The Myth and Reality of Scientific Consensus: a Physicist's Perspective

Paolo G. Radaelli, Dr Lee's Professor of Experimental Philosophy
Clarendon Laboratory, University of Oxford, OX1 3PU, United Kingdom

Drawing mainly from examples in solid-state physics, I argue that consensus in healthy, non-pathological science can evolve very rapidly, following neither a traditional 'knowledge accumulation' model nor a social constructivist model. Rather, it is found that revolutions can be driven by individuals who were previously rather marginal within the scientific community, and may happen to stumble upon radically new facts, sometimes without the guidance of correct intuition. Although scientific consensus seems to be robust against fraud and genuine mistakes, the unpredictability of this pattern of evolution poses a challenge for rational governance. I propose a minimal approach to establish what is scientifically consensual and how to incorporate it into policy, based on the concept of 'actionable evidence'.

I. INTRODUCTION

The application of scientific principles is an essential tool of political and economic governance, at least if one wants to govern rationally. Although the existence of scientific consensus does not, in itself, dictate policy decisions, establishing whether there *is* a scientific consensus and, if so, what it is, are the first steps towards rational governance. Yet, the very authority of science to speak objectively about *any* issue is currently being challenged. On one hand, the association between science and technology means that the man-made disasters produced by the abuse of the latter (such as the Chernobyl, Bhopal and Fukushima incidents, to mention just a few) taint the reputation of the former; at a more philosophical level, the Enlightenment idea that science represents Nature, rather than the specific culture of those who produce it (typically white European males) is called into question [1]. In this brief contribution, I set myself three challenges:

1. To illustrate how scientific consensus is reached in 'normal science', and to highlight some of the risk factors that can undermine the integrity of the scientific process.
2. To discuss examples of how scientific consensus changes over time, often very rapidly, and to explain what happens when consensus collapses.
3. Zooming out to the more general issue of science in governance and society, to discuss a working approach to establishing the strength of consensus on a given issue at any particular time, and to link it to policy-making through the concept of 'actionable evidence'.

In developing these themes, I will draw on my 25-year experience in solid-state physics and on my position as a senior academic at a global and prestigious university like Oxford. The attentive reader cannot fail to notice that, in doing so, I am establishing myself in the very position of authority that I have called into question in the previous paragraph. I humbly ask the reader to suspend disbelief for now, whilst proposing to return to this issue in the final part of this contribution. I wish to add a few words about my subject and its status, representative or otherwise, in the more general landscape of the sciences. Physics contends to chemistry the title of the oldest of the Natural Sciences, and is *the* mature scientific subject *par excellence*, in which the scientific method is expressed at its very pinnacle. More than any other branch of physics, solid-state physics has transformed the world by spawning a multitude of technologies, from the transistor to the LED. Sources of many ethical dilemmas in the past, physics in general and solid-state physics in particular seem today less ethically contentious than, say, genetic engineering. For all these reasons, my subject is arguably not entirely representative of the whole spectrum of the sciences, especially if we wish to include the human sciences. However, some aspects are indubitably common at least to all the natural sciences, and I hope you will appreciate the 'cold' analysis method of an experimental physicist, combined with the reflective and narrative approach that, I hope, I can personally contribute.

