S|S|S

ABSTRACT Russian measurements of the quality factor (Q) of sapphire, made 20 years ago, have only just been repeated in the West. Shortfalls in tacit knowledge have been partly responsible for this delay. The idea of 'tacit knowledge', first put forward by the physical chemist Michael Polanyi, has been studied and analysed over the last two decades. A new classification of tacit knowledge (broadly construed) is offered here and applied to the case of sapphire. The importance of personal contact between scientists is brought out and the sources of trust described. It is suggested that the reproduction of scientific findings could be aided by a small addition to the information contained in experimental reports. The analysis is done in the context of fieldwork conducted in the USA and observations of experimental work at Glasgow University.

Keywords experiment, international trust, measurement of skill, natural science, repetition of experiments, writing conventions

Tacit Knowledge, Trust and the Q of Sapphire

H.M. Collins

What is Tacit Knowledge?

Despite some years of effort, measurements of the quality factor (Q) of sapphire made 20 years ago in Russia were successfully repeated in the West only in the summer of 1999. The failure to transfer the 'tacit knowledge' of how to make the measurements has been responsible for at least some of this delay. The idea that scientists have 'tacit knowledge' was first introduced by the physical chemist Michael Polanyi.[1] Tacit knowledge has been shown to have an influence in, among other things, laser-building,[2] the development of nuclear weapons,[3] biological procedures,[4] and veterinary surgery.[5] There is also a burgeoning literature on tacit knowledge and expert systems and other 'intelligent machines', and the philosophy of tacit knowledge and the notion of practice in general.[6] Here I apply the idea to the lived world of experimental scientists. The new categorization of tacit knowledge is intended not to deepen our understanding at a philosophical level, but to explicate the idea clearly and to draw out its implications for scientific practice. My analysis is based on fieldwork, much of it conducted in the USA, among internationally-based scientists working on the detection of gravitational radiation using laser interferometry, and on observations and interviews conducted at Glasgow University in the summer of 1999.[7]

Social Studies of Science 31/1(February 2001) 71–85
© SSS and SAGE Publications (London, Thousand Oaks CA, New Delhi)
[0306-3127(200102)31:1;71–85;017581]

A New Categorization

Some scientists can do certain experiments while others cannot. We might say that this could be because the latter are bad at hand–eye coordination or related skills; it might be because the unsuccessful scientists do not have the right equipment or specimens to hand; or it might be because they lack tacit knowledge. Here we will define 'tacit knowledge' as 'knowledge or abilities that can be passed between scientists by personal contact but cannot be, or have not been, set out or passed on in formulae, diagrams, or verbal descriptions and instructions for action'.[8] Where transfer of tacit knowledge is a problem it can sometimes be solved by an exchange of visits: experimenter (B), who cannot accomplish a measurement or make a piece of apparatus work, will often succeed after spending time in the laboratory of an already accomplished experimenter (A), or after having A work for a period in B's laboratory. At least five kinds of knowledge can be passed on by such personal contact:

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 reveals these things.

Concealed Knowledge is not very interesting as a 'philosophical' category since the limitations have to do with logistics or deliberate concealment. The next four kinds of tacit knowledge apply even when A has no intention to conceal, and there is no shortage of space.

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 A and B watch each other work.
3. *Ostensive Knowledge*: Words, diagrams, or photographs cannot convey information that can be understood by direct pointing, or demonstrating, or feeling.
4. *Unrecognized Knowledge*: A performs aspects of an experiment in 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.[9]
5. *Uncognized/Uncognizable Knowledge*: Humans do things such as speak acceptably-formed phrases in their native language without knowing how they do it. Such abilities can be passed on *only* through apprenticeship and unconscious emulation. Aspects of experimental practice are similar. Uncognizable Knowledge is the most philosophically contentious case: 'reductionists' will want to say that all our abilities will one day be understood at the level of the physics and chemistry of the body and brain, so that this category (5) will collapse into category 4; others believe that abilities such

as language are irreducibly social accomplishments,[10] which means that they will never be understood at the level of brain functioning. The debate about whether some or all Uncognized Knowledge is Uncognizable need not concern us when we study the way tacit knowledge works in experimentation, for two reasons. First, the fact that language and similar human accomplishments are currently not fully understood means that now, and for the foreseeable future, even that which can be articulated in language rests on a foundation of uncognized abilities, even if they are not for ever uncognizable. Second, so long as science continues to develop, new experiments will be continually passing through a stage in which they are not fully understood, and certain aspects of the skills required to do them will be passed between experimenters only tacitly.

The category of Uncognized/Uncognizable Knowledge plays no major part in this case study, but I have included it here for clarity and completeness; one cannot define 'tacit knowledge' exhaustively without including it. It is also a vital element in debates such as that over whether computers will ever fully mimic the achievements of social beings.[11]

Toward Routine

Given the above classification, there are four ways in which procedures that were once esoteric and difficult because of their tacit component become routine.

