Rethinking Expertise

HARRY COLLINS AND ROBERT EVANS

The University of Chicago Press
Chicago and London

CHAPTER ONE

The Periodic Table of Expertises 1: Ubiquitous and Specialist Expertises

The Periodic Table Introduced

Table 1 is the Periodic Table of Expertises—a table of the expertises that might be used when individuals make judgments. The table is drawn in two dimensions, but every now and again, as indicated below, a third dimension would be useful. This chapter deals mainly with the first three rows. The next chapter deals mainly with the bottom two rows—Meta-Expertises and Meta-Criteria. We begin, however, with a summary explanation of the whole table that can act as a map or "ready reference"—a quick reminder of the whole structure and the meaning of any particular category. At the end of chapter 2 we will provide another summary for those who have read through the details.

Working from the top of the table, *ubiquitous expertises* are those, such as natural language-speaking, which every member of a society must possess in order to live in it; when one has a ubiquitous expertise one has, by definition, a huge body of tacit knowledge—things you just know how to do without being able to explain the rules for how you do them. This row of the table also includes all those expertises one needs to make political judgments. Below this line the table is exclusively concerned with technical expertises—those that have science and technology content.

Dispositions are not very important to the conceptual structure of the table; they are personal qualities—the ones we discuss are linguistic fluency and analytic flair.

The next row deals with the *specialist expertises.* Low levels of specialist expertise are better described as levels of knowledge—like knowledge of the kind of facts needed to succeed in general knowledge quizzes. One may be able to recite a lot of such fact-like things without being able to do

Table 1: The periodic table of expertises

UBIQUITOUS EXPERTISES					
DISPOSITIONS				Interactive ability	Reflective ability
SPECIALIST EXPERTISES	*UBIQUITOUS TACIT KNOWLEDGE*			*SPECIALIST TACIT KNOWLEDGE*	
	Beer-mat knowledge	Popular understanding	Primary source knowledge	Interactional expertise	Contributory expertise
				Polimorphic	*Mimeomorphic*
META-EXPERTISES	*EXTERNAL (Transmuted expertises)*		*INTERNAL (Non-transmuted expertises)*		
	Ubiquitous discrimination	Local discrimination	Technical connoisseurship	Downward discrimination	Referred expertise
META-CRITERIA	Credentials		Experience	Track record	

anything much as a result except succeed in quizzes. Three low levels of expertise are listed on the left-hand side of the specialist expertise rows of the table. It is important to note that acquiring low levels of expertise seems like a trivial accomplishment only to those who already possess ubiquitous expertise; acquisition of even these low levels rests on the prior acquisition of a vast, but generally unnoticed, foundation of ubiquitous expertise.

To acquire higher levels of specialist expertise, more than ubiquitous expertise is needed. To go further along row three it is necessary to immerse oneself in a domain so as to acquire *specialist tacit knowledge*, not just learn more facts or fact-like relationships. Two categories of higher level expertise are found at the right hand end of the specialist expertise rows. The highest level is *contributory expertise*, which is what you need to do an activity with competence. Just below this, however, is *interactional expertise*, which is the ability to master the language of a specialist domain in the absence of practical competence. The idea of interactional expertise is immanent in many roles, from peer reviewer to high-level journalist, not to mention sociologist or anthropologist, but it seems not to have been discussed before in an explicit way. A good proportion of the book is taken up with explaining the notion of interactional expertise because it is a new concept.

Moving down to the fourth row we encounter *meta-expertises*. The first set of two meta-expertises are the prerogative of judges who, not possessing the expertise in question, make judgments about experts who do possess it. This is done by judging the experts' demeanor, the internal consistency of their remarks, the appropriateness of their social locations, and so forth. These are "transmuted expertises" because they use social discrimination to produce technical discrimination. The first kind of discrimination depends on the kind of ubiquitous expertise one gains in a democratic society as one learns to choose between politicians, salespersons, service providers, and so forth. The second kind of discrimination depends on local knowledge about those around you. The second set of three meta-expertises do not depend on transmutation, as they are based on possessing one level or another of the expertise being judged. *Technical connoisseurship* is like the expertise of art critics or wine buffs who, crucially, are not themselves artists or wine-makers. The middle of the three categories relates to what we most naturally think of as skillful judgment—where one specialist judges another. There are three directions in which this middle category of judgment can be made: an expert can judge someone who is still more expert, an expert can judge someone equally expert, or an expert can judge someone less expert. Mostly experts think they are pretty good at judging in any of the three directions, but we argue that only the downward direction is reliable, the other directions tending to lead to wrong impressions of reliability or irresolvable disputes. The one reliable category which appears in the table is, therefore, labeled *downward discrimination*. *Referred expertise* is the use of an expertise learned in one domain within another domain. In the chapter we use examples drawn from the management of large scientific projects, where a manager moves from one to another, to illustrate the concept.

The final row of the table refers to the criteria that outsiders try to use to judge between experts to avoid having to make the more difficult kind of judgments described above. They can check expert's qualifications, they can check expert's track records of success, or, what we argue is the best method of the criterion-based judgment, they can assess the expert's experience.

Ubiquitous Expertises

We now move on to a more detailed analysis of each category of expertise starting with the top row. 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 realize how hard a task it would be.[1] 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 trivial accomplishments.[2] 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 a ubiquitous expertise. Others include moral sensibility and political discrimination. These are abilities that people acquire as they learn to navigate their way through life. In the case of ubiquitous expertises the Problem of Extension ceases to have any practical significance because almost everyone is a genuine contributory expert.[3] Thus, when we say that the folk wisdom view is often misplaced 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. What

1. We have already mentioned how this mistake was made by the pioneers of natural-language-handling computers.

2. For a discussion of the fact that the so-called unskilled persons who work in McDonald's restaurants actually have a huge depth of ubiquitous expertise, see Collins and Kusch 1998, chap. 8. Here is a clear case where the lack of scarcity value of an expertise has been taken to mean that it is easy.

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. 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 nonprofessional person; someone who is not an expert. But all laypersons 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.

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 patronizing—treating the public as incapable of learning the rules of ordinary living through the normal processes of socialization.

Technical expertises can be ubiquitous in certain societies and not in others—for example, the ability to wire a plug or mend a fuse. In the Society of Social Studies of Science meeting in Pasadena in 2005, Wiebe Bijker suggested that the Dutch population as a whole had sufficient expertise in dam-building to contribute to public debates about new projects in contrast to what everyone agreed was a low level of expertise in the working of levees among the general population of Louisiana. Whether this claim is true or not remains to be proved, but it is an intriguing suggestion that is central to the concerns of this book.

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 central 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. We can 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.

As explained, our model is human-centered. At some stage all human expertise touches on tacit knowledge, that is, an understanding of rules that cannot be expressed. It is the inexpressibility of the rules of ubiquitous expertises that make them so hard to capture in computers. The idea of tacit knowledge will be discussed further in chapter 3, but for now we need to note only that tacit knowledge enters into knowledge acquisition in two ways. Some kinds of knowledge acquisition amount 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 or listening without interactive discourse. 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 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 internalization of the unspoken rules of road-craft. Likewise, becoming a

full-blown specialist in a scientific or technical domain requires immersion in the society of the domain specialists. This gives us an initial division of types of expertise: those that can be acquired using only ubiquitous tacit knowledge and those that involve specialist tacit knowledge.[4]

Expertises Involving Only Ubiquitous Tacit Knowledge

Beer-Mat Knowledge

So far we have established than an essential component of any ladder of human expertise is *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.
>
> 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 analyzing 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 beer mat or something similar, is "the bishop may move, only diagonally, any distance, backwards or forwards." But it is possible to know this rule in more than one way. One might know

4. 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.

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 seven 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 chessboard, that the bishop is restricted to squares of only one color (nor even that there are two colors on the chessboard). 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 chessboards and the game of chess is a rather different thing even for a novice chess-player than for a Trivial Pursuit player. The novice player who knows the rule knows how to move the bishop on a chessboard. (However, it is only with further experience that the novice learns how to recognize when, and in what circumstances, moving the bishop might be a good idea. And it takes the skill of a chess master to see that making *this* move with the bishop is the turning point of the game.)

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 correct anyone's mistakes about the nature of holograms, or to make a sensible decision about the long-term dangers associated with the unrestrained spread of holograms, or convey any information about holograms other than the formula itself.

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."[5] 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

5. We are grateful to Matthew Harvey for helping develop this popular understanding category.

which bodies such as the Royal Society's Committee for the Public Understanding of Science (COPUS) once directed its efforts.

Popular understanding does involve a deeper understanding of the *meaning* of the information than does 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 because both give rise to the same judgments. 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 have shown us that, in disputed science, detail, tacit knowledge, and unspoken understandings of who is to be trusted among those who work in the esoteric core of the science are vital components of decision-making at the technical level. Popular understanding hides detail, has no access to the tacit, and washes over scientists' doubts. The consequence, well-established in the sociology of scientific knowledge, is summed up in the phrase mentioned in the introduction: "distance lends enchantment." Special cases aside, the more distant one is from the locus of the 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.[6] Any rede-

6. For "distance lends enchantment," see Collins 1992. For a modification, see MacKenzie 1998. Ludwik Fleck, who was a sociologist of scientific knowledge before the term was invented, wrote in the 1930s:

> Characteristic of the popular presentation is the omission both of detail and especially of controversial opinions; this produces an artificial simplification [and] . . . the apodictic

scription 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 judgments, or at least *informed* judgments, 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 is likely to yield poor technical judgments.

The problem of judgments based on popular understanding applies whether the conclusion is positive or negative—whether the consumers of the simplified version accepts everything they read and hear (for example, they might accept that Stephen Hawking's utterances about black holes are revealed truth), or rejects the claims (for example, 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 understanding.[7]

One of the troubles with the old officially sponsored 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.[8] 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, popular understanding is a big step up from beer-mat knowledge but a long way from deep understanding of scientific matters. The

---

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 1979 [1935], 112–13, emphasis in original)

7. 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 way certainty increases as the bandwidth narrows (for example, details of the time, place, and personal involved in an experiment are successively removed as accounts enter more public domains). What they do not explain, however, is how the certainty comes to be that the account is unproblematically true, on the one hand, or unproblematically false, on the other.

8. A 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.

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 crucial.

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 be that the feeling of confidence that comes with a mastery of the primary literature is a factor feeding into the folk wisdom view.[9]

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. Many 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 from understanding a scientific dispute.[10] 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

9. 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. For a discussion of expertise in the medical context, see Collins and Pinch 2005.

