming 'the emperor has no clothes.' The idea of a hoax of this type is all in the revelation; if the perpetrator can show it is easy to pretend to have the expertise in question then the expertise is made to look more like lying in bed or Rogerian psychotherapy than solo violin-playing. Sokal's hoax certainly did expose lax editorial practices at Social Text but not much more since hoaxes are not so hard to pull off even in theoretical physics.⁴³

⁴¹ Kosinski 1971.

⁴² To be more exact, ELIZA's mistakes tended to be in its language handling rather than the substantive content of its output. For discussion of this case see Weizenbaum, 1976; Collins, 1990.

⁴³ Sokal's more extended conclusion, and that of many of his admirers, namely there is a large gulf between the integrity of the social and the natural sciences, was not borne out by events. Not much later a number of papers by the Bogdanov brothers on string theory were published in a variety of physics journals and a long argument followed about whether they were genuine or a hoax; not being to pin down whether a paper really is a hoax after extended examination is, perhaps, still more embarrassing than the hoax carried off by Sokal. It is likely that at the cutting edge of all disciplines there are areas where no-one is really sure about what the new conventions should be just as in the case of the avant-garde in the arts.

The original Sokal hoax is Sokal, 1996. For more references see <http://www.physics.nyu.edu/faculty/sokal/#papers>. For the Bogdanov brothers event see <http://math.ucr.edu/home/baez/bogdanov.html>.

Hoaxes and frauds are more easy to carry off than they should be because of the well-know tendency of their targets to `repair' deficiencies in the skill of the perpetrator, especially when they have something to gain by believing in the performance. Part of the skill of the professional con-artist is to make the victim believe that he or she can bring them great financial gains but the principle is universal: in nearly every case of a bogus performance, life for those around the fake will be much more inconvenient if they have carry out a complex investigation and perhaps replace a hitherto trusted colleague with someone else. That is why, even in the case of solo violin-playing, we have to say that the musical piece must be a `standard.' If it is not a standard then the audience, who have paid good time and money to hear a virtuoso performance, will be all too ready to believe that is what they are hearing; they might, for example, believe they are present at a rendition of some avant-garde composition, or a piece of `conceptual art' the non-melodiousness of which is the very point since it asks questions about the meaning of music. The role of avante-garde artist has been nicely satirised in the 1961 comedy film, `The Rebel.'⁴⁴ In the film the comedian Tony Hancock plays an incompetent and untrained `artist' who finds himself sharing a garret-flat in the Bohemian quarter of Paris in the 1950s. As a result of a series of accidents, he is `taken up' by Bohemian society, and for a while his smears are treated as great works of art.⁴⁵ We are amused but not surprised at this outcome because we know that with the avant-gard, as the very name implies, there are no established conventions of artistic practice upon which to base a judgement.⁴⁶ That is also why we, or at least some of us, are relieved when we learn that Picasso was a brilliant artist within the conventions of realist depiction before he began to push forward the

Note that there is a difficulty for those who would want to defend a journal with a `post-modernist mission,' such as Social Text, from the predations of hoaxers such as Sokal. The difficulty is that even to accept that Sokal has transgressed a boundary they have also to accept that there is a difference between the genuine exercise of an expertise and its attribution. They have to agree, then, that even if there is nothing but attribution to everyone else's expertise there is something real about their own. Social scientists are quite ready to make claims to their own expertise in their codes of professional conduct (e.g. ASA 1999, BSA 2002).

⁴⁴ `Call Me a Genius' in the USA.

⁴⁵ `The Rebel' seems at first sight to be an exact parallel to Being There but the difference is that there are kinds of painting that cannot be faked without artistic talent (successful art forgers must be technically accomplished) whereas that may not be true of the US Presidency because of the way the President is embedded in a body of advisors.

⁴⁶ In the analysis of technical decision-making in the public domain a term used for the setting of new conventions is `framing.' The debate about who `frames' technical judgements, which we will discuss in another part of the book, is conceptually homologous with the debate about who sets the standards for the avant-garde in art.

frontiers of art; knowing that he was so talented in ways which are relatively easy to judge we can feel more secure in our appreciation of his less conventional works. In the same way we might like to know whether Tracy Emin and Damien Hirst can draw well: if we knew that we could use their more easily understood skills as a proxy for their talent in a world where conventions provide no scaffolding for judgement.⁴⁷

Interactional expertise as a bridge between the public and expertise

We need now to solve a philosophical problem which lurks beneath our descriptions of judging expertises that lie within accepted conventions -- expertises such as solo violin-playing and realist art. A long tradition of thought holds that it ought not to be possible to judge an expertise without actually having the expertise. Philosophers and anthropologists have claimed that it is not possible even to understand the meaning of an expertise without taking part in the 'form-of-life' (or 'paradigm,' or 'language,' or historical 'epoch'), in which it is embedded.⁴⁸ According to this tradition, without having access to the inside of the form-of-life its categories would appear meaningless; a fortiori even to try to make judgements between what counts as skillful performances within an alien form-of-life would be senseless.

⁴⁷ It is not unreasonable to feel uncomfortable when all the virtuosity lies seems to lie in establishing a new convention rather than in executing a skill because it seems to reduce art to marketing. Thus, we do well to be concerned that the Satchi brothers, who run the UK's most successful advertising agency, are also the country's most successful collectors of avant-garde art and, co-extensively, the most powerful definers of its value. Obviously there is healthy ground for concern about this relationship if one believes that art is more than advertising.