- First, as we interact socially, that which was not obvious becomes obvious; this is what happens in the case of Concealed Knowledge, Mismatched Salience, and Ostensive Knowledge.
- Second, as we understand more science we learn to make explicit elements of our knowledge which we did not know we knew: Unrecognized Knowledge becomes recognized, and can then be passed on without personal contact.
- Third, social contact between scientists spreads knowledge that is still tacit throughout the community; that is, more scientists learn the new experimental language even though no-one can set it out. This mechanism applies to Unrecognized Knowledge so long as it remains unrecognized, and to Uncognized/Uncognizable Knowledge.
- Fourth, mechanical or 'turnkey' methods for packaging the experiment are worked out, replacing the need for tacit knowledge (which is the direction in which the case discussed below is now heading).[12]

Trust and Tacit Knowledge

The existence of tacit knowledge makes it hard to know how much time and effort will be required to copy a new piece of apparatus, or to check a measurement that has been reported elsewhere. If A's result is hard to repeat, B has to choose whether to give up on that type of work, do more experiments, try to learn more by arranging visits, or announce publicly that the original result cannot be confirmed. These options have different risks and costs. The choices sometimes arise out of the experimenters' regress,[13] but the need for them is still present in uncontested fields when no regress is apparent. Other things being equal, the more certain B is that A's result is genuine, the longer B will press on. One physicist described to me the problem as encountered in the case of the helium–neon laser:

> We regularly tried to build helium–neon lasers in the lab for staff projects. And, if you didn't know that this laser could lase, you would never believe it; it requires such patience to get it started. It makes you wonder how he [the inventor . . .] ever got it to lase because it requires so much patience to line up. Once you know it will go you can do it.

Confidence in a result may be increased or decreased as a result of familiarity with A and his or her laboratory. Thus social contact between B and A can transmit not only tacit knowledge but trust in a result, even before it has been accomplished or witnessed. I now show how this analysis applies to a current case.

The Q of Sapphire

For about 20 years, the team led by Vladimir Braginsky at Moscow State University, as part of a larger programme on low dissipation systems, has been claiming to have measured quality factors (Qs) in sapphire up to 4×10^8 at room temperature.[14] The 'quality factor' of a material indicates the rate of decay of its resonances – how long it will 'ring' if struck. (A bell that rings for a long time has a high Q, and *vice-versa*.) The number, 4×10^8 (that is, 400 million) relates to the time (in seconds) that the ringing in an object takes to die to half its original amplitude. A high Q therefore indicates a long 'ringdown' time. Reliable measures of the rate of decay (from which the ringdown time, and hence the Q, can be deduced by extrapolation) typically take a matter of minutes (see below); the ringdown times themselves can be of the order of tens of years. A long ringdown implies that the object must ring with a very pure tone – to use the jargon, it has a 'narrow resonance band'. The mirrors for the next generation of laser-interferometer gravitational wave detectors are to be made from a material with a very high Q, so that the tails of the resonance band are narrow: these tails are then less likely to spread into the frequency range of gravitational waves and become mixed up with the signals that the interferometer is designed to detect. The higher the Q – that is, the purer the tone of the mirror materials – the more sensitive will the interferometers be, because the noise levels will be lower; this is one of the most crucial features of a highly sensitive interferometer. The Russian measurements suggest that sapphire would be the best currently known material; it appears that it will soon be possible to grow sapphire crystals of sufficient size for the mirrors (about 30kg).

Because sapphire looks the most promising material, efforts have been made at other universities, including Caltech, Stanford, Perth (Western

Australia), and Glasgow (Scotland, UK) to repeat the Russian measurements.[15] But until the summer of 1999, no one outside Moscow State had succeeded in measuring a Q in sapphire higher than about 5×10^7. One American scientist told me that 'there had been a certain amount of doubt in the [Western] community because the only really high Qs that had been measured above 10^8 at room temperature had been in Moscow', while a scientist from Moscow State told me that certain Western universities had indicated that they did not trust the Russian findings.

Building Trust

Prior to the summer of 1999, what would affect Western confidence in the Russian results?

First, the result is not *a priori* improbable: it violates no scientific laws, expresses nothing radically discontinuous with what is already known, nor does it suggest improbably high levels of energy exchange; in these respects it is not like anti-gravity, water-memory, cold-fusion, or the initial claims to have seen high fluxes of gravitational radiation, all of which violated some basic or widely accepted theory or model of the contents of the universe.