10. 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, 2004a).

bishop's move. Our claim is that in the case of scientific disputes primary source knowledge is not much better in respect of the science than a chess novice's understanding in respect of the bishop's move.

Expertises That Involve Specialist Tacit Knowledge

Over the last half-century, the most important transformation in the way expertise has been understood is a move away from seeing knowledge and ability as quasi logical or mathematical and toward a more wisdom-based or competence-based model. As has been intimated, expertise is now seen 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. Polanyi, who invented the term "tacit knowledge," has also been influential, especially among scientists and philosophers of science, 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.[11] The Wittgensteinian frame of mind (as interpreted here) leads us to expect to find specialist knowledge located in specialists' practices rather than in books.[12] 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 culture rather than by extended study of dictionaries and grammars or their equivalents. 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

11. 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 that of Bloor 1973 and 1983.

12. We must not pass this point without noting that the logic of what is currently the most dominant trend in science studies, the work of Bruno Latour and Michel Callon around so-called "actor network theory," includes absolutely no role for expertise or any other special property that pertains to human societies or the particular capabilities of humans. On the contrary, actor network theory takes even inanimate objects to be ontologically indistinguishable from humans. Thus, while we talk of a change in our understanding of knowledge, this does not apply to the dominant part of science studies as it is practiced today. For a critique of actor network theory, see Collins and Yearley 1992.

the subject matter of the domains; the acquisition of the first three kinds of knowledge (though it depends on ubiquitous expertises), involves reading rather than immersion in the specialist culture. "Enculturation" is the only way to master an expertise which is deeply laden with tacit knowledge because it is only through common practice with others that the rules that cannot be written down can come to be understood.

What is new about our analysis of expertises learned through immersion in a culture is that we split them 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. This is the traditional way of thinking about this kind of expertise, and we discuss it first before moving on to the new category of *interactional expertise,* an idea which we consider to be a significant contribution to the understanding of expertises in general.

Contributory Expertise

The five-stage model of acquisition of contributory expertise is one of the more well-known and influential schemas, and we will use it to stand for all the important approaches—those that stress the importance of the "internalization" of physical skills.[13] It could be usefully represented on a third dimension of the table. According to the five-stage model, only at the early stages of skill acquisition is there need for calculation or even self-conscious rule-following (the left-hand side of the specialist expertise row in our table); self-conscious application falls away as a skill becomes "embodied"; this is essential for efficient performance. The five stages can be exemplified by the process of learning to drive a car:[14]

- Stage 1 is the *novice.* The novice driver will try to follow explicit rules and as a result the performance will be labored, jerky, and unresponsive to changes in context. The skill will be exercised "mechanically," following rules such as "change gear when the car reaches 20 mph as indicated on the speedometer." These are "context-free" rules, because in applying

13. Dreyfus and Dreyfus 1986. There are, of course, other kinds of approach to expertise and apprenticeship, many of which emerge from the field of education research. See for example, Ainley and Rainbird 1999; Coy 1989; Pye 1968; Lave 1988; and Lave and Wenger 1991.

14. Dreyfus and Dreyfus 1986, 21–36.

them the learner does not take into account the nuances implied by different conditions of application.

- Stage 2 is the *advanced beginner.* As more and more of the skill is mastered, however, more unexplicated features of the situation start to play their part in the performance, such as using the sound of the engine as an indicator of when to change gear, which will in turn mean gear changes at different speeds according to whether, say, the car is going uphill or down.
- Stage 3 is *competence.* Here the number of "recognizable context-free and situational elements" becomes overwhelming, and expertise becomes much more intuitive rather than calculating. "Problem solving" is no longer the predominant motif.
- Stage 4 is *proficiency.* The proficient driver recognizes whole problem situations "holistically" in the same way as the advanced beginner recognizes specific features of the environment. In the advanced beginner stage it is, say, the sound of the engine that is recognized from experience; in the proficient stage it would be a complete traffic scenario. Nevertheless, some elements of conscious choice and analysis remain to guide the proficient driver's decisions.
- Stage 5 is *expertise.* When expert status is achieved, complete contexts are unselfconsciously recognized and performance is related to them in a fluid way using cues that it is impossible to articulate and that if articulated would usually not correspond, or might even contradict, the rules explained to novices. Hence the common experience of driving a familiar journey to work and being unable to remember anything about it when one arrives; something else was occupying the mind during the journey and the unselfconscious self was left to cope with controlling the car just as it normally copes with walking or chewing (again, activities about which we know nothing in a self-conscious way). It is also the case that when experts attempt to revert to a more self-conscious way of tackling a task they do it less well—like the mythical centipede that would trip over its legs if it thought about where it was putting them. In sum, skills practiced by individuals have to be "internalized" if they are to be practiced efficiently.[15]

15. For pedestrians, one may think of the way we learn to cross the road. Explicit rules—"look left, look right, look left again and if nothing is coming walk rapidly across"—disappear as they become absorbed into the generalized skill of crossing roads; this is a skill that suddenly reappears, and has to be relearned starting with a conscious routine, when we go to a country where they drive on the other side of the road. Knowing how to cross the road is known by the novice as a set of explicit and fixed rules, but by the experienced road-crosser as an unexplicated skill which is acted out in different ways as each new and

The five-stage model would be represented in a three-dimensional table by columns coming out of the page wherever a practical skill was under examination, but it has no bearing on many of the categories. A problem with the five-stage model, even as a discussion of contributory expertise, is its individualistic nature. Bicycle riding has a venerable history as an example in the debate about the nature of skill, and we will switch for a moment from cars to bikes. Michael Polanyi introduced the bicycle example, pointing out that the physics of riding a bike is exceptionally complex and counter-commonsensical and certainly of no use to those wishing to learn to ride. But imagine that our brains and nerve impulses were speeded up a millionfold: Would things change? We can ask the same question in reverse, as it were, by slowing everything down. Suppose the loss of balance happened much less quickly (as in bicycle-riding on the Moon or on an asteroid with a still lower gravitational field). The bike might fall over so slowly that there would be time to read a book of balancing instructions and follow them in the new, much slower, real time. Bike-riding would then become more like assembling flat-pack furniture: you hold the instructions in one hand and obey them without any significant time constraints.[16] The physics of bike-riding is not, then, as forbidding as it seems. Though humans cannot master it, there seems no a priori reason why a much faster non-human machine could not master it.[17]

Crucially, what Polanyi was discussing was not "bicycle-riding" but "bicycle-balancing." *Bicycle-riding* has two components: the first is balancing upright; the second is negotiating traffic. Car-driving has the equivalent two components: the first is control of the gears, steering, etc. and the second is, once more, negotiating traffic. Negotiating traffic is a problem that is *different in kind* to balancing a bike or using the clutch in that it includes understanding social conventions of traffic management.

---

unanticipated circumstance is encountered. As ability to cross the road increases, the pedestrian seems to *know* less and less about it. Actually, the experienced road-crosser uses a non-machine-like set of procedures having to do with making eye contact with the driver, and so forth. Learning to cross the road fits the Dreyfus model quite well.

16. See Collins 2007 for a discussion of bicycling on the Moon. Or think about the human-specificity of the skill of sportsmen and women. Batting at baseball or cricket would be an entirely different proposition if the batter's brain worked 1,000 times as fast. It would be a matter of hitting what would be, essentially, a stationary ball—the skills required would be more like those of stationary ball games, such as pool or golf—something far more calculative. (Collins and Kusch 1998 point out that golf-ball-striking can be done better by machine.)

17. As a matter of fact, bike-riding has been accomplished by mechanical means, but this seems to be a matter of analogue feedback from gyroscopic sensors. There seems no reason why the analogue device could not be reproduced by a sufficiently complex digital version.

These are the property of social groups; they vary from place-to-place and time-to-time. To master them requires not embodiment of the skill but being socialized into the relevant group practices.

The difference between bicycle—balancing and negotiating traffic has been described in terms of the difference between mimeomorphic and polimorphic actions.[18] Mimeomorphic actions, however complex, and however hard to master, do not turn on social understanding and can, in principle, be reproduced by mimicking fixed behaviors—though sometimes these will be too complex in practice to be accomplished. It is for this latter reason that automation of factories and so forth has to start with standardization of the whole manufacturing process, not just the replacement of individual machines within the chain of production. Furthermore, humans (working in a normal time-frame), cannot master complex mimeomorphic actions in a machine-like way for the reasons explained by Polanyi. In most cases humans have to internalize the abilities and the process of learning the new abilities in ways that appear similar to the learning of a social skill. Closer analysis, as in the case of bicycle-balancing, shows that this is a matter of the limitations of humans rather than the intrinsic nature of the expertise. It follows that sometimes machines that do not have human limitations can master the skills; we can easily imagine a very fast computer being constructed that would use the explicit physics of bike-riding along with an array of feedback devices to balance a moving bike. On the other hand, polimorphic actions, which do depend on social understanding, require that behavior fits changing social circumstances, and they cannot be mastered by machines failing a way of making machines that fit as smoothly into social life as humans.

We indicate the difference with the two boxes beneath the contributory expertise box in the Periodic Table but one might think of this distinction too as something that should be represented by a third dimension coming out of the page since a similar analysis could be conducted for every box. We will not say much more about mimeomorphic and polimorphic actions here as they are the subject of an entire book—that referred to in the last footnote. The distinction is crucial for understand the relationship between humans and machines but also for the proper understanding the relationship of the human body and brain in the acquisition of expertises.[19]

18. Collins and Kusch 1998.

19. The famous example of the breadmaking machine discussed by Nonaka and Takeuchi (1995) would also have benefited from breaking down the actions of human breadmaking into its mimeomorphic and polimorphic components. If bread, like music, is sometimes

Interactional Expertise

The overlooked second type of deeply tacit-knowledge-laden expertise is interactional expertise. This is expertise in the *language* of a specialism in the absence of expertise in its *practice*. This may seem contrary 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, 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, on the other. The formal can be expressed in rules, formulae, and facts, and can be encapsulated in computer programs, books, and the like. The informal or tacit, insofar 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. They can be recognized as rule-like because it is easy to see when they have been broken. That is, it is easy for those who have internalized the rules, by being enculturated into the form of life that expresses them, to see when they have been broken.[20] The perennial question that emerged with particular clarity in the debate about "artificial intelligence" is whether the informal can be reproduced by sets of formal rules if the set of rules is made large enough. This question has tended to polarize analysts.

