

School of Social Sciences



Working Papers Series

Paper 1

2000

TACIT KNOWLEDGE, TRUST, AND THE Q OF SAPPHIRE

HARRY COLLINS

Centre for the Study of Knowledge, Expertise and Science (KES), School of Social Sciences, Cardiff University, Cardiff, UK

Russian measurements of the quality factor (Q) of sapphire, made twenty years ago have only just been repeated in the West. Shortfalls in tacit knowledge have been partly responsible. 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 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 findings could be aided by a small addition to the information contained in experimental reports.

1. WHAT IS TACIT KNOWLEDGE?

Measurements on the quality factor of sapphire made twenty years ago in Russia were repeated in the West only this summer. The failure to transfer the 'tacit knowledge' of how to make the measurements has been responsible for at least some of the 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 laser-building and the development of nuclear weapons,^{2,3} and the idea has been refined by considering its importance for expert systems and other 'intelligent machines.'^{4,5}

Some scientists can do certain experiments while others cannot. This might be because the latter are bad at hand-eye co-ordination 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 defined as knowledge or abilities that can be passed between scientists by personal contact but cannot be, or has not been, set out or passed on in formulae, diagrams, or

verbal descriptions and instructions for action. 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 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 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 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 realise 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 and watch each other work.
3. Ostensive Knowledge: Words, diagrams, or photographs cannot convey information that can be understood by direct showing or pointing.

4. Unrecognised Knowledge: A performs aspects of an experiment a certain way without realising their importance; B will pick up the same habit during a visit while neither party realises that anything important has been passed on. Much Unrecognised Knowledge becomes recognised and explained as a field of science becomes better understood but this is not necessary. (The importance of the characteristics of the varnish used by Stradivarius and passed on to his apprentices has only just been recognised!)
5. 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.^{4,5} 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 category 5 will collapse into 4; others believe that abilities such as language are irreducibly social accomplishments⁶ which means they will never be understood at the level of brain functioning. But the philosophical argument is irrelevant to the understanding of 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.

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 that of which we were once unaware: Unrecognised Knowledge becomes recognised and can 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 know its 'grammar' in a conscious way. This mechanism applies to Unrecognised Knowledge so long as it remains unrecognised and to Uncognized/able 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 headed).

2. TRUST AND TACIT KNOWLEDGE

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 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; in highly contested fields, such as the

early days of gravitational wave detection, a dilemma sometimes referred to as the `experimenter's regress is apparent,'⁷ but even in less contested fields the choices are difficult. Other things being equal, the more certain is B that A's result is genuine, the longer will B press on. A physicist described 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.

3. THE Q OF SAPPHIRE

For about 20 years the team led by Vladimir Braginsky at Moscow State University, as part of larger program on low dissipation systems, have been claiming to have measured quality factors (Qs) in sapphire up to 4×10^8 at room temperature.⁸ The quality factor of a material indicates the rate of decay of its resonances and relates to the width of the resonance band; a high Q indicates a slow decay and a narrow band. The mirrors for the next generation of laser-interferometer gravitational wave detectors are to be made of a material with a

very high Q so that the resonance band is narrow and the tails of the resonance are less likely to overlap the frequencies of potential gravitational waves. The Russian measurements suggests 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 (c 30kg).

Because sapphire looks the most promising material efforts have been made at universities including Caltech, Stanford, Perth, and Glasgow to repeat the Russian measurements. But until the summer of 1999, no one outside Moscow State had succeeded in measuring a Q higher than about 5×10^7 in sapphire. 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 implied that they did not trust the Russian findings.

4. 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 improbable 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.

Second, it is easier to get a false low reading in Q-measurement experiments than a false high reading. The measurement of Q involves energising a crystal and watching its vibrations decay using a laser

interferometer to monitor movements of the end face of the crystal. 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 driven 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 - making accidental driving still less likely. False positives might arise from faults in the laser-interferometer or other parts of the measuring system but they are not strong possibilities (though cheating would be very easy - eg by registering a false time-scale on the decay profile). In this, the experiment is more like building a successful laser than like, say, paranormal experiments, 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 none of its energy of vibration will be dissipated in the suspension (see below). One scientist 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, offering the suggestion 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 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. The social class structure of England provided a proxy for more direct sources of trust in the early days of experimental science;⁹ 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.

Sixth, however, the leader of the relevant group in Moscow State University, Vladimir Braginsky, was well known in laser-interferometer gravitational wave detector circles. Caltech theorist, Kip Thorne, 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 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.

5. 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 the Glasgow University group visited Moscow State University for a week to learn the Russian technique. Shortly thereafter, a member of the Moscow team - who I will refer to as 'Chekhov,' - worked in the Glasgow laboratory for a week. In neither case was a high-Q measurement

achieved. Nevertheless, the Glasgow team had become convinced that the Russian results were correct after only a few days in Russia. 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 trustworthiness of the Russian experimenters as individuals. The new sense of trust was very robust and it stood up to continued failures to measure a high-Q from summer 1998 to early 1999.