Second, it is easier to get a false low reading in Q-measurement experiments than a false high reading. The measurement of Q involves setting a crystal of the material in question into vibration and watching the decay of the ringing. (In the experiments I watched, a small laser interferometer was used to monitor the decreasing movements of the end face by using the face as one of the interferometer mirrors.) There are many ways in which energy can dissipate unwantedly and unknowingly from the system, but few ways in which such a crystal can be unknowingly energized at its natural frequency so as to decay more slowly than it otherwise would. In the case of these measurements, the small crystal samples had high natural frequencies – around 40Khz (that is, 40,000 vibrations per second) – making accidental driving still less likely. That is why it is easier to damp the energy of the crystal accidentally than enhance it accidentally. False high readings might arise from faults in the laser-interferometer or other parts of the measuring system, but these are not strong possibilities (though cheating would be very easy – for example, by registering a false time-scale on the decay profile). The experiment, then, is, in this respect, more like building a successful laser than like, say, detecting telepathy, in which there are many possible sources of leakage for sensory information which could account for the results. Here, mistakes tend to produce poor results rather than positive results.

Third, however, the measurement of Q is currently very much a 'craft'. It turns on methods of suspending a crystal so that little or none of its energy of vibration will be dissipated in the suspension (see below). One scientists described the Russian experiments to me as involving a great deal of 'black magic'.

Fourth, crystals vary, and non-Russian scientists could not be sure that it was not the Russian crystals that were special rather than the Russian techniques. Apparently the Russians did nothing to clarify the situation, allowing it to be thought that they may have had special crystals developed for military purposes. (The early work on sapphire in Russia was done in connection with gyroscopes for cruise-missile guidance systems.)

Fifth, because of the Cold War, and the more recent financially impoverished state of Russian science, social ties between Russian scientists and Western scientists, and knowledge about Russian science, remain weak. Certain aspects of Russian science have long been accepted as being first class, while others – such as Lysenkoism – engender distrust; it is difficult for a non-Russian to know how to rank Russian universities and research groups. In the early days of experimental science, the social class structure of England provided a proxy for more direct sources of trust;[16] nowadays the hierarchy of universities and research groups has become a proxy for the confidence that might otherwise be inspired by social class, personal contacts, or shared membership of dense social networks. This proxy becomes less effective when the academic structure of a nation is unfamiliar.

Sixth, however, the leader of the relevant group in Moscow State University, Vladimir Braginsky, was well known in laser-interferometer gravitational wave detector circles. Kip Thorne, the Caltech theorist, had effectively been his 'sponsor' in the West for two decades, and Braginsky's quantum-level analysis of gravitational wave detectors, after initially being received with incomprehension or scepticism, has come to be an important theme in the field. On the other hand, it was widely known that at least one of Moscow State group's early experimental results – not to do with the gravitational wave field – had caused the famous American experimentalist, Robert Dicke, to disagree with Braginsky in public, and the subsequent debate has never been fully resolved to the satisfaction of the whole American community.

How the West was Won

In the summer of 1998, after a series of failed efforts to measure Qs comparable to the Russian claims, members of a Glasgow University group visited Moscow State University for a week to learn the Russian technique. Shortly thereafter, a member of the Moscow team – whom I will refer to as 'Checkhov' – worked in the Glasgow laboratory for a week. In neither case was a high-Q measurement achieved. Nevertheless, after only a few days in Russia, the Glasgow team had become convinced that the Russian results were correct. They were convinced as a result of experiencing inexactly describable features of experimental practice – the care and integrity with which the Russian experiments were done, and the apparent trustworthiness of the Russian experimenters as individuals. The new sense of trust was very robust: it stood up to continued failures to measure a high Q from summer 1998 to early 1999, even though the Glasgow team never actually witnessed any high-Q measurements.

In particular, Checkhov had left a piece of Russian sapphire with the Glasgow laboratory (after doing experiments on other crystals with them for a week), but the highest Q the Glasgow group could obtain with this specimen was around 2×10^7. And this was after attempting to match the Russian measurement over three weeks, during which they tried 20 different suspension combinations, each with a number of ring-downs at different vacuum pressures (see below for experimental details). When they finally emailed Checkhov to explain their problems, he reported that he had checked back in the Moscow laboratory notebooks and discovered that the Q of that particular piece of sapphire was not as good as he had said! Such a sequence would be taken almost to 'disprove' the existence of, say, a paranormal effect. I discussed this incident with the leader of the Glasgow team, whom I will refer to as 'Donald':

Collins: So at this point – January 1999 – you'd never seen a measurement of a high Q and you had no evidence that sapphire had this, over 10^8 Q, except from what the Russians had said. It had never been done outside Russia and you had not seen it done in Russia and you had then tried to do it on a piece of sapphire which you had been told by the Russians was capable of exhibiting 10^8 and you failed. You then got in touch with Checkhov who said 'Ah – well that bit was the wrong bit anyway'. OK – but you still did not doubt him [*Donald*: No] – because of the skills that he'd exhibited [*Donald*: Yes] because of the personal contact [*Donald*: Yes].