This is also the easily understood and very reasonable reaction of critics to modern studies of science which take themselves too seriously as epistemology rather than methodology. In sociology of scientific knowledge much of what used to be taken to be the exercise of skills within a convention is now understood to be coextensive with the establishment of a convention; this, for example, is the consequence of the experimenter's regress for the use of replicated experiments in the establishment of the existence new phenomena -- those experiments that are counted as well-executed are those experiments that produce what are taken to be valid findings under the new conventions of seeing (Collins, 1985/92). The hostile reaction is quite unreasonable to, say, sociology of scientific knowledge as a methodology, however.

It is, of course, very dangerous when movements in the sociology of science which proclaim that there is nothing beyond the establishment of fashions in science take themselves so seriously that they set out to establish their own position by advertising, the corraling of powerful allies, and the like.

⁴⁸ For example, refer to the works of Wittgenstein, Winch, Kuhn, Dreyfus, Quine, and Collingwood, not to mention the writings of anthropologists and sociologists of scientific knowledge that take their lead from such ideas.

Luckily we already have the bridge we need to span the philosophical chasm between having an expertise and judging it; the bridge is interactional expertise. The gap spanned by interactional expertise is that between full-scale physical immersion in a form-of-life (which gives rise to contributory expertise), and full-scale understanding of the language pertaining to a form-of-life, learned by immersion in the relevant linguistic community alone.⁴⁹ Let us see how it works.

Technical connoisseurship

The exercise of expertise within an established convention is, of course, convention-bound. It is not that, say, realist painting contains within itself some universal standard obvious to all: the conventions of what we understand to be realism in art have had to be established and change from epoch to epoch.⁵⁰ Nevertheless, within a stable convention virtuosity can be recognised.

The ability to recognise skillful practice itself improves through practice and that this is the case is recognised in terms such as 'connoisseurship'.⁵¹ Connoisseurship is a meta-expertise. A connoisseur is,

⁴⁹ The arts already have an answer to what we present as a problem in the hermeneutic tradition; according to that tradition the job of the critic is not to judge the expertise of the artist but to reinterpret it within the current set of conventions (or forge a new set of conventions for its interpretation). Indeed, the author of a work of art may be said to have little influence on the interpretation of the work, a view encapsulated in Roland Barthes' slogan 'the author is dead.' The realist theory of expertise discussed here does not allow for such an easy solution, however, so for us the puzzle remains. We will discuss this issue again in a later part of the book.

⁵⁰ Difficulty of execution of a skill is in the last resort independent of convention. For example, suppose I decide to express my private artistry by peeling an apple in a spiral such as to produce a long unbroken ribbon of apple skin. To make a very long unbroken ribbon (imagine it only a couple of millimeters wide), might take months and months of practice but there is no existing convention in which this expertise would be valued. To develop this skill would be like inventing a 'private language' (though one could imagine an entrepreneur finding a way to have it taken up).

⁵¹ Nelson Goodman's (1969) Languages of Art describes the importance of practice in observing art; it is the promise of future skills honed by practice that enables him to explain why art forgeries, even though they may be good enough to fool a professional, should still be separated from genuine works of art. He claims that observational skill may become more refined in the future so that what once fooled the professional will no longer do so.

Carlo Ginzburg's (1989) paper 'Morelli, Freud and Sherlock Holmes: Clues and Scientific Method' treats connoisseurship as the ability to detect the author of a work of art and associates its method not with the 'Galilean' sciences but with that of Sherlock Holmes, of medicine, and of historical scholarship in that all of them deal with specific instances of events rather than general relationships. Ginzburg's,

according to the Chamber's Dictionary, 'a person with a well-informed knowledge and appreciation,' the dictionary definition tells us that the knowledge and appreciation is usually applied to fine food, wine, or the arts. But connoisseurship -- that is, judgement honed by exercise -- can be applied to all expertises and we describe it when applied to such other expertises as 'technical connoisseurship.'

Consider, for example, that a builder is employed to make major alterations to a house. At various stages, and particularly at the end of the job, it has to be decided whether the work has been finished satisfactorily. Imagine that some tiling has been done in the new bathroom. How even should the tiling be? How clean and square should the grout lines be? When has the job been finished?⁵² One can see immediately that there are conventions that give meaning to bathroom tiling, conventions that would be unknown to someone who had never lived in a society with tiled bathrooms. Some of the conventions can be set out as formal standards. For example, should the tiles be cut round curves with a diamond saw, or cut square to an approximation to the curve, or chipped and hacked, the gaps being filled with grout? Somehow one must 'negotiate' with the builder over what will count as a satisfactory job in terms of both the formalised standards and the unformalised conventions. Interestingly, one may employ a professional -- an architect -- to do the negotiations for you.⁵³ The fact that it is possible to employ a professional who has probably never actually done any tiling to make these judgements shows that the crucial thing is experience within the conventions of judgement rather than experience of the skill itself. There is a connoisseurship of tiling. The judgement being exercised by the architect, or the houseowner, who themselves may not be capable of tiling (who have no contributory expertise), but has seen and discussed many bathrooms, is based on interactional expertise.⁵⁴

however, seems confused. Physical and biological sciences, and the identification of the provenance of paintings, and the method of Sherlock Holmes, are all typically scientific activities since though they deal with specific instances it is specific instances of general laws -- general laws are applied to the specific case which is under examination (just as, say, building a rocket to fly to the Moon applies general physical laws to a specific instance). History, on the other hand, is different, as Popper points out, since the broad flow of history is a one-off event; to believe that science can predict one-off events is what Popper calls 'historicism.'