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

- *Informal view:* Full immersion in an entire form of life would be needed to master a language.

made to fit a context, then making bread using a machine is like listening to a recording of a concert rather than listening to a live performance. The former is always exactly the same, the latter varies subtly each time. In fact, even when we move to the apparently mimeomorphic aspects of breadmaking we are likely to find that the inputs and outputs will have to be more standardized than would be necessary in the case of a human breadmaker (Ribeiro and Collins 2007). For an extended discussion of the relationship between Dreyfus's analysis of expertise and that discussed here see Selinger and Collins 2007.

20. For example, I may not be able to say what the rules for proximity to others are in various societies, but with a little habituation I will be able to accomplish them and I will also soon be made to know if I break them, and can easily recognize if someone in my own society breaks them (say, by standing too close to me).

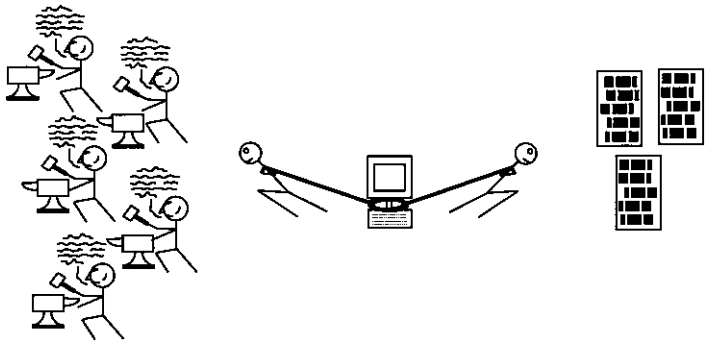

Figure 1. Is the "intelligent computer" text or practice?

- *Formal view:* 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.

The second of these possibilities—that of the "formalists"—has, we believe, been thoroughly exploded. It has been shown to be wrong by theoretical analysis, and it has been shown to be wrong by the failure of the long-running experiment with intelligent computers. This experiment shows that the language of a domain, like any spoken language, consists of more than propositional knowledge. The problem this has left 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 most informalists, to amount to the claim that the formal view is true since the only possibilities on the table, as it were, are the informal and formal views. This excludes analysis of 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 the two views.

Figure 1 illustrates the point in cartoon form. On the left we have a group engaged in the practical activity and the 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 what is, essentially, a set of propositions or

symbols whose natural home is with the printed artifacts and the like on the right. The existence of this 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 between the skilled group and the books. For example, one of the standard motifs in critiques of artificial intelligence is the mistakes made by coaches of any practical activity, the point being that the coach, whether mechanical or human, cannot express in words what the athlete or other kind of learner can only master by action.[21] The idea of interactional expertise gives us a better idea of what some human coaches might be doing and how they succeed despite the gap between language and practice. As we now see, the human coach *can* teach some things through the medium of spoken language because the coach shares some of the nonexplicit skills of the student: the shared linguistic skills can transfer mutually understood tacit meanings that would not be available to those with levels of expertise below interactional.[22]

Interactional expertise, then, is found in this middle ground between practical activity and books, computers, and so forth. Interactional expertise is, however, nearer to the informal than to the formal view. Interactional expertise is far from 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 cannot be expressed in propositional terms. The computer, no current or 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 the informal view—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-

21. For typical critiques of coaching see Dreyfus 1972, Dreyfus and Dreyfus 1986, and Collins 1990.

22. Coaches can also transfer tangential rules, such as "hum the Blue Danube when you swing your golf club," and "second order measures of skill." Second order measures of skill are such as: "If you are a surgeon, intending to spay a ferret, and you cannot find its uterus first time, look again but don't look more than about six times." These specific sets of rules do not comprise the skill but do give valuable guidance to a human trying to master a skill. In the same way it is good to know that (1) a human can ride a bicycle; (2) it is likely to take several hours to learn; (3) a human can learn to play the piano; (4) it is likely to take at least a year to learn; (5) a human cannot learn to fly unaided. Collins 1990, chap. 6, discusses tangential rules. Pinch, Collins, and Carbone 1996 deals with second order measures.

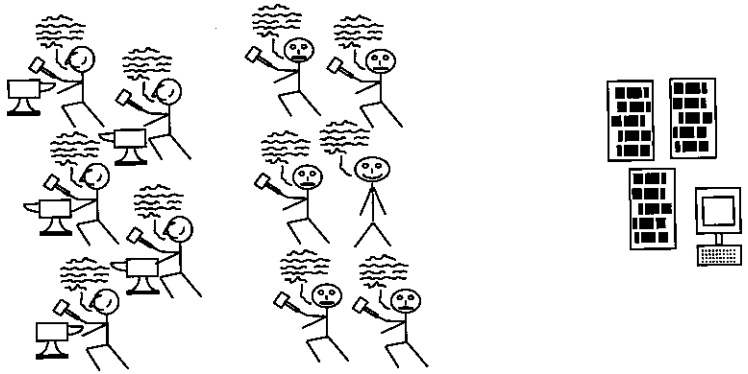

Figure 2. The sociologist delighted with interactional expertise

blown physical immersion in the form of life. To try to be as sharp as possible, we make a "bold conjecture": the testable "strong interactional hypothesis." The strong interactional hypothesis states that, in principle, the level of fluency in the language of a domain that can be attained by someone who is only an interactional expert is indistinguishable from that which can be attained by a full-blown contributory expert. Figure 2 illustrates the point.

Figure 2 shows the computer in its proper place, while interactional expertise occupies the left center 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 talk about their form of life rather than in 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 bold conjecture of the strong interactional hypothesis.

Origins of Interactional Expertise

The idea of interactional expertise has wide application. As well as being needed in some approximate form by successful participatory sociologists, ethnographers, and social anthropologists, mastery of interactional expertise is also the goal of specialist journalists; it is needed by

salespersons and, as we will argue, by managers; it is often the medium of specialist peer review in funding agencies and journal editing where the reviewers are only sometimes contributors to the narrow specialism being evaluated; it is the medium of interchange within large-scale science projects, where again not everyone can be a contributor to everyone else's narrow specialism; it is, *a fortiori*, the medium of interchange in properly *interdisciplinary*, as opposed to multidisciplinary, research; it is, arguably, the basis of the assumption of false identities in Internet chatrooms;[23] finally, on those occasions when activists or other concerned persons are driven to it, it can be the medium of interchange between scientists and groups of the public.

Though it has very wide application, the idea of interactional expertise emerged for us as a result of our experience of sociological fieldwork, and this is how we will introduce it. 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 attained.[24]

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

23. But confidence tricks of all kinds generally depend on a major contribution from the "mark"—the person being fooled (Maurer 1940). In the imitation game (see below) it is quite clear that the judge's job is to try to distinguish the true expert from the expert who has mastered only the language; the judge does not make the kind of contribution made by the mark.

24. In rare cases the sociologist might even progress to the level of contributory expertise. Contributory expertise can be attained only by practicing 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 prerequisite to certification. It is not impossible, however, in sciences which are not too difficult, as in Collins's contributions to parapsychological research, where he took an active part in designing and carrying out experiments (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.)

technical skill). Interactional expertise cannot always be attained—the science may simply be beyond the capacity of the analyst, as one of the authors, Collins, discovered when he attempted to do research in the field of the theory of amorphous semiconductors. After completing about thirteen hours of taped interviews with scientists, he had to concede 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 that 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 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" that marks the switch to mastery of interactional expertise, but its steady acquisition can nevertheless be recognized. 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 the practice of their science 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 respect of some common problem. Sometimes the analyst will be able to introduce a new piece of science to a scientist.[25] 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

25. For example, in May 2005 Collins found himself explaining the "Christodoulou effect," a technical wrinkle in gravitational wave physics (Collins 2004a), to a gravitational wave physicist during the course of a workshop—the physicist had never heard of it.

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-pulls are all recognized, and respondents will no longer be tempted to give a "pat" or formulaic answer drawn from a ready-made set of responses representing the canonical face of science. Mostly, respondents 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 recognize 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 semiconductors, 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 the theory of amorphous semiconductors was at the time of this fieldwork), the analyst cannot understand what the protagonists are disagreeing about, nor how deep the disagreements run, nor, with any certainty, who disagrees with whom! 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 of 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; that would demand 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.

The Parasitic Nature of Interactional Expertise

We can be fairly sure that a 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 socialization; 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, in practical fields, 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 communities who have contributory expertise in that specialism, not persons who have interactional expertise in that specialism. One would guess that, if the attempt were made to transmit interactional expertise in the absence of contributory expertise over several generations, it would rapidly become distorted as messages are distorted when they are passed on by word of mouth through many intermediaries. The point is that interactional expertise is skill in speaking a specialist language, and the nature of a whole language is a function of the whole environment, physical and social, in which it develops. Change the environment (e.g., 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, as we will argue at greater length in chapter 3.

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, at least in principle, all of the lower levels but not *vice-versa*. There are, however, a few practical exceptions to the transitivity. First, as we will discuss in the next section, a contributory expert's interactional expertise may be "latent," i.e., not realized. Second, contributory experts may know the journal literature only at second hand rather than have the firsthand acquaintance of those whose knowledge extends only as far as the primary sources. Experts' knowledge does not come primarily from an exhaustive knowledge of the literature but from a familiarity with a subset of the literature, often at second hand, and always modulated by the opinions of other experts. Third, and for similar reasons, it may well be that specialists in general knowledge quizzes and the like could have a greater breadth 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 expertise, and contributory expertise, 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

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 Repetitive Strain Injury patients (Arksey 1998); AIDS treatment activists (Epstein 1995, 1996) and nuclear protestors (Welsh 2000).

expertise are fewer in number still, since gaining interactional expertise requires crossing social boundaries and spending a long time in alien social environments to which there is restricted access. Finally, those with contributory expertise may, in highly technical sciences, be limited to somewhere between a half-dozen and a few hundred. (Remember that all contributory experts are counted as possessing interactional expertise by definition—the numbers of interactional-but-not contributory experts are very small.)

Interactional Expertise and Interactive and Reflective Ability

Interactional expertise looks similar to but is distinct from 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 smoothly about the domain which they have chosen to study or within which they exercise their judgment, to reflect upon their subject matter so as to articulate their findings or judgments, and sometimes to translate the expertise of one domain into the language of another insofar as this can be accomplished. 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 one has contributory expertise in a domain one also has interactional expertise, but, if one 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. What we mean by this is that, in order to realize the latent expertise, nothing new pertaining to the specific domain in question has to be learned. The things that have to be learned to realize the latent expertise are to do with the domain of talking, reflecting, translating, and so forth, not laser-building, or gravitational wave physics, or car driving, or whatever. 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 unreflective contributory expert by skilled and persistent probing—from skilled interviewing one 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.[27]

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 realize the interactional expertise it is also necessary to possess interactive ability.