Donald: Occasionally you meet somebody and you just know – if you work with someone for a week, you either trust them or you don't. With Checkhov it was clear that the guy was just superb, and everything he said would turn out to be right.

Collins: Let's push this: can you really tell me how you came to this conclusion?

Donald: Well, sitting in front of this apparatus to a large extent – him looking at what we were doing and he would say 'I want to try something and modify something slightly' and you'd see improvements taking place. And he would say if you changed something you'd make it worse, and, right enough, you would change it and it would get worse. And also, you know, you hardly needed to exchange words – it was one of these things. You were thinking the same way and that is how we made such enormous progress. Because the interactions were very good with the man – you could tell how he was thinking and he could understand how you were thinking.

Collins: And there was no way this could have happened unless he'd actually been here, or you'd been there.

Donald: No – you need to have someone actually working in the lab; we were just gathered round this machine. This summer when he was across, we spent 90 hours in the lab from starting on a Sunday and finishing on the following Sunday. And he didn't want to go out and eat. He much preferred just to quickly get a sandwich and come back, and just keep going, and so we worked like that for seven days, and it is very impressive when you have a small group working like that. You get a lot done.

In the summer of 1999, Checkhov again visited the Glasgow group, bringing another piece of sapphire with him. After another week of effort, in mid-June 1999, a Q of over 10^8 was measured in the West for the first time; a similar result was achieved for a sample of American-grown sapphire. Subsequently, the measurements have been repeated with no Russian present by a member of the Glasgow team (who was present during Checkhov's visits to Glasgow), working in Stanford on an American-grown sample.

Components of Tacit Knowledge in Q-Measurement

The method of measuring Q is to suspend the crystal – which might be a cylinder 5–10 cms long and 1–10 cms in diameter – in a sling about its mid point. The sling is a single thread or wire which is looped around the crystal, the ends being held by compressing them in a clamp above the crystal. The crystal is thus balanced at the end of a pendulum which helps isolate it from vibrations transmitted from the apparatus. The suspended crystal is loaded into a vacuum chamber which is pumped down. One end of the crystal is painted with a dot of aluminium so that it acts as a mirror for the laser interferometer used to measure the vibrations, through a porthole. The crystal is driven up (set ringing) by an electrostatic end-plate generating an AC field at the crystal's natural frequency. The field is switched off and the decay of the vibration, measured by the interferometer system (which can compensate for gross movements of the whole crystal), can be seen on a chart recorder or fed directly to a computer for analysis. The rate of decay can be converted into the Q of the sapphire. For a high-Q crystal it might take 20 minutes or so to register sufficient decay to provide a good measurement. A lower-Q crystal requires only a minute or so to give an easily measurable result.

Cylindrical resonators have no perfect modes,[17] so even if one of these sapphire crystals is suspended exactly around its mid-point some movement will be transmitted to the suspension fibres; therefore it is effectively the Q of the crystal/pendulum system that is being measured. A false low reading will result from losses of energy in the system. Significant energy can be transferred from crystal to suspension if the pendulum's natural frequency of vibration is similar to that of the crystal – that is, to make the system work properly, the pendulum length must be 'anti-matched' to the crystal frequency. Friction losses between fibres and clamp must be avoided by making the clamp contact the fibres sharply where they first enter the clamp area – but not so sharply that the fibres are severed. Energy can also be lost in friction between crystal and fibre, and there are potential friction losses within the fibre itself – thus the choice of fibre and the preparation of the fibre are both important. There are also thermodynamic losses between the vibrating elements and the residual air in the vacuum chamber. The art of the experiment is to minimize all these losses.

By watching Checkhov work, the Glasgow group learned that good measurements had to be accomplished by trial and error over many repeated runs – they learned that the experiment remained difficult even *after* a first success had been achieved. As Donald put it:

> I think the thing that we learned most of all was patience. [We] would experiment away for a morning, perhaps, and after several runs we would end up with the same Q; in the past we would have been tempted to say that *was* the Q. What we learned from Checkhov was that he was much more patient than that. He would go for days before he would believe [such a result]. He would keep varying the parameters by tiny amounts, because he knew to do that from the work he had done previously. And there would be enormous time put into it. And we would be sitting watching . . . And once you know to do that [you can succeed] – but until you know that, it's hard.