⁵² We are grateful to Kevin Parry and Mike Bergelin for executing the example.

⁵³ As with the Wittgensteinian description of following a rule -- it is not possible to completely describe following the rule but it is possible to know when the rule has not been followed properly.

⁵⁴ This is not meant to imply that sociologists with interactional expertise should replace peer reviewers in the scientific fields they study -- far from it -- the sociologist's experience is merely used to prize open a conceptual space. Nevertheless the discussion does open the possibility that there could be opportunities for people in positions such as that of sociologists to contribute to scientific judgements. (see also Collins 2004b, Chapter 42)

Just as the general public cannot have expertise in all domains of specialisation neither can any single person. There may be some who would claim that refined judgement in all things is the inheritance of the members of a well-bred aristocracy but if the aristocracy do have special qualities in this regard it is actually a matter of training in the specific domains of food, wine, or art. When they take it to be their birthright to extend that refinement of taste to more technical domains things usually go wrong.⁵⁵ Nevertheless, it is at least possible to judge and expertise without being able to practice it.⁵⁶

We can now see more clearly why it was that in order to be sure to recognise that a solo violin-player was a fraud it would have to be a familiar piece of music that was being performed; it would have to be a piece of the general type in respect of which we were experienced listeners -- the musical equivalent of bathroom tiling. Only this way would the foundation of interactional expertise on which judgement must be based be distributed among the population of non-musician listeners.

Downward discrimination: peer review and its variants

The claimed superiority of peer review as a method of judging scientific papers, grant applications, and the like, is based on the idea that the best judges of an expertise are those who share the expertise: only those with contributory expertise should judge those with contributory expertise. But, and this may be little more than a truism, the medium of judgement, even when contributory expertise is used to judge contributory expertise, is interactional expertise. Quite simply, the reviewer of, say, a paper or grant application in, say, gravitational wave physics, is not exercising contributory expertise -- he or she is not engaged in the physically involving act of detecting gravitational waves at the time the review is being written -- the reviewer is exercising interactional expertise -- the ability to talk or write about gravitational wave physics.⁵⁷ Luckily, as we have seen, the relationship between contributory and interactional expertise is transitive: to have contributory is to have interactional expertise. If the interactional expertise is latent it will have to be realised to the extent that the reviewer is going to make a useful comment on the paper or application.

Now, the possession of contributory expertise can be taken to guarantee that the maximum possible (latent) interactional expertise has been acquired and that is a very good reason for taking contributory expertise to be a sound basis for judgement. A transitive relationship works one way; the possession of interactional expertise does not guarantee the possession of contributory expertise but, so long as the

⁵⁵ But see Shapin and Schaffer 1987 for the class basis of satisfactory of what was counted as 'witnessing' in early scientific experiments.

⁵⁶ And likewise that one might be able to teach or coach an expertise without being able to practice it.

⁵⁷ Of course, in a wider sense this is a contribution to gravitational wave physics per se, but only in a wider sense.

interactional expertise is realised to the full, someone who possesses it ought to be as good a judge of the contributory expertise to which it pertains as someone who has the contributory expertise itself. In practice, however, the relationship is going to be more complicated. It is very hard in practice (though not in principle), for someone with no contributory expertise to master the same level of interactional expertise as a fluent person with contributory expertise. So, in the main, but not always, those with contributory expertise will be (potentially) better judges than those without. The issue is confounded if the interactional expertise of those with contributory expertise remains largely latent -- that is they lack interactional and reflective ability. In such a case a person with a great deal of interactional and reflective ability and a modicum of interactional expertise may turn out to be the better (though less than optimum) contributor to a decision-making panel.⁵⁸

A still stronger claim is likely to be made by art critics and the like. They will say that a level of connoisseurship (which, as we have established, is itself interactional expertise), can be developed through assiduous viewing and discussing of art. They would claim that this makes for judgements that are superior in principle to those of artists. For example, it might be argued that artists generally work in a narrow genre whereas critics have wide experience. Artists give implicit support to this kind of claim by refusing to exercise their interactional and reflective abilities, preferring to 'let the art speak for itself.' (Later in the book we will try to turn this difference between the sciences and the arts into a general principle of demarcation based on the different loci of meaning of the two cultural endeavours.)

In writing the above passage we have talked not just about the broad boundaries of potential expertise but about what might make one expert better than another. We have more or less said that, other things being equal, in the matter of judging an expertise 'E,' the more (realised) interactional expertise in E the better. This leads us onto dangerous ground but it is ground that cannot be circumvented. It is dangerous because three decades of research in science and technology studies has shown us that internal judgements made by one expert about another are always contestable.⁵⁹

Does this mean we have fallen into an epistemological trap? The answer is that if it is a trap it is a shallow trap. It must be possible to make certain internal judgements about expertise. If it were not, none of the comparisons we have discussed in the section on hoaxing and faking would make sense. In the absence of internal judgement it would make no sense to say that solo violin-playing is more difficult to fake than ability in avante-garde art because the difference between skillful and unskillful performances of all kinds would be

⁵⁸ Here, as we shall argue later in the book, lies an important difference between the sciences and at least some of the arts. In the arts the locus of judgement favours realised interactional expertise and so it favours interactional and reflective ability more highly than in the sciences.