A lack of realized 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 by 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 needed to interact with others are 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).

An important difference between interactional expertise and interactive ability is that the latter, unlike interactional expertise, is not parasitic—it 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 generalized abilities, not an expertise *in a special domain*. It is because interactional

27. The above explanation of how we use the term "latent" is a response to Evan Selinger who argues (personal communication) that you cannot have a latent expertise; a latent expertise is no expertise at all. Carolan (2006) also discusses these issues in the context of an interesting analysis of the role of interactional expertise in the development of farming skills.

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 anything 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 generalized 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 professionalized and specialized ability than interactive ability because it is taught, quite self-consciously, in sociology and philosophy courses and the other critical disciplines. Like interactive ability, reflective ability is enormously useful in the building of interactional expertise. It is sometimes what makes the difference between the analyst and the specialist scientist when they are talking about an esoteric domain. Some scientists are actually proud of their lack of reflective instincts, boasting that "Philosophy of science is about as useful to scientists as ornithology is to birds."[28] This is perfectly correct, but it 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 figure 3 represent respectively the *concepts* of contributory expertise, interactional expertise, and interpersonal and reflective ability (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 realized (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

28. Commonly attributed to Richard Feynman.

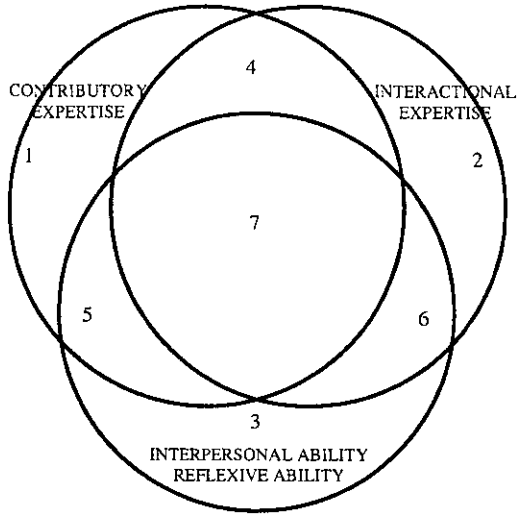

Figure 3. Relationships among some expertises

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.

Acquiring Expertise: Five Kinds of Face-to-Face Knowledge Transfer

What is involved in the acquisition of an expertise under our model? We take as our example the ability to build a working model of a piece of novel scientific apparatus. We ask the following question: Suppose scientist B from B-lab wants to learn to build one of those heavy and nontransportable A-mators so far only built successfully by scientists in A-lab. We know from previous studies that, barring reinvention "from scratch," the right thing for scientist B to do is visit A-lab for a period and spend time talking with the A group as they work on the A-mator. Alternatively, a group from A-lab can visit B-lab for a period and hold extensive discus-

sions. We then ask what scientist B learns from a visit.[29] Five problems related to knowledge transfer can be identified that could be resolved by extensive meetings between scientists from different laboratories:

1. *Concealed Knowledge:* A does not want to tell "the tricks of the trade" to others, or journals provide insufficient space to include such details. A laboratory visit can reveal these things but only if it is in the direction B to A.
2. *Mismatched Salience:* There are an indefinite number of potentially important variables in a new and difficult experiment, and the two parties focus on different ones. Thus, A does not realize that B needs to be told to do things in certain ways, and B does not know the right questions to ask. The problem is resolved when B watches A work so, again, a B to A visit is needed.
3. *Ostensive Knowledge:* Words, diagrams, or photographs cannot convey information that can be understood only by direct pointing, or demonstrating, or feeling. This can be accomplished by B visiting A-lab.
4. *Unrecognized Knowledge:* A performs aspects of an experiment a certain way without realizing their importance; B will pick up the same habit during a visit, while neither party realizes that anything important has been passed on. Much unrecognized knowledge becomes recognized and explained as a field of science becomes better understood, but this is not necessary.[30] Again, the B to A direction is best.
5. *Uncognized/uncognizable Knowledge:* Humans do things such as speak acceptably formed sentences in their native language without knowing how they do it. There are experimental counterparts to such abilities. They are passed on only through apprenticeship and unconscious emulation. At the end of the transfer neither party can describe what has been transferred; they may not even notice that anything has been transferred. Insofar as the emulation is of practices, it will be necessary, once more, for B to visit A.[31]

Concealed knowledge is to do with lies and secrecy, not the nature of knowledge transfer, so it is not central to this discussion. Turning to the other four problems, we have discussed them in terms of the transmis-

29. The elements of this discussion are originally worked out in Collins 1974 (or see 1992), and 2001a (or see 2004a, chap. 35).

30. For a discussion of the problem with reference to Italian violins, see Gough 2000.

31. For the first use of this list, see Collins 2001b.

sion of contributory expertise. But what if it was interactional expertise alone that was to be transferred? Would an A-group's visit to B-lab be as good as B's visit to A-lab? How important to language acquisition is it that the talk be conducted in the presence of the apparatus? In chapter 3 we will return to the problem, but it may be that we already know the answer. The point is that none of those things that A thinks are salient and B does not realize are salient (2), and none of the things that A can only point to (3), and none the things that A does without realizing it (4), and none of the things that no-one can articulate (5), figure in A's *language*. Therefore, if B is interested in acquiring only A's language, then B has no need to acquire the understandings that don't have any counterpart in the language. Therefore, at a first approximation, a visit of an A-group to B-lab should be just as useful for acquiring interactional expertise as a visit of B to A-lab.[32] This is but a first approximation, since it is probably the case that a visit of a single individual to B-lab is not going to be a very good way of transferring interactional expertise—this would be like our trying to learn fluent French from the visit to our house of a single French person; what would be needed would be extended visits of a group of scientists from A-lab to B-lab.[33]

If it is true that, in principle, so little in the way of knowledge is necessary for the acquisition of interactional expertise that it can be acquired without being in the presence of the apparatus, then interactional expertise is very limited. But that only makes it the more interesting, and, at first glance, counter-commonsensical, that interactional expertise is so important in science and technology. Remember, if, as per the strong interactional hypothesis, a person with contributory expertise cannot say anything in virtue of having that expertise that a person with interactional expertise cannot say, then interactional expertise is just as good in forums that work through the medium of language as contributory expertise.

It now seems a shame that the analysis of the immediately preceding paragraphs has not been tested when studies of scientific knowledge transfer have been carried out. For example, in an early study of the trans-

32. Even though this is a "first approximation," it is true only in principle. In practice, having an apparatus present is likely to make the conversation easier, even if it is only interactional expertise that is at stake. The point remains, and it will be discussed at length in chapter 3, that where circumstances do not permit proximity to the apparatus or its equivalent, interactional expertise can be acquired even though it will need extra effort.

33. We are grateful to an anonymous reviewer of the manuscript for pointing this out to us.

fer of laser-building ability no attention was paid to how much of what was being learned by one scientist in another scientist's laboratory was learned by talk and how much by watching and doing.[34] There would be very many confounding factors in a naturally occurring situation, but perhaps this kind of question could be the subject of an experimental study.

The same oversight can be found in a 1990 discussion by the same author. Collins asked how a spy, who pretends to be a native of a town he has never visited—Semipalatinsk—would be caught out in conversation with a native of Semipalatinsk. He claimed there were two ways for a spy who had never visited the town to learn about Semipalatinsk: learning from books, photographs, and so forth, and learning through conversation with an expatriate native of the town. He claimed that, in contrast, genuine natives would also learn from experiencing the physical reality of the town and from immersion in the linguistic culture. Collins argued that the spy would fail a "Turing test" when compared with a native if the interrogator was also a native because of the things natives could say to each other as a result of the second two kinds of experience but which the spy would not be able to say.

We can now see that Collins failed to consider the extent to which the spy could have been trained to pass the interrogation by immersion in the language of Semipalatinsk alone if, say, that immersion was arranged by his spending a year with a group of exiled Semipalatinskians. Such a group could have transferred the capacity to handle nuances of pronunciation or vernacular speech and all that can be acquired through the development of interactional expertise. On the arguments set out in respect of the transfer of scientific knowledge, physical experience of the town would not be important since everything that could be said as a result of physical immersion could be said by the expatriates. The spy, then, could have mastered the interactional expertise of Semipalatinskianism to the level that the expatriates still possessed it. The spy interrogation test, which is very like an imitation game (see below), is a test of interactional expertise not contributory expertise.

As it happens, Collins and Kusch, in their 1998 book, did better (though without realizing it at the time). They discuss the role of Cyrano in the play *Cyrano de Bergerac* by Edmund Rostand. Cyrano agrees to write love letters on behalf of another suitor. They write:

34. The study in question is Collins 1974 (or see 1992).

> . . . suppose Cyrano had never known love but, notwithstanding, had the skill to write the prose? We may imagine he had read the relevant literature and poetry and had frequented the society of those who did know love. In other words, Cyrano would have understood the institution of love and love letters even though he had never felt the individual emotion himself. (94)

In retrospect we can now describe this as an instance of the acquisition of the interactional expertise of love in the absence of the contributory expertise of love and an illustration of the individual embodiment thesis (to be discussed in chapter 3).

CHAPTER TWO

The Periodic Table of Expertises 2: Meta-expertises and Meta-criteria

We now turn to the lower section of the Periodic Table. Meta-expertises are expertises used to judge other expertises. There are two kinds of meta-expertise: *external meta-expertise,* which does not turn on acquisition of the expertise itself, and *internal meta-expertise,* which does involve an acquaintance with the substance of the expertise being judged.

External Judgment: 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 judgment the boundaries extend to the very edge of the general public. Just as there is ubiquitous expertise, there is also ubiquitous meta-expertise, which we will call "ubiquitous discrimination." In such a case the Problem of Extension dissolves, as with other ubiquitous expertises. Ubiquitous discrimination, like other ubiquitous expertises, is acquired as part-and-parcel of living in our society.

For example, those with little scientific knowledge can sometimes make what amounts to a *technical* judgment on the basis of their *social* understanding. The judgment turns on whether the author of a scientific claim appears to have the appropriate scientific demeanor and/or the appropriate location within the social networks of scientists and/or not too much in the way of a political and financial interest in the claim.