Checkhov's approach, however, also revealed ways in which each of the many runs could be done more efficiently. To pump a vacuum chamber down to a very low pressure takes a long time. The Glasgow group had been pumping down for about two and a half hours prior to each measurement, while Checkhov's practice cut the time in half, sacrificing an order of magnitude or two in vacuum. Checkhov's practice showed that most of what needed to be learned could be learned at a higher pressure, reserving the lowest pressure runs for a final measurement only. Checkhov also used very short suspensions. The Glasgow group had used suspensions comparable in length to those that would eventually be employed in full-scale laser interferometers, but Checkhov used as short a length as possible, so as to make frequency matching with the crystal less likely (the nodal frequencies are further apart in short strings). Thus, with Checkhov's approach, fewer set-ups were wasted and less time and care had to be spent on getting the length of the pendulum right so as to make sure it was anti-matched to the crystal.

Social science is untidy compared to a controllable laboratory science, but we will try to describe what was going on in terms of the five-fold classification of tacit knowledge. In this case there does not seem to have been any category 5 (Uncognized/able Knowledge) transferred between Moscow and Glasgow, because both groups already shared the same broad 'language of science'. Differences in this kind of knowledge show themselves only where very big differences in scientific world view are juxtaposed.

Knowledge about the degree of vacuum and the length of the suspension belong to categories 1 and/or 2 (Hidden Knowledge/Mismatched Saliences). This is because the degree of vacuum in exploratory runs is not likely to be noted in a published report; likewise, gross pendulum length seems like a choice that would be made on grounds other than experimental efficiency. Yet, with trial and error, efficiency is very important if enough runs are to be carried out to press the measurements to the limit. Certainly, the most appropriate choices became clear to the Glasgow group only through watching the Moscow practices.[18] Though the importance of the clamping could be described, and has been described, it was Checkhov's way of working that revealed the possible importance of repeated minute adjustments to the clamp, should high-Q not be achieved. To describe the principle of clamping, and to mention its importance, is not the same as revealing its importance through the care that is taken in practice; we do not have an exact language for describing 'degree of care that needs to be taken', so coming to understand it is a matter of Ostensive Knowledge – category 3.

Something similar applies to the material of the suspension fibres. Checkhov used very fine Chinese silk thread which he supplied to the Glasgow group (who had earlier used steel piano wire). Trial and error had shown the Russians that other kinds of silk thread gave lower Qs. It was also known that fine tungsten wire gave still better results, but that it had to be polished carefully to just the right (indescribable) degree, and that the clamping problem was particularly acute with tungsten. Donald believed that it was the hardness of the tungsten that made the clamping so critical – the compressibility of silk allowed a certain leeway in the design of the clamp. Thus silk was used for most runs, with tungsten (which might improve the Q by a factor of 2) being preserved for a final measurement once the general area of the expected result had been defined by the easier method. The nature of suspension materials and clamping seem to belong in categories 2 (Mismatched Salience), 3 (Ostensive Knowledge), and 4 (Unrecognized Knowledge): they are matters whose salience became clear for the Glasgow group only after working with Checkhov. For both parties the science was slowly emerging and turning knowledge that no one knew they could or should express, into something that could be articulated as the importance of previously unnoticed parts of the procedure became revealed.[19]

Polishing of tungsten (as described above), and greasing of both tungsten and silk, had been found to be vital. In Braginsky, Mitrofanov and Panov's book, we find this claim: 'The presence of a fatty film (e.g., pork fat) at the points of contact between the suspension fiber and the resonator is important'.[20] It was believed that grease between fibre and crystal prevented frictional losses. Greasing turned out to be critical, but there is no vocabulary to describe exact amounts of pork fat (after watching Checkhov, the Glasgow group used commercially available lard, whereas they had previously used 'apiezon' grease).

Working with Checkhov revealed two methods of greasing a fine silk thread. A thicker Italian silk thread was first greased with a 'daud' (a Scots dialect word) of lard and wiped with a cloth until most of the lard had been absorbed or rubbed off. The crystal was then mounted and balanced in a loop of this thread. The greased Italian thread would leave a thin track on the crystal. The crystal was then dismounted and re-hung on fine, Russian-supplied, Chinese thread, which would now be sitting in the thin ring of grease left by the thicker Italian thread. The run I witnessed produced a slightly lower Q than had been expected, and the reasons described to me indicate a nice case of Ostensive Knowledge:

Ericson: It's very difficult to be precise about the amount of grease you apply because you're just applying grease to the thread. If you apply too much the Q tends to fall off because it's too loose and it will wobble and you will get an erratic ringdown. But if you have too little grease then the thread may stick and slip rather than sit smoothly on the mass. In this case I think there probably wasn't quite enough grease, which is why it [the Q] is slightly lower than what I thought it might be. But if you get it spot on you can usually get a very high result. . . . I think there's not quite enough.

Collins: And that's just from your looking at it.

Ericson: Yeh – that's just empirical – from my experience of doing this before, I can sort of tell. When you take off the greased thread and you see this band of grease, there's a feel for what's enough and what's too much. And that looked less – but not too far off.