⁵⁹ It is the very contestability of such judgements that makes the attributional theory of expertise, that we will discuss in later in the book, so attractive.

impossible to notice. Life would be one long gamble with chance when it came to judging even a difficult expertise like violin-playing and clearly there is more to life than this. Confidence trickery and the like would cease to be a puzzle in need of explanation because it would come as no surprise that an unskilled person could pass themselves off as a skilled person -- there would be no trick in it -- there would be nothing to be explained. In other words, we could not make sense of the way we live our lives without some notion of internal assessment of expertises. So how do we make internal judgements about expertise that are, on the whole, more likely to be right than wrong.⁶⁰

The principle toward which we are working is what we will call 'downward discrimination.' Even if one's level of expertise is low one can still sometimes make the internal judgement that another is acting as though they have a still lower level of knowledge or understanding. It is such a judgement that we will apply when we claim that the focus group members and the GM Nation participants (both discussed later in the book), did not understand the technicalities of the GM debate. We claim we can make a judgement within a discipline even though our expertise within the discipline is very low; we can do it because those we are judging have, recognisably, a still lower level of expertise.⁶¹

There is an important difference between external and internal judgement in this regard. External judgement does not have a preferred direction: it can be applied equally well upward, downward, or horizontally. That is to say, one can reasonably distrust the demeanour or interests of a Lord Marshall, or a Sellafield spokesperson, even if that person has much more expertise in the technical domain in question, but one cannot reasonably distrust the technical understanding of an expert unless they have substantially less technical knowledge. While the epistemological problem may run all the way through, the difficult practical problem of argument between experts shows itself only when levels of expertise lie in the same broad range.

⁶⁰ As always, we are not trying to provide an infallible formula, merely some general principles. Also, we should note that most of our judgements of professional expertise are not really judgements of expertise at all but merely the normal taking on trust the efficient workings of society and the correspondence between the role that someone occupies and their capabilities; external judgements in other words. What we have said is that failures of internal judgement would not need to be explained by reference to external judgement were there no such thing as internal judgement.

⁶¹ Social scientists also have some referred expertise (see below) that is applicable to the GM case and similar in that they learn something about the nature of sound argument and statistical inference through their professional training.

Internal horizontal discrimination does not exist; where expertises are similar there is argument and negotiation rather than clear judgement.⁶²

Obviously, internal technical discrimination cannot be applied in the upward direction: it is impossible to make a reasonable assessment of the technical understanding of an expert with far more expertise. A corollary is that those downwardly discriminated against may have too little technical understanding to recognise the validity of the judgement however well-founded it is; higher levels of expertise may not be recognised for what they are. Well-founded downward technical discrimination is all too easy to confound with bias.⁶³

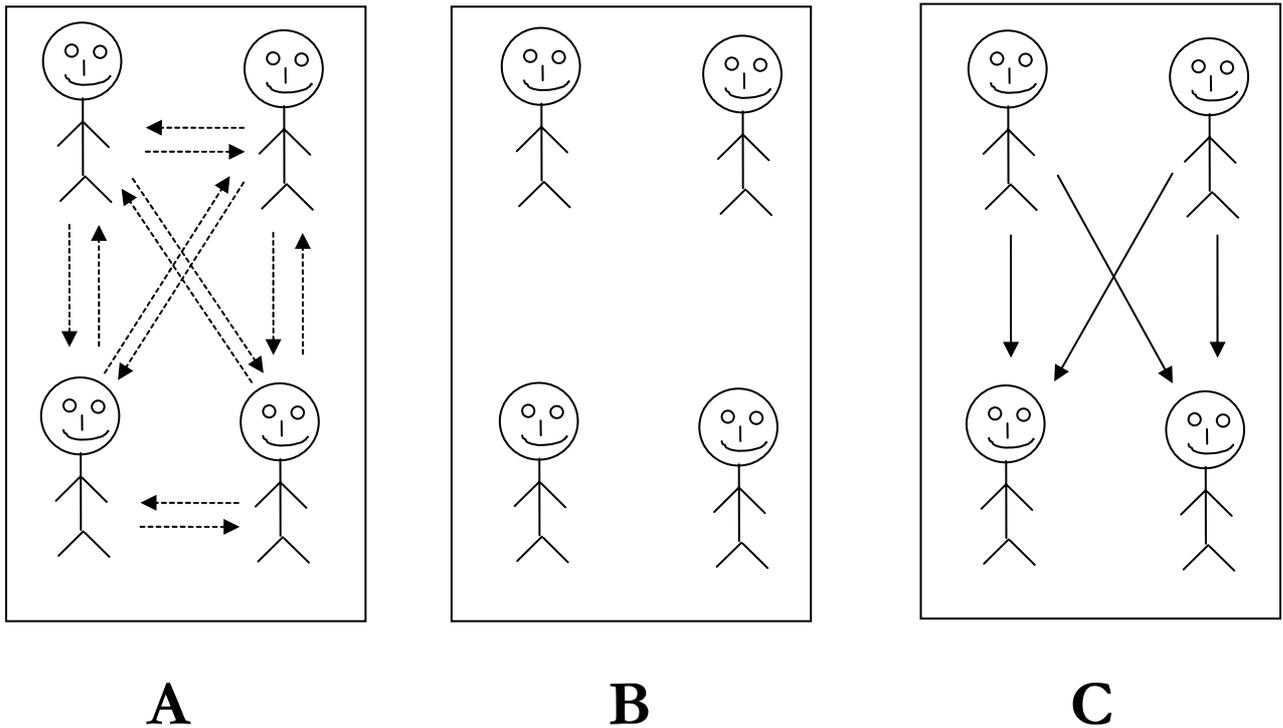
Why doesn't the idea of downward discrimination simply return us to the old fashioned view of top-down scientific authority? Because it works only where there is a settled consensus. Thus, in our critique of the public understandings of GM foods we will exercise downward discrimination only in respect of those aspects of the argument where debate is long settled, not in respect of the technology of GM as a whole. For example, we will criticise the public's invocation of what is 'natural' and we will criticise such things as the claim that GM foods expose the consumer to radioactivity because radioactive isotopes are sometimes used as markers in GM laboratory experiments. But we will not criticise the suggestion that insufficient testing has been done to guarantee that herbicide resistance will not spread to weeds and the like. From our position we are able to see that the argument from naturalness is too indiscriminating to use as a basis for decisions and we can see that to associate GM foods with radioactivity is a technical mistake, but there is insufficient consensus to criticise the 'insufficient testing' argument.