Ubiquitous discrimination is what we have all been learning since we could speak, and it is just a particular application of our regular judgments about friends, acquaintances, neighbors, relations, politicians, salespersons, and strangers, applied to science and scientists within Western

Table 2: The lower part of the periodic table of expertises

META-EXPERTISES	*EXTERNAL (Transmuted expertises)*		*INTERNAL (Non-transmuted expertises)*		
	Ubiquitous discrimination	Local discrimination	Technical connoisseurship	Downward discrimination	Referred expertise
META-CRITERIA	Credentials		Experience		Track record

scientific society. For example, many members of such a society, just by being members, are able to discriminate between what counts as science and what does not. This is the ubiquitous judgment on which we rely when we dismiss certain beliefs such as astrology from the list of contributors to the *scientific* element in technical decision-making. Most members of our society have sufficient social understanding to know that the standards and social and cognitive networks of astrologers do not overlap with the standards and social and cognitive networks of scientists.[1]

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.[2] 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. The astronaut David Scott reports his reaction at the time as follows:

> 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?"

1. Poor social judgments 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 chapter 5 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.

2. Susan Carter brought the relevance of the Moon landings to our attention.

> 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.[3]

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

What stops us doubting their validity is, once more, ubiquitous discrimination, which is a social judgment. 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 organized to lie so constantly and consistently; we know that if there were any possible credence to the story of the fake, the Russians, deeply involved in the cold war as they were, would have exploited the doubts—yet they did not. Our ubiquitous social discrimination allows us to be sure about the Moon landings, even while the technical discrimination of 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. As the cold fusion saga drew to a public conclusion around the last years of the twentieth century, most reasonably literate members of Western society, who knew nothing of cold fusion beyond what they had seen on the news or read in the newspapers, "knew" that cold-fusion had been tried and found wanting. Though there was a time when cold fusion was contiguous with science as we knew it, it was now understood that its cognitive and social networks no longer overlapped with those of legitimate scientific society. This knowledge had nothing to do with scientific competence. On the contrary, it was vital to ignore scientific credentials, and even track records of success, if a socially appropriate judgment was to be made. Thus Martin Fleischman, the co-founder of the cold fusion field, had an enviable track record for success in the sciences, was immensely well qualified, was honored as a Fellow of the Royal Society, and had both interactional and contributory expertise in cold fusion, yet still believed in the effect, contrary to the scientific consensus. What people in Western societies had in common is what they had heard about cold fusion in the broadcast media (popular understanding). Their consensual view, insofar as they had one, emerged from making *social* judgments about *who* ought to be agreed with, not

3. Scott and Leonov 2004, 124.

scientific judgments about *what* ought to be believed. To expect the citizen to be sufficiently educated in science as to be able to make scientific judgments in disputed esoteric matters like this was the futile 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.)[4] The crucial judgment, however, is to "know" when the mainstream community of scientists has reached a level of 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 judgment. Note that this is not the sort of judgment that we would expect even an immaculately qualified scientist from "another planet" to be able to make. A scientist from another planet, reading published papers for and against cold fusion, would have difficulty working out who was right; the scientifically ignorant citizens of this planet, in contrast, had a relatively easy decision to make.[5]

External Judgment: Local Discrimination

There is a quite different version of discrimination that pertains to more 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 1986. Wynne's work has been influential in both a positive and negative way. It has helped to establish the idea that technical expertise can be found beyond the normally recognized qualified groups, but it has also given rise to much confused thinking about the extent to which laypersons can be experts. Wynne looked at the interaction between the UK Ministry of Agriculture Food and Fisheries (MAFF) scientists and the Cumbrian sheep farmers after radioactive fallout contaminated their pastures.[6] Wynne argued that the expertise of the sheep farmers in respect of sheep ecology should not have been ignored by the scientists—the sheep farmers had what he called "lay expertise." Though spotting the expertise

4. Which is the flaw in the "deficit model" of public understanding of science.

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

6. Wynne 1989, 1996a, 1996b.

of the sheep farmers was insightful, "lay expertise" was an unfortunate choice of term because of its potential to cause confusion. For example, the term has often been interpreted as meaning that laypeople possess specialist expertise. It would have been better if Wynne had talked of experts without formal qualifications. For example, the sheep farmers are not laypersons; they are experts in sheep farming who happen to have no paper qualifications. The sheep farmers have a specialist contributory expertise. Their expertise is esoteric, highly relevant to the ecology of sheep on radioactive fells, but, unfortunately, it went unrecognized by Ministry scientists.[7] As a group of contributory experts, discussion of the sheep farmers' expertise belongs in chapter 1.

But Wynne also found that in addition to contributory expertise the farmers possessed what we will call "local discrimination," which is more properly a topic for this chapter. Soon after the Second World War, the Windscale-Sellafield nuclear processing plant was built on the Cumbrian fells, and the farmers thus 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 been able to judge the pronouncements with such finesse. The sheep farmers were able to discount statements by the nuclear industry's spokespersons, but in this case it was a result of local experience rather then a more generalized discriminatory ability developed over a lifetime of social and political education. Local discrimination in such a case is analytically, and sometimes practically, distinct from contributory expertise. For example those long-term residents of the area surrounding Windscale, who knew nothing about sheep farming, almost certainly knew quite a lot about the nature of the assurances offered by the representatives of the local nuclear industry.

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."[8] Wynne

7. Or so we understand.
8. Wynne 1992, 39.

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."[9]

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 local 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 the 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 [our] social relationships and institutions." This 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 "esoteric" 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 as a result of long experience in their particular workplace—then that is a different matter.[10]

9. Wynne 1993, 328.

10. Gravitational wave scientists report that they use the following criteria to judge whether an experiment by another scientist needs to be taken seriously: faith in experimental capabilities and honesty, based on a previous working partnership; personality and intelligence of experimenters; reputation of running a huge lab; whether or not the scientist worked in industry or academia; previous history of failures; "inside information"; style and presentation of results; psychological approach to experiment; size and prestige of university of origin; integration into various scientific networks; nationality.

Problems of External Discrimination

Local discrimination, like ubiquitous discrimination, is a meta-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*. It is, as explained at the beginning of this chapter, a way of reaching technical conclusions via nontechnical means. In general it is very unreliable because of the temptation to read too much into stereotypical appearances and stereotyped behavior. It was this tendency to read too much into appearance that was exploited by the "scientists in white coats" who, for many years, assumed and were given a license to speak with authority on almost any subject. It could be said that the stereotype of the scientist was what gave rise to the misleading picture of the power of logical thought and experimental genius. It could well be argued that the public's misunderstanding of the MMR controversy (see chapter 5) was partly a product of the demeanor of Andrew Wakefield—a young, handsome, kindly and, seemingly caring doctor up against the establishment. Again, the remarkable Dr Fox Lecture should warn us against judgments based on demeanor. Dr Fox was an actor hired to present a lecture to university audiences. The content of his talks were impressively technical-sounding gobbledygook. Answering a questionnaire after the performance, large proportions of the audience expressed themselves well satisfied with what they had heard.[11] Of course, there is also an equally misleading counterstereotype: the mad, monster-creating scientist which gave rise, at the time of the controversy over genetically modified organisms, to the famous British newspaper headline referring to "Frankenstein Foods." Nevertheless, judgments on the basis of demeanor and social position are often made within science (see footnote 10), so we should not dismiss them out of hand when they are made by the public. Furthermore, there are restricted and special circumstances where public discrimination seems sound. Some examples have been provided above: the Moon landings, cold fusion, and circumstances where members of the public were able to induce from past misleading statements that current statements about safety and the like were also likely to be misleading or, perhaps, vice-versa in the case of Wynne's apprentices. It should be noted that these examples turn on either a fairly good understanding

---

The first, second, sixth, eighth and tenth items could be said to be matters of local discrimination with the remaining items matters of ubiquitous discrimination.

11. Naftulin, Ware, and Donnelly 1973.

of social life or enough relevant experience to give rise to fair inductive inferences and do not imply the reliability of less well-founded discriminations. When the conditions are met, we seem to have what we might call "transmuted expertise"—a transmutation of social knowledge into technical knowledge. What we can say for certain, a point we will return to in the final chapter, is that discrimination of this kind is not part of the *legitimate* methods of science. When scientists discriminate in these ways, they do not trumpet the method in publications nor in any other part of the "constitutive forum."[12] Likewise, when the method is used by the public, it does not, and should not, be accepted as a legitimate input to scientific method.

Internal Judgments of Expertise and Their Problems

What other ways are there of judging between experts? The standard method of choice is by reference to the qualifications of the expert.[13] This we now know to be inadequate because it is possible to have expertise, and that includes specialist expertise, in the absence of qualifications. Thus, Wynne found that the uncredentialed Cumbrian sheep farmers knew a great deal about the ecology of sheep, the prevailing winds, and the behavior of rainwater on the fells that was relevant to the discussion of how sheep should be kept so as to monitor and reduce the impact of the radioactive fallout.

A still better documented case of uncredentialed persons gaining expertise is Epstein's study of AIDS treatment activists in the 1980s.[14] In 1985 a new drug, AZT, was about to be subjected to double-blind randomized control trials. AIDS sufferers were concerned that too many of those who were assigned to the placebo groups 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

12. Collins and Pinch 1979.

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

14. Epstein 1996. See also Collins and Pinch 1998, or 2005 for a summary of Epstein's study.

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 1996, 116). The activists, however, undertook an arduous course of self-education and learned the language of medical discourse. Most importantly, 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 randomized control trials. They knew that these demands were unrealistic: since death was in constant prospect, the groups regularly smuggled untested cures from Mexico, continued to take other drugs which were banned and which had the potential to confound test statistics, and even shared placebos and trial drugs between experimental groups. ACT UP knew that the randomized control trials were not working as the scientists assumed they were.

Eventually, the activists gained so much interactional expertise in research design that, allied with their experience, they were able to make real contributions to the science that 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 an "unbelievably high" standard. He said: "It's frightening sometimes how much they know and how smart some of them are" (Epstein 1996, 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, not least because it enabled them to do better science.

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.

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 unfit for the workplace, where a different kind of expertise—the level of expertise associated with doing rather than knowing—is required. It is at these higher levels of expertise that the problems of judging expertise, both practical and conceptual, arise. We begin our investigation of the deeper problem by looking at 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 Roman Catholic priest, and bogus CIA men. It seems, then, given the right conditions, people are ready to attribute almost unlimited expertise and authority to people whose behavior can, given a stretch of the imagination, be interpreted as indicating that they have it.