II. HOW IS CONSENSUS REACHED IN 'NORMAL SCIENCE'?

In his landmark 1962 book 'The structure of scientific revolutions' [2], Thomas Kuhn contrasts the 'paradigm shifts' that characterise scientific revolutions with much longer phases of what he calls 'normal science' in which science operates within established paradigms. Yet, even in normal science, consensus is constantly challenged and sometimes undergoes sudden collapses and reconstructions. How is this done? Normal science follows processes that are largely standardised across very different fields and different parts of the globe, so as to appear culturally non-specific. One learns about these processes as a students and takes them for granted thereafter, often forgetting that most people outside science have very little idea of how they work. ‘My’ normal science deals with the properties of certain magnetic oxides that may be useful for future IT applications (I remain deliberately vague here), and is largely created in the lab (either in my institution or at large-scale laboratory infrastructures) and in discussions with my students and junior collaborators. We normally start off with a working hypothesis on how a certain material would behave, often exploiting the power of generalisations and analogies, sometimes with some very distant field of knowledge. The hypothesis leads to the design of an experiment, which is put together, in some cases, with the help of collaborators in many different institutions across the world. The results are analysed, often leading to surprising discoveries, and this information is used to propose new hypotheses and new experiments. The culmination of this process is a ‘paper’, i.e., a scientific essay directed at a specialised audience and intended to be published in a leading journal. This is fundamentally different from other academic subjects, especially in the humanities, which disseminate their results and ideas mainly through books [3]. Nevertheless, it is important to emphasise the central role that creative writing plays in drafting a scientific paper: I never tire to tell my students that the core of a paper is always a *story*, and that the success of a paper depends on how interesting the story is and, particularly, on how tightly one can grip the reader’s attention in the first few lines. Most papers, these days, report original research; papers reviewing the state of the field or summarising the contribution to a conference used to be much more common, but they are increasingly regarded as a waste of time by most. Once sent to a journal editor with the all-important ‘cover letter’ explaining why it is absolutely essential for the journal to publish that piece of work (...), the paper is usually forwarded to anonymous referees (i.e., other scientists in the same field — the so-called ‘Peer Reviewers’), with very little editorial pre-filtering. If the paper is not rejected outright, one or more rounds of revisions usually follows, and the whole process is completed, typically in a few months, resulting in the actual publication in print and on the journal web site. Except for the very last part, the process has remained basically the same for the last three generations [4]. It is also very common today, at it was then, to discuss unpublished results at conferences and workshops. One aspect has changed dramatically, however, and concerns the dissemination of manuscripts prior to publication — the so-called *preprints*. Since time immemorial, *preprints* were produced, usually after the last round of refereeing and before the actual publication in print, and were either posted to colleagues or handed in at conferences. The function they served at that time, that is to avoid the lengthy delays associated with the printing and shipping of the journal, is now entirely obsolete, because papers usually appear on the journal web site very soon after acceptance. At least in some communities, *preprints* are now usually published in specialised on-line archives at the time of *submission*, and their function is to establish priority on a particular discovery or important result (“We don’t want to be scooped by X whilst the editors are twiddling their thumbs, do we?”) In fact, some subfields now function entirely by *preprints*, without any need for actual publications.

Regardless of the mode of dissemination, some papers are clearly more influential than others in establishing a new consensus, in undermining an old one or, quite simply, in provoking discussions among interested colleagues. On average, scientists claim to read about one paper per day, but I suspect it is actually less, so how do they choose which few papers they actually read? The answer is that this is decided largely by the reputation of the authors, but that the kind of journals that accept these authors’ papers is a large contributing factors in establishing their reputations, which is later consolidated by invitations to speak at prestigious conferences, national and international prizes, membership of learned societies etc. [5] All these factors determine not only how widely a certain result is disseminated, but also which topics (theories or experimental subjects) are ‘in fashion’, and therefore worth investigating and pursuing by less prominent scientists. The reader will certainly be aware of the fact that the kind of research I am discussing here costs a great deal of money, which is used to pay for an army of people and for expensive kit. Unsurprisingly, the same mechanisms are at work to determine what and who gets funded, because senior scientific advisors to government and members of grants-awarding panels are also senior members of the scientific community. Politicians like to think that they are ‘directing’ the flow of funds towards priority areas of national need, and this certainly has an effect at very high level (e.g., in deciding whether a given country needs more battery research rather than particle accelerators). At a lower level, it is unclear whether government prioritisation can have an effect in promoting research that results in tangible benefits, because it is extremely difficult to predict which areas will result in ‘impactful’ outcomes (after all, who could have predicted that the most significant practical outcome of CERN would be the World Wide Web?). For this reason, until recently, British research policy was governed by the so-called Haldane principle — the idea that decisions about what to spend research funds on, at least below the very highest tier, should be made by researchers rather than politicians [6]. The perception in recent times is that in the UK this principle has been significantly eroded, but it is still alive and well elsewhere. Most importantly, the top-tier European funding organisation (the European Research Council – ERC) operates entirely according to the Haldane principle [7].