The second method of greasing thread demonstrated by Checkhov, and used interchangeably with the first method, was direct greasing of the fine thread with human body grease. Checkhov would run the fine Chinese thread briefly across the bridge of his nose or behind his ear. The ear method was adopted by the Glasgow group, though it turned out that only some people had the right kind of skin. Some, it transpired, had very effective and reliable grease, others' grease worked only sporadically, and some experimenters' skins were too dry to work at all. All this was discovered by trial and error, and made for unusual laboratory notebook entries such as: 'Suspension 3: Fred-greased Russian thread; Suspension 12: switched from George-grease back to Fred-grease', and so forth. As with James Joule's famous measurement of the mechanical equivalent of heat,[21] it seems that the experimenter's body could be a crucial variable. Knowledge of how to apply the right amount of grease to the system has aspects that belong in categories 2, 3 and 4.

Conclusion and Recommendation

A difficult measurement can be repeated by inventing a new method or reinventing the old one. In the case of quality measurements of crystals, it seems that one American group managed to measure high Qs in glass by an entirely different method, and in July 1999 an Australian group briefly mentioned an independent replication of the Russian results using a tungsten wire support. However, if B is to repeat a difficult measurement using A's (the original experimenter's) method, three kinds of things have to be established: B needs to master A's explicit and tacit knowledge; B needs to be certain that the result really has been achieved by A; and B needs to know how difficult the procedure is, as this indicates how long it will be necessary to persevere to have even a chance of repeating the result.

On the learning of explicit and tacit knowledge, there is little to add except to re-emphasize the importance of laboratory visits, and to hope that recognizing and understanding the importance of tacit knowledge might ease its transfer – especially to new recruits to science who have not experienced the problems for themselves.

Being certain that a result has been achieved is a matter of trust. Replication of results leads to trust, but the case also illustrates the opposite point: it was only because the results emerging from Moscow State were trusted – for the reasons given in the section on 'Trust and Tacit Knowledge' – that Western laboratories thought it worthwhile to continue after a long period of failure. The still greater trust engendered by the exchanges of visits between the Glasgow and Moscow State groups led the Glasgow team to redouble their efforts. Thus, though successful repetition of a result leads to trust, more importantly for the confirmation and spread of new techniques, *trust leads to successful repetition.*

Knowing how difficult a skill is, is another important part of learning to master it.[22] If one believed that bike-riding could be mastered in one minute, a few minutes of falling off would lead one to distrust claims that bikes could be ridden at all, and one would never learn to ride – still more so with, say, playing a musical instrument. One important thing that the Glasgow group learned from Checkhov was what they called 'patience' which, in these terms, is a matter of learning that measuring Q is difficult and remains difficult (like, for example, golf, rather than bike-riding), even after one has first accomplished it.

Reporting a Second Order Measure of Skill

This kind of science could be made easier if the importance of knowing the difficulty of an experimental skill or procedure was recognized and emphasized. The conventional style of writing scientific journal papers (and even books) excludes details of this kind. Yet someone trying to rediscover how to produce a result in the absence of a laboratory visit could be helped by knowing just how hard the experiment or measurement was to carry out in the first place, and just how hard it continues to be. Such information could be roughly quantified – it is a 'second order measure of skill'.[23] Experimenters could record something along these lines:

> It took us some 17 months to accomplish this result in the first instance, during which time we tried around 165 runs with different set-ups, each run taking around a day to complete. Most successful measurements on new samples are now obtained in around 7 runs, but there is a range of approximately 1 to 13 runs; each run now takes about 2 hours. The distribution of numbers of runs on the last 10 samples we have measured is shown in the following diagram . . .

Information of this sort could be expressed briefly, without radically changing the conventional style of scientific paper-writing, and yet could be of significant benefit to those trying to repeat the work. It is just a matter of admitting that most things that seem easy now were very hard to do first time round, and that some remain hard even for the experienced experimenter. We concede, of course, that within the current conventions of scientific writing, setting out these difficulties would look like weakness;

science is conventionally described as though it were effortless, and the accepted scientific demeanour reinforces this impression. What we are suggesting is a slight transformation of convention and demeanour – with a view to improving the transmission of scientific knowledge.

Notes

This research was supported by UK ESRC Grant No. R000236826. The quotations in the text are from interviews carried out in that research. The names of those quoted have been anonymized, though their identities will be obvious to researchers in the field. At the request of respondents, quotations have been lightly edited for style, and the amended quotations have been checked by the speakers. I thank members of the Glasgow group for their time and patience and supplying most of my material, members of the Moscow State group for introducing me to the sapphire problem and reading and commenting on the paper, and Peter Saulson for helpful comments on the analysis, as well as on the physics.