The phenomenon has been marvelously exploited in fiction. To pick one example among many, Jerzy Kosinski's book, *Being There* (1971) subsequently made into a film of the same name, shows an educationally subnormal, sexually neuter, but well-dressed gardener, the eponymous "Chauncey Gardiner," rise from unemployment to become president of the United States complete with a reputation for innovatory sexual adventure. It happens because a band of hangers-on reinterpret his minimal vocabulary as the profound and gnostic discourse of a sage and rake.

In relational theories of expertise there is a problem about how to deal with frauds and hoaxes. If expertise is attributed to a hoaxer, there is little more to say about it—the relevant topological location in the network has been achieved. The problem for a relational or attributional theory is what is special about hoaxes and frauds (before they are exposed), as opposed to genuine exercises of expertise. Indeed, this problem has been made into a virtue—the study of the attribution of the label "hoax" being the very point of some analyses.[15] Since we are concerned with judgment 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 mis-

15. See, for example, Brannigan 1981.

take about, which are 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, there are no confidence tricks in which a solo performance with an orchestra playing a well-known 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. More recently, and uncomfortably "close to home" for some 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 clichés 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 may have exposed lax editorial practices at *Social Text*, but it reveals little more since hoaxes are not so hard to pull off even in theoretical physics.[16]

16. To be 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. 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.

Note that there is a difficulty for those who would want to defend a journal with a "postmodernist 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.

Hoaxes and frauds are more easy to carry off than they should be because of the well-known 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.[17] 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 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 nonmelodiousness of which is the very point, since it asks questions about the meaning of music.

The role of avant-garde artist has been nicely satirized in the 1961 comedy film, *The Rebel*. 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.[18] We are amused but not sur-

17. Collins and Pinch (2005, chapter 1), discuss fakes in general and a case study of bogus doctors carried out by Collins and Hartland. Bogus doctors are rarely unmasked as a result of medical mistakes because the surrounding team is ready to "cover" for them. Novice doctors, even though they have been through medical schools, have no experience and are expected to be unfamiliar with hospital conventions. As a result many bogus doctors have the opportunity to learn on the job and become quite successful. Experience, what we consider to be the most reliable criterion of expertise, stands up well when confronted with the bogus doctor case. A bogus doctor who survives long in the profession becomes a skilled doctor, even though he or she may have no certificates. As a result medical professionals who work with bogus doctors can be astonished when they discover that they have been fooled:

> I've never been so shattered in my life when a nurse came up to me and said the CID [the detective branch of the police] had been there . . . and I said, "What for?" and they said, [Carter] "had never qualified." I felt as if I had been hit with an atom bomb. . . .
>
> Had we been asked, and this was the general opinion of everybody when this came out, had we been asked to pick a doctor who was bogus, he would have been the very last of them all. (Quoted in Collins and Pinch 2005, 47)

18. *The Rebel* (*Call Me a Genius* in the USA) 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

prised at this outcome because we know that with the avant-garde, as the very name implies, there are no established conventions of artistic practice upon which to base a judgment.[19] 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 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 well-known avant-garde artists such as 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 judgment.[20]

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

---

without artistic talent (successful art forgers must be technically accomplished), whereas that may not be true of the presidency because of the way the president is embedded in a body of advisors.

19. The notion of the avant-garde helps to explain how bogus doctors maintain their presence in medical settings before they have learned on the job. Novice doctors are drawn from training hospitals in many different countries, and medical practice is sufficiently open to allow for a degree of variation which, while it may not be so great as in art, still provides a lot of leeway.

20. When all the virtuosity seems to lie in establishing a new convention rather than in executing a skill, art seems to be reduced to marketing, and it is not unreasonable to feel uncomfortable. Thus, we do well to be concerned that the Saatchi brothers, who run what was 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 1992). The hostile reaction is quite unreasonable to, say, sociology of scientific knowledge as a methodology, however.

from epoch to epoch.[21] Nevertheless, *within a stable convention* virtuosity can be recognized.

The ability to recognize skillful practice itself improves through practice and that this is the case is recognized in terms such as "connoisseurship."[22] Connoisseurship is a meta-expertise. A connoisseur is, according the Chamber's Dictionary, "a person with a well-informed knowledge and appreciation." The dictionary definition tells us that the knowledge and appreciation are usually applied to fine food, wine, or the arts. But connoisseurship—that is, judgment 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 contractor 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? 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, to go around curves, should the tiles be cut with a diamond saw, or cut square to an approximation to the curve, or chipped and hacked, the gaps being filled with grout? Some-

21. 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).

22. Goodman 1969 is relevant here but will be discussed at in chapter 5. Carlo Ginzburg's paper "Morelli, Freud and Sherlock Holmes: Clues and Scientific Method" (1989) 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 seems confused. Physical and biological sciences, the identification of the provenance of paintings, and the method of Sherlock Holmes are all typically scientific activities: though they deal with specific instances, these instances are specific instances of general laws. In other words, general laws are applied to the specific cases 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 (1957) calls "historicism."

how one must "negotiate" with the tiler over what will count as a satisfactory job in terms of both the formalized standards and the unformalized conventions. Interestingly, one may employ a professional—an architect—to do the negotiations. The fact that it is possible to employ a professional who may never actually have done any tiling to make these judgments shows that the crucial thing is experience within the conventions of judgment rather than experience of the skill itself. There is a connoisseurship of tiling. The judgment being exercised by architects, or homeowners, who themselves may not be capable of tiling (who have no contributory expertise) but who have seen and discussed many bathrooms, is based on *interactional expertise.*[23] Interactional expertise is the bridge between full-scale physical immersion in a form of life (which gives rise to contributory expertise) and non-expert acquaintanceship with the idea of tiling and the discourse pertaining to it. Interactional expertise enables architects to speak to both tilers and homeowners. The strong interactional hypothesis posits that by being immersed in the language community alone one may learn to "know what one is talking about" even if one cannot do the corresponding activity.

Is this, then, what Lord Campbell was getting at in his remarks about TV produces and sugar workers quoted in the introduction? No! Lord Campbell's view was different because he implied that he did not need interactional expertise or specialist experience of any kind (outside of management) to make the judgments. Just as the general public cannot have expertise in all domains of specialization, neither can any single person. There may be some who would claim that refined judgment in all things is the inheritance of the members of a well-bred aristocracy, but if the aristocracy does 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. The notion of connoisseurship does not, then, safeguard a Lord Campbell-type view, though it does make safe the idea that it is at least possible to judge an expertise without being able to practice it.[24]

23. We are grateful to Kevin Parry and Mike Bergelin for providing a lived example of tiling. 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.

24. See Shapin and Schaffer 1987 for the class basis of what was counted as legitimate "witnessing" in early scientific experiments.

We can now see more clearly why it was that in order to be sure to recognize 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 in this way would the foundation of ubiquitous, or at least relatively widespread, experience on which judgment 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; in these areas it is considered that only those with contributory expertise should judge those with contributory expertise. But, and this may be little more than a truism, the medium of judgment, 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 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; rather, the reviewer is exercising interactional expertise—the ability to talk or write about gravitational wave physics.[25] 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 realized 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 judgment. A transitive relationship works one way; the possession of interactional expertise does not guarantee the possession of contributory expertise, but, according to the strong interactional hypothesis, someone who possesses it in full ought to be as good a judge of the contributory expertise to which it pertains as someone who has

25. Of course, in one of the wider senses mentioned at the end of chapter 1, this is a contribution to gravitational wave physics.

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, as we have seen from Epstein's study) 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, if 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, say, a decision-making panel.[26]

A still stronger claim is sometimes made by art critics and the like. They say that a level of connoisseurship (which, as we have established, is itself based on interactional expertise) can be developed through assiduous viewing and discussing of art. They sometimes claim that this makes for judgments 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 sometimes 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."

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 (realized) 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* judgments made by one expert about another are always contestable.[27]

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 judgments about expertise. If it were not, none of

26. Here, as we shall argue in chapter 5, lies an important difference between the sciences and at least some of the arts. In the arts, the locus of judgment favors realized interactional expertise and so it favors interactional and reflective ability more highly than do the sciences.

27. It is the very contestability of such judgments that make relational theories of expertise so attractive.

the comparisons we have discussed in the section on hoaxing and faking would make sense. In the absence of internal judgment it would make no sense to say that solo violin-playing is more difficult to fake than ability in avant-garde art because the difference between skillful and unskillful performances of all kinds would be 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* judgments about expertise that are more likely to be right than wrong?

The principle toward which we are working is what we will call "downward discrimination." We claim that judgments can be made within a discipline even though the judge's expertise within the discipline is very low; it can be done when those being judged have a recognizably lower level of expertise. Consider a claim made by a participant in the UK debate about the safety of genetically modified foods. The participant insisted that since radioactive tracers are used to mark genes during genetic manipulation then the genetically modified foods that result would themselves be radioactive. Nearly all readers of this book will understand that this claim is incorrect. They will be able to exercise downward discrimination even though, in many cases, it will be on the basis of almost no specialist knowledge of the particular science under discussion. It is just that the person making the claim has recognizably less knowledge.[28]

There is an important difference between external and internal judgment in this regard. External judgment does not have a preferred direction: it can be applied equally well upward, downward, or horizontally. That is to say, an ordinary person can reasonably distrust the demeanor or interests of even highly technically experienced and highly qualified spokespersons for the nuclear industry, or the tobacco industry, or any other industry, even if he or she is in no position to question their claims on technical grounds. With internal judgments the epistemological problem means secure judgment can run only downward, not upward nor

28. We are grateful to Matthew Harvey for this example, which is taken from his fieldwork observations of the UK debate.

horizontally. Where the direction is horizontal, there are only arguments and negotiations.[29]

Because it is impossible to make a technical assessment of the technical understanding of an expert with more expertise, those downwardly discriminated against may not recognize the validity of the judgment; higher levels of expertise may not be recognized for what they are. Well-founded downward technical discrimination is all too easy to confound with bias by those on the receiving end of negative judgments.[30]

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 a critique of the public understandings of GM foods we can 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 can criticize 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 cannot criticize the suggestion that insufficient testing has been done to guarantee that herbicide resistance will not spread to weeds and the like; in the latter case there is no settled scientific consensus to draw on.