Social interactions among scientists and between scientists and policy-makers clearly play a major role in establishing the 'pecking order', and this may appear to be deeply worrisome: if sociological factors determine not only how much attention a given paper or group receive, but also which direction research funds take, are the answers we give as a scientific community not already contained in the questions we ask? Are we not bound to find what we were seeking in the first place? The philosophical discussion of these questions would take us very far into a field that I am determined to stay clear of. I would rather tell a story, which, I believe, is both exemplar and hopeful.

III. A HOPEFUL EXAMPLE: THE STORY OF HIGH-TEMPERATURE SUPERCONDUCTIVITY

Superconductivity is the curious phenomenon whereby certain metals and compounds lose their electrical resistance *completely*: below a certain temperature, known as the *critical temperature*, the resistance becomes pretty much indistinguishable from zero [8]. Superconductivity is one of the few macroscopic manifestations of quantum mechanics, and is entirely incomprehensible by classical physics. It has many practical applications — for instance, every time you undergo an MRI scan, you or parts of your body are inserted into a superconducting coil — and at least eight Nobel prizes have been awarded in relation to it.

The story of superconductivity begins in 1911, when Dutch physicist Heike Kamerlingh Onnes, using liquid helium that he had produced for the first time, measured the resistance of mercury at very low temperatures [9]. When cooling below 4.2 K (degrees Kelvin, i.e., above the absolute zero, corresponding to –269 Celsius) he found something very strange, whence his famous logbook entry "Kwik nagenoeg nul" ('Mercury almost zero'). Like all good history of science stories, that of superconductivity is full of interesting adventures but is way too long to be recounted here in full. One crucial milestone is, however, worth describing: it is the famous 'microscopic theory of superconductivity', published in April 1957 in the journal Physical Review by American physicists John Bardeen, Leon Cooper and John Robert Schrieffer (all 1972 Nobel laureates for this work). Known after the authors' initials, 'BCS' theory describes electrons in the superconducting state as bound together by a sort of 'glue', which is provided by the crystal vibrations (known since the 1930s' as 'phonons'). BCS' simple and elegant quantum mechanical treatment described pretty much all superconducting phenomena known at that time, and made several important prediction that were accurately verified. By the late 1980s', superconductivity had developed into a sizeable business (mainly in science instrumentation) but was held back by the fact that the highest known critical temperature was a paltry 23.2 K. At the time, this seemed consistent with the theoretical developments of the BCS theory by Soviet theorist Gerasim M. Eliashberg, who calculated an absolute upper limit of $\sim$30 K for the critical temperature. The vast majority of known superconductors were either pure metals or metallic alloys, and all the best ones had the same cubic crystal structure. Working at Bell Labs and, later, at several US universities, German-born American physicist Bernd T. Matthias was one of the most prolific 'superconductor hunters' (he discovered hundreds of them), and had developed several empirical rules to help in his quest. Roughly stated, these rules were: high symmetry is good with cubic best, high density of electronic states is good, stay away from oxygen, stay away from magnetism, stay away from insulators, and stay away from theorists. With the exception of the last (only semi-serious) one, *from roughly 1960 to 1986, Matthias' rules represented a very good example of solid scientific consensus.*

Enter K. Alexander Müller. Born in Basel (Switzerland) and based at the IBM Zürich Research Laboratory, Alex was a well-respected physicist, especially in the fields of oxides and magnetic resonance, but in superconductivity he was regarded by most as a bit of an outsider. He espoused a rather obscure theory, according to which a hypothetical type of 'glue' (known as Jahn-Teller polarons [10]), could overcome the Eliashberg limit yielding much higher critical temperatures. Therefore, amidst general skepticism, Alex Muller and his collaborator J. Georg Bednorz set out to measure the electrical resistance of several oxides at low temperatures. In 1986, they struck gold: the critical temperature of a copper-containing oxide was found to be nearly 35 K, shattering the previous record, and a slight variant of the same oxide exceeded 42 K. It is noteworthy that these oxides were not particularly novel or rare — many chemist had them in their drawers for years, as one of them later confessed to me almost in tears; nor was Müller's apparatus particularly unique, although his previous work had equipped him well for this research. Berdnorz and Müller were promptly awarded the Nobel prize in 1987 [11], and their work spawned a veritable frenzy of research activities, which let to breaking one record after another. By the early nineties, the critical temperature of copper-containing oxides reached 135 K — a record that is still unbeaten today at ambient pressure [12]. An even greater effort was devoted to try to understand what was going on with these wonder materials. Two things were established rather quickly: that superconductivity in copper oxides had nothing to do with the Jahn-Teller effect (although Alex would probably disagree to this day), and that the 'glue' in this case seems to be magnetic. What about Matthias' rules? They are pretty much a pile of ruins: superconducting copper oxides have rather low symmetries (not a single cubic one in sight), they have very low density of electron states, they are — well— oxides, and indeed very close in chemical compositions to well-known magnetic insulators. Only the last rule still stands as a memento: "Stay away from theorists".