1. The classic text is Michael Polanyi, *Personal Knowledge* (London: Routledge & Kegan Paul, 1958), esp. Part Two, 'The Tacit Dimension', 69–245.
2. H.M. Collins, 'The TEA Set: Tacit Knowledge and Scientific Networks', *Science Studies*, Vol. 4, No. 2 (April 1974), 165–86; Collins, *Changing Order: Replication and Induction in Scientific Practice* (Chicago, IL: The University of Chicago Press, 2nd edn, 1992), Chapter 3, 51–78.
3. Donald MacKenzie and Graham Spinardi, 'Tacit Knowledge, Weapons Design and the Uninvention of Nuclear Weapons', *American Journal of Sociology*, Vol. 101, No. 1 (July 1995), 44–99.
4. Kathleen Jordan and Michael Lynch, 'The Sociology of a Genetic Engineering Technique: Ritual and Rationality in the Performance of the "Plasmid Prep"', in Adele Clarke and Joan Fujimura (eds), *The Right Tools for the Job: At Work in Twentieth-Century Life Sciences* (Princeton, NJ: Princeton University Press, 1992), 77–114; Alberto Cambrosio and Peter Keating, '"Going Monoclonal": Art, Science, and Magic in the Day-to-Day Use of Hybridoma Technology', *Social Problems*, Vol. 35, No. 3 (June 1988), 244–60. Oddly, Jordan and Lynch do not use the term 'tacit knowledge' in their discussion of continuing variation of craft practices in biological preparations, nor do they refer to the previously existing and well-known literature on the same topics as explored in the physical sciences. Cambrosio and Keating stress the way that scientists themselves use similar categories to describe the 'artistic' and 'magical' aspects of their work as those developed here. They claim that scientists' own categories are sharper and more useful than those developed by sociologists. Certainly one criterion of success in the description of tacit knowledge is whether the outcome seems plausible to scientists – that is, whether it matches their world as they already understand it. In this paper I try to provide a more systematic exploration of that world than is common in the physical sciences, and then to draw out some implications which certainly vary from scientists' current practice. Should all this be a failure in the eyes of scientists, I believe that it is still a worthwhile exercise, for the object of much sociology of scientific knowledge is to explore the world of scientific knowledge for the sake of non-scientists.
5. Trevor Pinch, H.M. Collins and Larry Carbone, 'Inside Knowledge: Second Order Measures of Skill', *Sociological Review*, Vol. 44, No. 2 (May 1996), 163–86.
6. For the idea applied to expert systems and machines in general, see H.M. Collins, *Artificial Experts: Social Knowledge and Intelligent Machines* (Cambridge, MA: MIT Press, 1990); Collins and Martin Kusch, *The Shape of Actions: What Humans and Machines Can Do* (Cambridge, MA: MIT Press, 1998); and, for example, the articles in Bo Goranzon and Ingela Josefson (eds), *Knowledge, Skill and Artificial Intelligence* (London: Springer-Verlag, 1988). For a critique of the idea, see Stephen Turner, *The Social Theory of Practices: Tradition, Tacit Knowledge and Presuppositions* (Oxford: Polity Press, 1994). For a discussion of this critique and an analysis of the way in which the idea of tacit knowledge has been used in sociology of science, see H.M. Collins, 'What

is Tacit Knowledge?', in Theodore R. Schatzki, Karin Knorr-Cetina and Eike von Savigny (eds), *The Practice Turn in Contemporary Theory* (London & New York: Routledge, 2001), 107–19; and for more of this, and a philosophical treatment of the notion of practice, see the other articles in this book.