To sum up, we all tend to believe we can make internal judgments of expertises upward, downward, and horizontally. The sociology of attribution is the study of the way actors negotiate the right to judge expertise; public legitimacy can be assigned to judgments made in any direction, and those judgments 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 upward judgments while reducing the potency of downward judgments. The normative view that we are developing here is that internal technical judgments, which are of a good enough quality to contribute to science and technology policy, can be made only when they run downward.

29. 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 recognizing 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.

30. All the judgments we describe may, of course, be wrong. That is the nature of judgment.

Referred Expertise

Another kind of reasonable internal judgment is based on referred expertise. Referred expertise is expertise taken from one field and indirectly 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. Thus, Gary Sanders was first a professor in the field of high energy physics, then became the project manager of LIGO (the Laser Interferometer Gravitational-Wave Observatory), which turned on the very different science of interferometry, and at the time of writing has taken up the post of director of a major new telescope-building project—again a very different field. He remarked to Collins: "They give you the keys to the Thirty Meter Telescope on Day One and say, 'Drive it.' I found myself making key design decisions, not really knowing the history, the lore, the tradition, the lessons learned in the telescope." He added: "I'm not an observing astronomer. I have to listen to the arguments of 'planets versus galaxy formation versus stellar populations,' and 'this instrument should be a first light instrument and that instrument should be first,' and the campaign has already begun. . . . In the end, guess what?—The guy who's never spent a night on a mountain opening the shutter and doing an astronomy observation is going to say 'I selected [this approach] and these are my reasons.' What the hell is that?" Sanders explained the way he had learned during his first eighteen months on the project, using the vocabulary he had discussed with Collins:

> I was concerned that I just would not understand it. But I've found that, remarkably, what you call interactional expertise was not hard to achieve. I couldn't design an adaptive optics system but I really do, after six to nine months in the field, I really do understand the different kinds of adaptive optics and the way that they work and I can draw a schematic and define the algorithm, and understand the technological readiness of the different techniques—which ones are really ready to apply to the sky and which ones need to be demonstrated and certain components have to be developed. . . .
>
> I can sit down with a bunch of adaptive optics experts who will come to me and say "Gary you're wrong—multi-object adaptive optics will be ready at first light and it will give the following advantages . . ." and I shall say "No, it's multi-conjugative adaptive optics" and I can give them four

> reasons why we should go with multi-conjugative adaptive optics based on the kind of science we want to do, the readiness of the technical components, when we need them, and so on, and I will see as I am talking about it that the room is looking back at me and saying "He does have a line, he's thought it through."
>
> [But] if someone said to me, "OK Sanders, we agree with you, now go and design a multi-conjugative adaptive optics system," I couldn't do it. I couldn't sit down and write out the equations. . . . But I can draw a diagram of what each part does, where the technological readiness of each one is—what the hard parts are—I know the language and I actually feel qualified to make the decisions.

Looking back to his period at LIGO, he said:

> I can't design the LIGO interferometer. I can't sit down and write down all the transfer functions and work out the noise budget like [named scientist] can. But if he gave a talk on it I could follow it. I can understand the important parts and the hard parts, partly by listening and partly by quantitatively understanding, but I couldn't come back and compose the symphony. But I was in a position where I had to decide. So it's a matter of who I listen to and which parts seem like they carry the argument—what it is that we want. . . . That's more than interactional but it's not quite contributory in, I think, your usual sense of the word.[31]

In most specialist domains in the field they have to manage, the managers, then, have interactional expertise but not contributory expertise.[32] 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—as Sanders says, there is something going on that is a bit more than interactional expertise. The resolution seems to be that, although, as we can see, contributory expertise is not required to manage even the science of a scientific project, management does need kinds of expertise that are referred from other projects. The managers must know, from their 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

31. Interview, 22 October 2005, Laguna Beach.

32. For examples of disagreement over whether managers from high energy physics had the competence to manage LIGO scientists, see Collins 2004a.

scientists they are leading at one remove, as it were. Managers of scientific projects with referred expertise would manage better (as well as with more authority and legitimacy) than those without it.[33]

The experience in other fields is applied in a number of ways. For example, in the other sciences they have worked in, they will have seen that what enthusiasts insist are incontrovertible technical arguments turn out to be controvertible; this means they know how much to discount technical arguments. They will know how often and why firm technical promises turn out not to be delivered. They will know the dangers of allowing the quest for perfection to drive out the good enough. They will have a sense of how long to allow an argument to go on and when to draw it to a close because nothing new will be learned by further delay. They will have a sense of when a technical decision is important and when it is not worth arguing about. They will have a sense of when a problem is merely a matter of better engineering and when it is fundamental. Interactional expertise is the medium through which this kind of expertise is made referable from one field to another.

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.[34] 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 manage anything. 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." Does this work for Lord Campbell's position? Again, it probably does not because making sugar and making television programs are not closely related in this sense. Indeed, that is the very point of Lord Campbell's outburst.

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; some of these will comprise the contributory expertise of management itself. Crucially, a manager of a scientific or technological project will need local discrimination; they

33. Though in the case of LIGO, 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: "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" (quoted in Collins 2004a).

34. See Thorpe and Shapin 2000.

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 as specialists as well as judging between their arguments.[35]

Meta-criteria: Criteria for Judging Expertises

Our goal, as explained at the outset of this discussion, is to find ways 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 such as certificates attesting to past achievement of proficiency. Possession of certificates will define a number of kinds of expert, but note that there are not credentials that indicate 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 judgment, nor for political judgment. 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 architect). Above all, there are no credentials for 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.

35. 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 1992; for a scientist's 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 2004a). For further analysis of exactly what it is that the managers of big science projects do, see Collins and Sanders 2008 (forthcoming).

Track Record

Track record is a better criterion than credentials. The philosopher Alvin Goldman argues that track record of success in making sound judgments is a way for laypersons to choose between experts.[36] 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. Even when we get to qualified scientists and technologists, disputed expertise often concerns new fields for which there are no track records, fields in which track records take decades to establish, and fields in which the meaning of success is ambiguous. Track records, then, are only sometimes better than qualifications, and the "sometimes" are likely to be those where disputes are shallow rather than deep.

Experience

A criterion that does seem to set the boundary in a better 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 judgments 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 on "a transformative Hermeneutics of Quantum Gravity" than superficial peer review, and the same applies to Chauncey Gardiner, Hancock's avant-garde artist, and any number of confidence tricksters.[37] Confidence tricks and other such scams work when experience is attributed to the fraud on the basis of short acquaintance, whereas knowledge of their lack of experience would elimi-

36. Goldman 2001. See also Kusch 2002 on testimony.
37. See Sokal 1996.

nate them from the class of experts. (Though, of course, if they did have a lot of experience, that would not guarantee that they *were* competent.)

Periodic Table of Expertises Summarized Again

Let us now summarize what we have said about kinds of expertise. Once more we read down The Periodic Table of Expertises—Table 1.

Ubiquitous Expertises are acquired by all members of human societies in 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 realized 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;[38] the latter is *contributory expertise.* The bridge between experts with contributory expertise and people who are not experts in the domain 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 judgments are made. In logic, there is a transitive relationship between the five specialist expertises, though it may not always be realized.

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 does 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

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

lower level. *Technical connoisseurship* turns on interactional expertise alone, which may have been specially refined for the purpose as in the case of certain kinds of professional or critic. *Referred expertise* depends on the indirect application of domain-specific contributory expertise from one domain within 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.

Problems of Categorization

The Borderline between Interactional and Contributory Expertise

We now have to deal with a conceptual problem. Interactional experts are continually making contributions to sciences in which we say they have no contributory expertise. Examples include the contributions of philosophers and sociologists, who have never written a computer program, to the science of artificial intelligence; the contributions of social scientists and statisticians, who have never examined a fingerprint, to fingerprint identification; the contributions of project managers, who have never designed or built an interferometer or a telescope, to the technology of interferometric gravitational wave detection and the design of large telescopes; and, more generally, the contribution of experts in science studies to scientific and technological debates in the public domain. Indeed, we will argue that there are cases when the potential contributions of those with interactional expertise but no contributory expertise are not sufficiently well recognized. So, when do the contributions of interactional experts turn them into contributory experts?[39]

A first step in the analysis is to distinguish between "making a contribution" and "being a contributory expert." Thus, those who drill the wells for the Shell Oil Company make a contribution to those who work in science studies by providing the fuel that gets them from home to office. Nevertheless, the well-drillers are not contributory experts in the field of science studies-they are contributory experts in oil production.

It may be that this is too easy a case because the contribution of the well-drillers does not touch on the core discipline of science studies. We can see this because their contribution would be the same irrespective of

39. This section arose out of the persistent queries of Simon Cole. Similar points are made in Selinger and Mix 2004.

the domain of the experts to whom the fuel was delivered; Shell's employees make similar contributions to nuclear power station engineering, ballet dancing, zoo keeping, and so forth. Nevertheless, the idea of "making a contribution" can also cover cases where there is a more direct link to the core domain but the relationship is marginal, or sporadic. Thus, based on his interactional expertise, Collins has occasionally made suggestions—rarely taken up but at least provoking discussion—about aspects of the science and technology of gravitational wave detection.[40] But a characteristic of interactional expertise is that it is parasitic on contributory expertise rather than freestanding, and Collins's rare contributions seem more like this than the kind of self-sustaining practical experience that could be passed on to others. As we have argued, even if there were a whole body of interactional experts like Collins (and nowadays there do seem to be whole communities of social scientists parasitical on the new genetics), their understanding and discourse would diverge from that of the practitioners as time passed unless it was continually maintained and refreshed by contact with the world of practice to which it refers.

Some of the other cases mentioned above might, however, have more of a contributory component. A discussion with Simon Cole about fingerprint identification revealed some of the possibilities. Simon Cole, an expert in science studies, has appeared as an expert witness in fingerprint cases where his evidence has been called to throw doubt on the certainty of identifications made by fingerprint examiners. In the cases in which Cole has been involved, his "book learning" has frequently been unfavorably compared to the practical experience of the fingerprint examiners. Cross-examinations of Cole exemplify the problem. For example:

Q: Your working knowledge of latent prints is actually minimal, isn't that right?

A: My knowledge is in how the profession developed and what's in their literature.

Q: I am going to ask the question again: Isn't it true that your working knowledge of latent prints is minimal?

A: If by that you mean by knowledge of how to examine latent prints and make comparisons the way that fingerprint examiners, do, yes it is minimal.[41]

40. Actually, on at least one occasion a suggestion was taken up.

41. People v. Hyatt, #8852/2000, Tr. Trans. 37 (Sup. Ct. N.Y. Kings Co. — Part 23 Oct. 4, 2001). The case is discussed at length in Lynch and Cole 2005.