Let us sum up these arguments with another cartoon (see top of next page).

In the figure box A indicates what we tend to believe about our powers of judging others' expertise. We believe we can make internal judgements of expertises upward, downward, and horizontally. Box B is the open space available to attributional theories of expertise. The sociology of attribution is the study of the way

⁶² This is the situation that holds between bathroom tiler and householder or architect. The architect is brought in not because he or she is better at recognising good tiling but because his or her professional status can be invoked to settle what might otherwise be an endless argument about standards. The point is that the interactional expertise of householder or architect will not in itself settle the issue, it being applied horizontally at best, but it does give its possessor a place at the negotiating table.

⁶³ Remember, all judgements may be wrong; we are just describing what is reasonable. Note also that similar judgements are made in social studies of science; the symmetry principle is generally used only in cases of roughly equal levels of expertise and experience (eg, among university scientists). In most of these cases attempts by untrained or inexperienced persons to make their opinions count are usually dismissed out of hand using downward technical discrimination. (See, for example, Collins's 2004b discussion of letters in 'green ink' but cf Malcolm Ashmore's approach.)



actors negotiate the existence of arrows between actors; public legitimacy can be assigned to any arrow in any direction and those which do in fact gain public legitimacy gain it as an outcome of the interplay of power, alliance-building, and so forth. For example, in recent years the folk-wisdom view has given a great deal of legitimacy to arrows going in the upward direction while reducing the potency of downward arrows. Box C shows the normative view that we are developing here for the purposes of policy-making. Internal technical judgements, we say, can be made with a reasonable degree of soundness only when they run downward. (Each box represents only one specialism. According to us there would be no solid arrows between specialisms.)

Referred expertise

Another kind of reasonable internal judgement is based on 'referred expertise.' Referred expertise is expertise taken from one field and applied to another. The term is taken from the idea of 'referred pain' -- for example when a back injury results in pain in the leg. Consider the managers and leaders of large scientific projects. In general they will not possess contributory expertise in respect of the many fields of science they must coordinate. In the field study of one of the authors this became a bone of contention. As one expert scientist put it:

... What I found disappointing was that after two years the project manager still didn't really know what it meant to do interferometric detection of gravitational waves.

Whereas a manager saw it this way:

Once you professionalize, the guys who are very good in the lab where you control everything, no longer have their arms around it all. Other people can work very well in that environment. They interface with the experts who built the electronics and understand what they need to do that; they interface with the computer people and do very well at that; and some people can work in this broader environment technically. Some people make the mistake of saying that as soon as you are in this broader environment it's a management problem; it's not a management problem! The technical part is actually more technical and more sophisticated.

If we stay with the terminology we have developed so far and we did not have referred expertise as a category we would have to say that in respect of the science they are managing, the managers have, at best, interactional expertise. Does this mean that their technical expertise is no greater than that of, say, a sociologist who has developed interactional expertise? To say 'yes' seems wrong and disagrees head-on with the last sentiment expressed in the above quotation. The resolution seems to be that to manage the science of a scientific project requires not contributory expertise in the sciences in question but contributory expertise in some related science. In other words, the managers must know, from their own work and experience in other sciences, what it is to have contributory expertise in a science; this puts them in a position to understand what is involved in making a contribution to the fields of the scientists they are leading at one remove, as it were. It would be quite reasonable to expect that managers of scientific projects with referred expertise would manage much better (as well as with much more authority and legitimacy), than those without it.⁶⁴ We may add that to make the expertise referrable from one field to another managers have to develop at least interactional expertise in the new specialism.

We know that not all managers of scientific projects have referred expertise. General Groves, who ran the Manhattan Project, seems to be a case in point.⁶⁵ The question of whether you need referred expertise to manage a science is, presumably, related to the question of how much specialist knowledge you need to

⁶⁴ Though in the case in question, some scientists thought that the referral was from too distant a site. They thought that high-energy physics, from where the managers came, gave them a misleading picture of the skills required to do interferometry.