The characteristics of Berdnorz and Müller's revolution, — without a doubt, the most profound in solid-state physics since the discovery of the transistor — are extremely interesting. It was not a Kuhnian revolution at all, in that there were no previous state of crisis, no schools of thought vying for supremacy and no paradigm confrontation. In fact, to this day, there is no unified theory of high-temperature superconductivity commanding the universal consensus of BCS. Nor was the revolution produced by the stroke of genius of one isolated scientist, as it was the case for the general theory of relativity, since, as we have seen, Alex' intuition was most likely wrong. Instead, in many ways, it was similar to Columbus' 'discovery' of America: regardless of what he or his opponents thought, America *happened* to be there, and, once known, this simple *fact* changed the world forever. As a professional scientist, I find this observation to be a deep source of reassurance: if the humble facts, unloaded by any epistemological significance, can still produce revolutions of this kind, then, it seems to me, there is hope. Scientific communities may be fragmented, confrontational and swayed by the opinions of the 'great and good', but, unlike Cesare Cremonini with Galileo, they no longer refuse to look through the telescope. When they see facts such as "Kwik nagenoeg nul" they bow to them without hardly any resistance — in fact they embrace them enthusiastically. I conclude this section with a word of warning: you should not read this as an endorsement of some kind of logical positivist philosophy of science. All I am stating here is that facts change the way science is done as much as any social interactions between scientists, and sometimes much more.

IV. CHALLENGES TO SCIENTIFIC INTEGRITY

From what I wrote in the previous paragraphs, it should be evident that I hold a rather optimistic view of the integrity of science. This does not mean that I am inclined to ignore the sociological driving factors, both internal and external, that shape what the scientific community as a whole considers important, as well as the threats, actual and potential, to the integrity of the whole process. To simplify at the extreme, I believe that a good scientific culture should be one that facilitates both the kind of 'fact-driven' revolutions that I just described and also the other, 'theory-driven' kind [13]. One common characteristic of both kinds of revolutions is that they tend to unfold quietly, at least initially. In fact, heated scientific controversies about matters of grand principles are actually much rarer than people may think, and they certainly do not play out in public in the manner of the 'evolution controversy' that followed the publication of Darwin's *On the Origin of Species* [14]. Moreover, scientists responsible for these revolutions, though usually well respected, are generally not 'community movers and shakers' — in fact the overlap between the two subsets is arguably rather small. It follows that the integrity of the scientific process depends on the survival of the kind of scientists who can produce unexpected revolutions, but may not be the best at getting funded. I will briefly describe three kinds of threats to this 'endangered species', though there are undoubtedly more: over-management of science by politicians, sociological distortions and the effect of scientific frauds and mistakes.

A. Over-management of science by politicians

Perhaps surprisingly, science over-management results in *more* power (not less) being given to scientists who have the ear of the politicians and tell them what they want to hear. The outcome is usually that vast sums of money are directed to solving known problems with few, if any, low-hanging fruits, at the expense of the true pioneers who are exploring virgin territory. Because this has to do with money, it is more of an issue for expensive experimental subjects than for theoretical ones, and it is felt more harshly at smaller institutions with little independent means.