7. See *www.caltech.ligo.edu* for details of the US Laser Interferometer Gravitational-Wave Observatory (LIGO) and links to other international projects, and *www.cardiff.ac.uk/socsi/gravwave* for the author's work on the sociology of gravitational wave detection.
8. A relatively 'low' epistemological level for the analysis is selected as it is intended that this paper make contact with the world of scientists as well as the world of epistemologists. Thus, there is no discussion here of the nature of knowledge in the natural sciences. For example, we do not discuss whether, say, the efficacy of the skills involved in the preparation of a ritual sacrifice to appease an angry god is more or less readily verifiable than that of the skills exhibited in the scientific laboratory. Harry Collins and Steven Yearley have argued the need for 'meta-alternation' – that is 'alternation', as discussed by Peter L. Berger in his *Invitation to Sociology* (Garden City, NY: Anchor Books, 1963), esp. 65–66, 77–78 – between epistemological levels: see H.M. Collins and S. Yearley, 'Epistemological Chicken', in Andrew Pickering (ed.), *Science as Practice and Culture* (Chicago, IL: The University of Chicago Press, 1992), 301–26. The approach of meta-alternation should be contrasted with the more epistemologically uniform but, as Collins and Yearley argue, less methodologically powerful, approach of Bruno Latour and Michel Callon and their followers: see M. Callon and B. Latour, 'Don't Throw the Baby Out With the Bath School!', in Pickering (ed.), op. cit., 343–68.
9. For a discussion of this problem with reference to the construction of Italian violins, see Colin Gough, 'Science and the Stradivarius', *Physics World*, Vol. 13, No. 4 (April 2000), 27–33.
10. The classic text here is Ludwig Wittgenstein, *Philosophical Investigations* (Oxford: Blackwell, 1953).
11. See, for example, Collins & Kusch, op. cit. note 6. For a different kind of discussion of Uncognizable Knowledge, see Cambrosio & Keating, op. cit. note 4.
12. For an extended discussion of the transformation of tacit knowledge into 'turnkey' knowledge, see Collins (2001), op. cit. note 6, passim, and Collins & Kusch, op. cit. note 6, Chapter 9.
13. Collins (1992), op. cit. note 2, 79–106. The 'experimenters' regress' occurs when a series of experimental replications is used to *test* a disputed claim. The usual criterion of successful execution of an experimental skill – an outcome in the right range – is absent, because the nature of the right outcome is exactly what is under dispute. Therefore experimenters can argue indefinitely over the question of which set of experiments with conflicting outcomes were performed adequately. The answer to this question provides the answer to the question about the correct outcome to such experiments. But the only way to decide which experiments have been performed adequately is to decide what the correct outcome is, and then to see which experiments produce it. Hence the regress.
14. Vladimir Borisovich Braginsky, V.P. Mitrofanov and V.I. Panov (trans. Erast Gliner, ed. and intro. Kip S. Thorne and Cynthia Eller), *Systems with Small Dissipation* (Chicago, IL: The University of Chicago Press, 1985).
15. For the sake of completeness, I should add that the Russians more recently discovered that sapphire has some less desirable properties that may make it less suitable as mirror material. Among the community this revelation is known as 'the Braginsky bombshell', but it does not affect the argument presented here.
16. For an extended exploration of this point, see Steven Shapin, *A Social History of Truth: Civility and Science in Seventeenth-Century England* (Chicago, IL: The University of Chicago Press, 1994).
17. Braginsky, Mitrofanov & Panov, op. cit. note 14, 25–26.
18. A diagram showing a crystal supported by a very short pendulum is shown on page 27 of Braginsky, Mitrofanov & Panov, op. cit. note 14, but it could easily be read as a

schematic representation rather than something to be interpreted literally. An impression of a detail of that diagram is reproduced here (Figure 1).

FIGURE 1
Apparently Short Suspension for Cylindrical Resonator

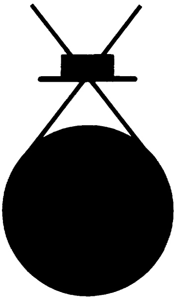

Source: adapted from Braginsky, Mitrofanov & Panov, op. cit. note 14, 27.

19. In some respects, the process could be compared with the experience of an anthropologist. Native members cannot describe the 'taken-for-granted' aspects of their world precisely because they *are* taken-for-granted, but these can become salient when they are viewed through the eyes of someone to whom they are not familiar.
20. Braginsky, Mitrofanov & Panov, op. cit. note 14, 29.
21. Heinz Otto Sibum, 'Reworking the Mechanical Value of Heat: Instruments of Precision and Gestures of Accuracy in Early Victorian England', *Studies in History and Philosophy of Science*, Vol. 26, No. 1 (March 1995), 73–106.
22. This point is argued by Pinch, Collins & Carbone, op. cit. note 5.
23. Ibid.

Harry Collins is Distinguished Research Professor of Sociology and Director of the Centre for the Study of Knowledge, Expertise and Science (KES) at Cardiff University. His published books include: *Changing Order: Replication and Induction in Scientific Practice* (1985; 2nd edn, Chicago UP, 1992); *Artificial Experts: Social Knowledge and Intelligent Machines* (MIT Press, 1990); with Trevor Pinch, *The Golem: What You Should Know About Science* (Cambridge UP, 1993; 2nd Canto Edn, 1998), which won the 1995 Robert Merton Prize of the American Sociological Association; also with Trevor Pinch, *The Golem at Large: What You Should Know About Technology* (Cambridge UP, 1998); and, with Martin Kusch, *The Shape of Actions: What Humans and Machines Can Do* (MIT Press, 1998). He won the 1997 J.D. Bernal Award of the Society for Social Studies of Science (4S), and is currently working on a longitudinal study of the history and sociology of gravitational wave physics.

Address: Centre for the Study of Knowledge, Expertise and Science (KES), Cardiff School of Social Sciences, Cardiff University, The Glamorgan Building, King Edward VII Avenue, Cardiff CF10 3WT, UK; fax: +44 2920 874175; email: CollinsHM@Cardiff.ac.uk