⁶⁵ eg see Thorpe and Shapin 2000.

manage anything; it relates, in other words to the judgement gap. If you believe that referred expertise is a good thing for managers then to manage the making of 'X' you need, at the very least, experience in making the closely related 'Y.'

Referred expertise, of course, is not the only kind of expertise needed by the manager of a scientific project. Such a manager also needs expertise in financial management, human resource management, networking skills, political skills, and so forth. Crucially, a manager of a scientific or technological project will need local discrimination; they will need to know how to judge, if not between the competing scientific arguments in the specialism, at least between the scientists in the specialism. The manager will have to listen to the competing claims of different specialists, each of whom will be more accomplished in terms of contributory expertise in the specialism, and judge between them.⁶⁶

META-CRITERIA: CRITERIA FOR JUDGING EXPERTISES

Our goal, as explained at the outset of this chapter, is to separate those who fall into the envelope of potential judges in respect of various expertises from those who fall outside that envelope. Another way to try to do this is by reference to externally measurable criteria.

Credentials

The standard way to try to measure expertise externally is by reference to credentials in the sense of certificates attesting to past achievement of proficiency. Possession of certificates will define a number of kinds of expert but note that not credentials exist to signify possession of many of the expertises we have discussed so far. There are no credentials for fluency in one's native language, nor for moral judgement, nor for political judgement. There are no credentials for ubiquitous discrimination, no credentials for the ability to distinguish between experts and novice violin-players, nor for the majority of other forms of connoisseurship (the exception being some kinds of professional roles that involve connoisseurship such as that of the

⁶⁶ Though bear in mind that, as sociology of scientific knowledge has shown, and as scientists acknowledge, judging the science even within an esoteric specialism often amounts to judging the scientists. For a sociological analysis see Collins 1985/1992; for a scientists remark, see Wolpert 1994 who says: '... scientists must make an assessment of the reliability of experiments. One of the reasons for going to meetings is to meet the scientists in one's field so that one can form an opinion of them and judge their work.'

In the management of large scientific project referred expertise can have advantages over contributory expertise; it carries less commitment to any particular way of doing things and can make for more unbiased decision-making. Collins 2004b discusses the problem in respect of Laser Interferometer Gravitational-Wave Project.

architect). Above all, there are no credentials for unqualified experts, such as the Cumbrian sheep farmers or the AIDS activists. Therefore we conclude that credentials are not a good criterion for setting a boundary around expertise.

Track record

Track record is a much better criterion than credentials. The philosopher Alvin Goldman argues that track record of success in making sound judgements is the best way for lay persons to choose between experts.⁶⁷ Reference to track record of success will certainly exclude a lot pseudo-experts but, again, it excludes too many. For example, it again excludes the sheep farmers and the like who might be applying their expertise to a technical debate in the public domain for the very first time. Likewise it excludes the ubiquitous and local discrimination of the public for which no track records of success are available.

Experience

A criterion that does seem to set the boundary in roughly the right place is experience in a domain. This nicely includes the sheep farmers, the AIDS activists, and all the other categories of expertise that we have described while excluding the general public from technical domains. We know from the outset that without experience within a technical domain, or experience at judging the products of a technical domain, there is no specialist expertise. Without experience of doing science, talking to scientists, playing or listening to violin-playing, or looking at and discussing bathroom tiling, the minimal standards for making judgements in these areas have not been met.

Thus, examination of the experience of, say, an Alan Sokal, would have been a better guide to the value of his work than superficial peer review, and the same applies to Chauncey Gardner, Hancock's avant-garde artist, and any number of confidence tricksters. Confidence tricks and other such scams work when experience is attributed to the fraud on the basis of short acquaintance whereas knowledge of their experience would eliminate them from the class of experts. (Though, of course, it would not provide assurance that they were competent).⁶⁸

⁶⁷ Goldman 2001.

⁶⁸ The story of Ramanujan, the poor Indian village dweller, who is said to have made original discoveries in number theory, argues that experience is too conservative a criterion. His case would suggest that discoveries (and therefore judgements), can be made by the inexperienced individual. But the Ramanujan case, as it is usually presented -- almost a case of a 'private language' -- also sits in opposition to the whole tradition of science studies which stresses embedding into paradigms/forms-of-life. Perhaps the Ramanujan case needs to be reexamined in the light of this work. Did Ramanujan really make his discoveries in such isolation? Perhaps his work was heavily edited. Are there more such cases in modern science?

PERIODIC TABLE OF EXPERTISES

Let us now summarise what we have said about kinds of expertise. The argument is summarised in the 'periodic table' of expertises that is presented at the head of this webpage

Ubiquitous Expertises are acquired by all members of human societies during the course of the normal 'enculturation' that takes place during upbringing. They include fluency in the natural language of the society and moral and political understanding. Ubiquitous expertises are the beginnings from which all other expertises are built.

Dispositions such as interactive ability and reflective ability convert latent interactional expertise into realised interactional expertise.