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 they could obtain with it was around 2×10^7 . And this was after attempting to match the Russian measurement over three weeks during which they tried twenty different suspension combinations each with a number of ring-downs at different vacuum pressures. When they finally emailed Checkhov to explain their problems he said 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, who 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 now 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. At the time of writing (10 September 1999), it is reported that the measurements have just been repeated with no Russian present by a member of the Glasgow team working in Stanford on an American-grown sample; this group member was present during Checkhov's visits to Glasgow.

6. 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 wraps round 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 which shines through a porthole. The crystal is driven up 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 face), can be seen on a chart recorder or fed directly for analysis in a computer. 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.

Sapphire crystals have no perfect modes,⁸ so even if they are suspended exactly around the 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 unless pendulum length is 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 minimise 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 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. The Glasgow group had been pumping down for about 2.5hrs 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 be employed in full-scale laser interferometers, but Checkhov used as short a length as possible so as to make frequency matching 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.

Social science is untidy compared to a controllable laboratory science but we will try to describe what was going on in terms of the fivefold classification of tacit knowledge. In this case there does not seem to have been any category 5 (Uncognizable Knowledge) transferred between Moscow and Glasgow because both groups already shared the same broad 'language of science.' Differences in Uncognizable 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 Salience). This is because 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. (A diagram showing a crystal supported by a very short pendulum is shown on page 27 of ref 8 but it could easily be read as a schematic representation rather than something to be interpreted literally.)

Though the importance of the clamping could be described, and has been described, it was Chekhov'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 Q's. 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 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 (Unrecognised Knowledge): they are matters whose salience became clear for the Glasgow group only after working with Checkhov, while for both parties the science was slowly emerging and turning inexpressible knowledge into something that could be articulated and revealing the importance of previously unnoticed parts of the procedure.

Polishing of tungsten (as described above), and greasing of both tungsten and silk had been found to be vital. In ref 8 it says (p29) 'The presence of a fatty film (e.g., pork fat) at the points of contact between the suspension fiber and the resonator is important.' 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 (the Glasgow group used commercially available lard after watching Checkhov, whereas they had used 'apiezon' grease previously).

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 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 thread, which would now be sitting in the thin ring of grease left by the thicker thread. The run I witnessed produced a slightly lower Q than had been expected and the reasons described to me indicates 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 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 turned out, 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,¹⁰ it turns out 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.

7. 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 one American group managed to measure high Qs in glass by a different method, and in July 1999 an Australian group briefly mentioned an independent replication of the Russian results using a tungsten wire support. In the normal way, however, to repeat a difficult measurement three kinds of things have to be known. 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-emphasise the importance of laboratory visits and to hope that recognising and understanding 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 section 2, above - 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 replication leads to trust, more importantly for the confirmation and spread of new techniques, trust leads to replication.

Knowing how difficult a skill is, is another important part of learning to master it. 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 (eg, like golf, rather than bike-riding), even after one has first accomplished it.

This kind of science could be made easier if the importance of knowing the difficulty of an experimental skill or procedure was recognised and emphasised. 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.¹¹ Experimenters could record something along the 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 hrs. The distribution of numbers of runs on the last 10 samples we have measured is shown on 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.

¹ Polanyi, M. Personal Knowledge. (Routledge and Kegan Paul, London, 1958).

² Collins, H. M. The TEA Set: Tacit Knowledge and Scientific Networks. Science Studies, **4**, 165-186 (1974).

³ MacKenzie, D., & Spinardi, G. Tacit Knowledge, Weapons Design and the Uninvention of Nuclear Weapons.. American Journal of Sociology, **101**, 44-99 (1995).

⁴ Collins, H. M. Artificial Experts: Social Knowledge and Intelligent Machines. (MIT press, Cambridge, Mass, 1990).

⁵ Collins, H. M., & Kusch, M. The Shape of Actions: What Humans and Machines Can Do. (MIT Press, Cambridge, Mass, 1998).

⁶ Wittgenstein, L. Philosophical Investigations. (Blackwell, Oxford, 1953).

⁷ Collins, H. M. Changing Order: Replication and Induction in Scientific Practice. (2nd Edition, University of Chicago Press, Chicago, 1992).

⁸ Braginsky, V. B., Checkhov, V. P., & Panov, V. I. Systems with Small Dissipation. (University of Chicago Press, Chicago, 1985). [Translated by Erast Gliner With an Introduction by Kip S. Thorne].

⁹ Shapin, S. A Social History of Truth: Civility and Science in Seventeenth-Century England. (University of Chicago Press, Chicago, Ill, 1994).

¹⁰ Sibum, H. O. 'Reworking the mechanical value of heat: instruments of precision and gestures of accuracy in early Victorian England. Studies in History and Philosophy of Science, **26**, 1, 73-106 (1995)

¹¹ Pinch, T., Collins, H. M., & Carbone, L. Inside Knowledge: Second Order Measures of Skill. Sociological Review, **44**, 2, 163-86 (1996).

ACKNOWLEDGEMENTS

This research was supported by UK, ESRC grant No. R000236826. The names of those quoted have been anonymised 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 the physics.