In this case the court concluded: "What Dr. Cole has offered here is 'junk science.'"

In our language, for Cole, the problem represented by this passage of discourse is that the court recognizes only the practical expertise of fingerprint examiners as making a contribution to the domain of fingerprint identification. The argument is not merely a matter of legal expediency—the practical experts believe in their craft. Thus Cole reports that, at a conference, he asked a relatively friendly fingerprint examiner how he knew what he claimed to know. The reply was along the lines: "I wish you could come to my laboratory and learn to do what I do and see what I see, and then you would see why I know that I know what I claim to know."[42]

What we would like to bring about is the establishment of a discourse of expertise that would enable Cole, if he wished, to replace his defensive responses under cross-examination with a confident: "I do not have contributory expertise in the matter of fingerprint identification but I do have interactional expertise in that domain and this enables me to make a contribution." In due course we may imagine it becoming the ordinary occurrence for interactional experts to be allowed to speak alongside contributory experts.

Now, setting Cole's expertise aside for the moment, consider statisticians who believe the have something to say about the likelihood of a fingerprint match being correct. Like Cole, their warrant for claiming that they have something to offer can only have to do with their understanding of how fingerprinting works. If the statisticians did not know quite a lot about fingerprint identification, they would not be in a position to argue that their expertise was relevant to the procedure of the courts. Their warrant, then, turns on their interactional expertise in fingerprint identification practice. What the statisticians want to bring to court is, however, a self-sustaining contributory expertise: statistics can be taught in classrooms and transmitted from generation to generation in the absence of intimate contact with any realm of practice to which it might be applied. The case is like that of Wynne's sheep farmers, who possessed a relevant contributory expertise but insufficient interactional expertise to ensure that it was recognized, and that of Epstein's AIDS sufferers, who did develop the interactional expertise to make sure they were heard.

When SSK is applied to artificial intelligence, the same thing seems

42. These are paraphrases generated by Cole's recollection of the conversation.

to be happening. A contributory expertise which provides understanding of the social nature of knowledge is being brought to bear on AI, but it is interactional expertise in the practice of AI that puts the outsider in a position to argue that the social analysis of knowledge should be brought to bear. The interactional expertise, then, makes a contribution, the contribution being to establish the value of a novel contributory expertise in respect of the esoteric domain. The role of the interactional expertise is to argue for the value of the new contributory expertise.[43]

Such a process can lead to the novel contributory expertise coming to be a regular part of the esoteric domain. But nothing philosophically profound happens when such a transformation takes place; it is just that the boundaries of the domain have shifted. The statisticians, for example, still have only interactional expertise in the practices of the domain with respect to which they were interactional experts in the first place. They are in no more anomalous a position than, say, a contributory expert in the calculation of the strength of gravitational wave emissions from inspiraling binary neutron stars and a contributory expert in the quality of coatings on interferometer mirrors who have only interactional expertise in each others' narrow domains even though both can be described as contributory experts in the wider domain of gravitational wave detec-

43. In the early days of critiques of AI by outsiders, the AI community appears to have been highly resistant to "outside interference." Philosopher Hubert Dreyfus wrote the first and, arguably, still the most definitive book-length critique in 1972 (see also 1992). He reports that Marvin Minsky and Seymour Papert among others (central figures in AI) attempted to prevent him gaining tenure at MIT in consequence. The intervention of MIT's president was needed to rectify the situation (email to Collins, 5 December 2003). Subsequently, however, Dreyfus was invited to advise the US military on their AI projects, and the AI community has become in general much more open to outside criticism. For example, anthropologist Lucy Suchman, who wrote a very well known (1987) critique of AI while at Xerox PARC, subsequently remained there and built up a research group *within* the organization. In 2002 she was awarded the Benjamin Franklin Medal in Computing and Cognitive Science, sandwiched between awards of the medal in 2001 and 2003 to two of the most prominent figures in the community itself, Marvin Minsky and John McCarthy. Harry Collins was encouraged to develop his AI work (e.g., 1990, Collins and Kusch 1998), when on his very first attempt to critique the AI community (Collins, Green, and Draper 1985) he was given a share in the prize for best *technical* paper at the 1985 meeting at the British Computer Society Specialist Group in Expert Systems. Collins was subsequently invited to sit on an ESRC review committee distributing grants for research in AI and a Ministry of Defence panel. Of course, the heartland of AI, which is designing and/or building programs, easily absorbs less critically minded philosophers and was founded by psychologists among others. Also, in contrast to fingerprint identification, it has a tradition of internal critique (Weizenbaum 1976; Winograd and Flores 1986), perhaps because it is a university-based rather than a craft discipline (but it should be borne in mind that this did not help Dreyfus in the early days).

tion.[44] It is the interactional expertise of one narrow domain specialist in another domain specialist's expertise that makes it possible for the larger field of gravitational wave detection to exist—otherwise it would just be collection of isolated groups of specialists. To refer to a member of one of these narrow groups as a contributory expert in gravitational wave detection is a matter of choice of focus not of ontology of knowledge, because it is what you say you are an expert *in* that determines whether that expertise is interactional or contributory. What a contributory expert can be said to be a contributory expert *in* is, then, to some extent arbitrary, because what is counted as a "domain" is to some extent arbitrary. But this does not mean that an interactional expert in some narrow practical domain can become a contributory expert in that same domain just by changing the attribution.[45] The above argument can be represented in the cartoon (figure 4).

In the cartoon, circles with spikes are areas of specialist contributory expertise, such as mirror polishing, the calculation of wave forms, or fingerprint identification. The irregular lines are the boundaries of domains of expertise. The left hand domain is gravitational wave detection physics. It contains many specialisms linked to each other by their members' interactional expertise (the patches of conversation). A purely interactional expert like Collins is shown as a stick figure with the ability to produce conversational performance pertaining to the field which is indistinguishable from that of the others, but no part of Collins is found within any of the contributory expertise icons. The right hand solid boundary contains the fingerprint identification domain with, at the bottom, someone like a statistician who wants to become a regular contributor to it. The statistician is bodily immersed in statistical contributory expertise, and capable of linking this contributory expertise into the domain of fingerprint identification (solid boundary) via interactional

44. And bear in mind that there will be narrower practical specialisms even *within* the mirror-coating and source-strength calculation domains. The notion of "domain" is fractal-like (Collins and Kusch 1998, chap. 1).

45. The same analysis shows us that even though the managers of big science projects make obvious and major contributions to the outcome of the science, it is still possible to distinguish among their skills between the interactional, the referred, and the contributory. For example, in Collins and Sanders 2008 (forthcoming), the authors argue that Sanders's ability to know how much weight to give to a technical argument is a referred expertise, his ability to make good decisions in respect of adaptive optics is an interactional expertise, while his ability to draw up progress charts is a contributory expertise (in management). Once more, however, there is no reason to say that the combination of all these expertises is not simply full-blown expertise in the management of science.

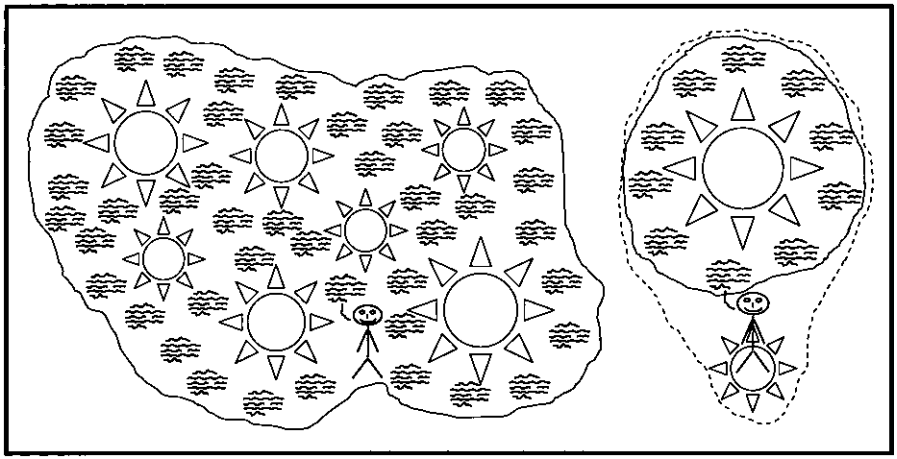

Figure 4. Contributory and interactional expertises in fields of science

expertise in fingerprinting. In the fullness of time the statistician's contributory expertise, along with the domain of statistics, will come to be seen as part of the field of fingerprinting (the ambition represented by the dotted boundary). The labels may change in such a case, and the statistician will become described as a contributory expert in fingerprinting, but his or her expertises will not have changed.

Exactly how a social scientist such as Cole should be represented on such a diagram depends on how his social science expertise makes its contribution. We can be sure that Cole can only justify any contribution he makes to the discussion of fingerprint evidence in court via his interactional expertise in the practice of fingerprinting, but whether his social science expertise might eventually become part of the domain after the fashion of the statisticians is less clear. Cole's expertise is more diffuse than that of the statisticians—it is a critique of the whole role of fingerprinting in court proceedings rather than a discussion of the correctness of any particular identification. Insofar as it has an impact on the court's decisions, it is the court that has to extract from his overall critique any contribution to the decision about guilt or innocence. Perhaps his kind of general critique is a one-off contribution which, once recognized, has no need to contribute actively to every case. In other words, while it may be that, although the critique arises from an expertise *sui generis*, it is something that is applied to different domains in turn, in the way that management consultancy is applied to different firms in turn. The social science critique is first applied to "DNA fingerprinting," then oral evidence,

then old-fashioned fingerprinting, then other aspects of forensic science, and so forth. Neither the social science critique nor management consultancy become permanent "living" features of the domains to which they are applied.

SEE as a Normally Flawed Science

We have embarked on a categorization of expertise and, no doubt, many readers will already be finding even more faults with it than we have spotted. But all categorizations of expertise will be flawed—for example, there will always be boundary problems. One reason is that, as with any other categorization, it is necessary 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.[46] 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. The point is to understand the need for a table of this kind, either this one, a modified or elaborated version of this one, or one based on a new conceptual framework.

46. For typical work by "science warriors," see *Social Studies of Science* 29, no. 2 (1999), and Dawkins 1999; Gross and Levitt 1994; Gross, Levitt, and Lewis 1996; Koertge 2000; Wolpert 1992.