Specialist, or domain-specific expertises include those with a relatively invisible component of ubiquitous tacit knowledge such as beer-mat knowledge, popular understanding, and primary source knowledge, and the full-blown specialist tacit knowledge-laden expertise which enables those who embody it to contribute to the domain to which it pertains;⁶⁹ this is contributory expertise. The bridge between experts with contributory expertise and non-domain experts is interactional expertise. Interactional expertise is tacit knowledge-laden expertise in the language of a domain and it is acquired through enculturation in the domain language. Interactional expertise is the medium of discussion where technical judgements are made. There is a transitive relationship between the five specialist expertises.⁷⁰

Meta-expertises are used for judging other expertises. External meta-expertises turn on the judging of skills through the judging of persons, or the more general characteristics of their discourse, rather than on domain-specific understanding. They include ubiquitous discrimination and local discrimination (which turns on local knowledge of people). Internal meta-expertise do depend on a degree of technical expertise within the domain. The most straightforward kind of internal meta-expertise depends on the application of contributory expertise to a domain through the mediation of interactional expertise. Downward discrimination applies even where a relatively low level of domain expertise is applied to a still lower level. Technical connoisseurship turns on interactional expertise alone which may have been specially refined for the

⁶⁹ Remember, we include widespread tacit-knowledge laden skills such as car-driving among the specialist expertises.

⁷⁰ If the table was three-dimensional we could add polimorphic actions on the other dimension; the ability to execute polimorphic actions is an expertise that, in essence, does not involve tacit knowledge at all. (We say 'in essence' because it may well be that human beings, because of the way they are made, learn to execute polimorphic actions in a non-explicit way, through practice and so forth. Therefore a human may experience the mastery and execution of polimorphic actions as tacit-knowledge rich though the action could be reproduced by mechanisms incapable of socialisation.)

purpose as in the case of certain kinds of professional or critic.⁷¹ Referred expertise depends on the transfer of domain-specific contributory expertise from one domain to another.

Meta-criteria are attempts to provide externally visible indicators of expertise. We have argued that experience is the best of the three possibilities presented.

To call the table a 'periodic table' of expertise carries the risk of hubris. The table has some structure but not enough to ensure that, like the periodic table of the elements which essentially describes the structure of the atom, it represents some kind of eternal truth. There are certainly other existing ways of classifying expertise, some of which we will look at in an Appendix to the book and there may be better classifications yet to be worked out. Nevertheless, it is vital, for the reasons we set out in other parts of the book, to make a start.

BIBLIOGRAPHY

Arksey, Hilary, (1998) *RSI and the Experts: The Construction of Medical Knowledge*, London: UCL.

ASA 1999,

Barnes, Barry, (1985) *About Science*, Oxford: Basil Blackwell.

Bloor, David, (1973) 'Wittgenstein and Mannheim on the Sociology of Mathematics', *Studies in the History and Philosophy of Science*, 4, 173-91.

Bloor, David, (1983) *Wittgenstein: A Social Theory of Knowledge*, London: Macmillan.

BSA 2002

Cole, Simon A., (2001) *Suspect Identities: A History of Fingerprinting and Criminal Identification*, Cambridge, Mass: Harvard University Press.

Collins, H. M., (1974) 'The TEA Set: Tacit Knowledge and Scientific Networks', *Science Studies*, 4, 165-186.

⁷¹ Of course, the sociology of scientific knowledge shows that, even in the heartlands of laboratory science, much of what passes for internal judgement of knowledge claims is external judgement, usually of the local discrimination variety.

- Collins, H. M., (1985/1992) *Changing Order: Replication and Induction in Scientific Practice*, Chicago: University of Chicago Press [First edition, Beverly Hills & London: Sage, 1985].
- Collins, H. M., (1988) 'Public Experiments and Displays of Virtuosity: The Core-Set Revisited', *Social Studies of Science*, 18, 725-48.
- Collins, H. M., (1990) *Artificial Experts: Social Knowledge and Intelligent Machines*, Cambridge, Mass: MIT press.
- Collins, H. M., (1999) 'Tantalus and the Aliens: Publications, Audiences and the Search for Gravitational Waves', *Social Studies of Science*, 29, 2, 163-197.
- Collins, H. M., (2001) 'Tacit Knowledge, Trust, and the Q of Sapphire', *Social Studies of Science*, 31, 1, 71-85.
- Collins, H. M. (2004a) 'Interactional Expertise as a Third Kind of Knowledge' *Phenomenology and the Cognitive Sciences*, 3 (2) 125-143
- Collins, Harry, (2004b) *Gravity's Shadow: The Search for Gravitational Waves*, Chicago: University of Chicago Press.
- Collins, H. M. (2004c) 'The Trouble with Madeleine' *Phenomenology and the Cognitive Sciences*, 3 (2) 165-70
- Collins, H. M., & Evans, R. J., (2002) 'The Third Wave of Science Studies: Studies of Expertise and Experience', *Social Studies of Sciences*, 32, 2, 235-296 [April].
- Collins, H. M., & Kusch, M., (1998) *The Shape of Actions: What Humans and Machines Can Do*, Cambridge, Mass: MIT Press.
- Dawkins, Richard, (1999) *Unweaving the Rainbow: Science, Delusion and the Appetite for Wonder*, London: Penguin.
- Dreyfus, Hubert L., (1972) *What Computers Can't Do*, New York: Harper and Row.
- Dreyfus, Hubert L., (1992) *What Computers Still Can't Do*, Cambridge, Mass: MIT Press.
- Dreyfus, Hubert L., & Dreyfus, Stuart E., (1986) *Mind Over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*, New York: Free Press.
- Epstein, Steven, (1995) 'The Construction of Lay Expertise: AIDS Activism and the Forging of Credibility in the Reform of Clinical Trials', *Science Technology & Human Values*, 20, 408-37.
- Epstein, Steven, (1996) *Impure Science: AIDS, Activism and the Politics of Knowledge*, Berkeley, Los Angeles & London: University of California Press.
- Fleck, Ludwik, (1979) *Genesis and Development of a Scientific Fact*, Chicago: University of Chicago Press [first published in German in 1935].
- Ginzburg Carlo, (1989) 'Morelli, Freud and Sherlock Holmes: Clues and Scientific Method', *History Workshop Journal*, 9, 5-36.
- Goldman, Alvin, I., (2001) 'Experts: Which Ones Should You Trust?' *Philosophy and Phenomenological Research*, 63, 1, 85-110.
- Goodman, Nelson, (1969) *Languages of Art*, London: Oxford University Press
- Gross, Paul, & Levitt, Norman, (1994) *Higher Superstition: The Academic Left and its Quarrels with Science*, Baltimore & London: John Hopkins University Press.

- Gross, P., Levitt, N., & Lewis, M. W., (1996) *The Flight From Science and Reason*, New York: New York Academy of Sciences.
- Irwin, Alan, (1995) *Citizen Science: A Study of People, Expertise and Sustainable Development*, London & New York: Routledge.
- Irwin, Alan, (2001) 'Constructing the Scientific Citizen: Science and Democracy in the Biosciences', *Public Understanding of Science*, 10, 1-18.
- Irwin, Alan, & Wynne, Brian, (1996) (eds.) *Misunderstanding Science? The Public Reconstruction of Science and Technology*, Cambridge & New York: Cambridge University Press.
- Koertge, Noretta, (2000) (ed.) *A House Built on Sand: Exposing Postmodernist Myths About Science*, Oxford: Oxford University Press.
- Kosinski, Jerzy, (1971) *Being There*, Orlando: Harcourt Brace Jovanovich, Inc.
- Latour, B., and Woolgar, S. (1979) *Laboratory Life: The Social Construction of Scientific Facts*, London and Beverly Hills: Sage.
- MacKenzie, Donald, (1998) 'The Certainty Trough', in Williams, R., Faulkner, W., & Fleck, J. (eds.), *Exploring Expertise: Issues and Perspectives*, Basingstoke: MacMillan, 325-329.
- Pamplin, B. R., & Collins, H. M., (1975) 'Spoon Bending: An Experimental Approach', *Nature*, 257, 8 [4 September].
- Selinger, Evans, & Mix, John (2004) 'On Interactional Expertise', *Phenomenology and the Cognitive Sciences*, 3 (2) 145-163.
- Shapin, Stephen, & Schaffer, Simon, (1987) *Leviathan and the Air Pump: Hobbes, Boyle and the Experimental Life*, Princeton: Princeton University Press.
- Sokal, Alan, (1996) 'Transgressing the Boundaries: Towards a Transformative Hermeneutics of Quantum Gravity', *Social Text*, 46, 47, 217-252 [spring/summer].
- Thorpe, Charles, & Shapin, Steven, (2000) 'Who Was J. Robert Oppenheimer? Charisma and Complex Organization', *Social Studies of Science*, 30, 4, 545-590 [August].
- Turner, Stephen, (2001) 'What is the Problem With Experts?', *Social Studies of Science*, 31, 1, 123-149.
- Weizenbaum, J., (1976) *Computer Power and Human Reason: From Judgement to Calculation*, San Francisco: W.H.Freeman.
- Welsh, Ian, (2000) *Mobilising Modernity: The Nuclear Moment*, London: Routledge.
- Winch, Peter G. 1958. *The Idea of a Social Science*. London: Routledge and Kegan Paul.
- Wittgenstein, Ludwig, (1953) *Philosophical Investigations*, Oxford: Blackwell.
- Wolpert, Lewis, (1992) *The Unnatural Nature of Science*, London: Faber and Faber.
- Wolpert, Lewis, (1994) 'Review of "The Golem: What Everyone Should Know About Science"', *Public Understanding of Science*, 3, 323-337.

- Wynne, Brian, (1989) 'Sheep farming after Chernobyl: A case study in communicating scientific information', *Environmental Magazine*, 31, 2, 33-39.
- Wynne, Brian, (1992) 'Public Understanding of Science Research: New Horizon or Hall of Mirrors?', *Public Understanding of Science*, 1, 1, 37-43 [January].
- Wynne, Brian, (1993) 'Public Uptake of Science: A Case for Institutional Reflexivity', *Public Understanding of Science*, 2, 4, 321-337 [October].
- Wynne, Brian, (1996a) 'May the Sheep Safely Graze? A Reflexive View of the Expert-Lay Knowledge Divide', in Lash, S., Szerszynski, B., & Wynne, B. (eds.), *Risk, Environment & Modernity: Towards a New Ecology*, London: Sage, 27-83.
- Wynne, Brian, (1996b) 'Misunderstood misunderstandings: Social Identities and Public Uptake of Science', in Irwin, Alan, & Wynne, Brian (eds.), *Misunderstanding Science? The Public Reconstruction of Science and Technology*, Cambridge: Cambridge University Press, 19-46.