B. Sociological distortions

If evaluation criteria for scientific excellence were to change radically, true genius could be weeded out early on, simply for lack of recognition. This is less far-fetched than it might appear: people who make decisions about early-career fellowships, tenures, etc. increasingly rely on proxy indicators of scientific excellence, such as bibliometric indices and journal impact factors [15], which can lead to significant distortions. I will cite another example: the increasing importance of editorial decisions in top-tier journals, which is largely driven by the success of the *Nature* business model [16]. Articles in *Nature* (and in its non-profit American competitor *Science*) are not refereed in the normal way, but are heavily pre-filtered by professional editors *before* they are sent out for review. For many years, these journals have been accused of favouring certain scientists or subfields that happened to be in fashion [17]; these were usually dismissed as unproven allegations by disgruntled scientists who only wished they could publish there. When it concerned only a couple of journals, albeit very prestigious, this controversy seemed just a quirk of the scientific publishing world. However, in recent times, the *Nature* group has exploited its successful brand by 'spawning' a myriad of mini-*Natures* (58 *Nature 'Something'* — journals like *Nature Physics*, *Nature Biomechanical Engineering*, etc.) and 23 Nature Partner Journals, which still retain the brand as part of their name. Although not all equally prestigious, publishing papers in one of them has become almost a *must*, especially for early career researchers. Is this particular trend good or bad? One could argue that the editorial pre-filtering helps the journal to publish papers that are actually read (and therefore subsequently cited) — whence the tremendous success of the *Nature* family in boasting bibliometric impact factors. Others argue that the kind of papers that pass the first filter are those that please the editors with a fashionable topic, nice figures, and a clear and comprehensible message, not always backed by rigorous analysis, and that a proliferation of this style of papers in the very top tier distorts the sociology of science in a negative way. Although I do not wish to take side, there is little doubt that excessive editorial influence has at the very least the potential to distort the science discourse in a profound way. Scientists are indeed reacting to this in a variety of ways — from plainly refusing to publish in particular journals to rejecting the editorial process altogether [18]. I personally believe that the editorial process, especially of the non-profit kind, still has a hugely important role to play in science. Ultimately, all papers are peer reviewed, and all appointments are made by human beings, not by machines. It is the referees' and panel members' responsibility to exercise due diligence, looking past the headlines and ensuring that the true scientific quality of people and papers is commensurate with the reputation of the journals, institutions or grant schemes.

C. Frauds and mistakes

Few things do more damage to the reputation of science than the exposure of scientific fraud. In physics, the biggest scandal of all has been the 'Schön affair', which unfolded from about 1997 at the prestigious Bell Laboratories in the US, and was exposed in 2002. Jan Hendrik Schön, then a junior scientists at Bell, claimed to have made a series of amazing discoveries at the interface between semiconductor and superconductor physics. If confirmed, these would have had very important technological implications, and his papers were published in prestigious journals, including *Science* and *Nature*. As it turned out, Schön had thoroughly fabricated his data, and has done so in a very naïve and clumsy way, so it was relatively easy to expose his fraud. Another Bell scientists told me later that the lab was under intense pressure [19], and that publishing in *Science* and *Nature* was no longer considered sufficient — it had to be the *New York Times*. This, combined with the fact that senior scientists had little time to check the results produced by their junior colleagues (though they still appeared as co-authors!) was the root cause of the affair.

Though sometimes less visible by the general public, genuine mistakes also have significant potential to disrupt science, quite simply, by wasting people's time and resources. A famous example occurred in 1989, when Martin Fleischmann (a well-known electrochemist) and Stanley Pons, working at University of Utah, claimed to have achieved "a sustained nuclear fusion reaction" at ambient conditions — a phenomenon subsequently dubbed 'cold fusion'. The work generated an enormous flurry of attempts to replicate the phenomenon, but to no avail. Although at the time there were hints that the claim was fraudulent, in all probability it was simply due to the fact that Fleischmann and Pons were not careful enough (many called their experiments 'sloppy'), and in fact it seems that the two scientists were pressurised by their university to go public with high-profile press releases before they were truly ready.

D. Fighting science's pathologies

When confronted with these instances of 'pathological science', one should ask the following question: does the scientific community possess the antibodies to fight these pathologies, and how effective are they? I found it particularly useful to employ bibliometric techniques to look for answers. One can track the number of citations associated with a person or keyword, taken as a proxy for interest (if not belief), and determine how long it took for the trends to dissipate. In the two cases I mentioned ('Jan Hendrik Schön' and 'cold fusion' — see Figure 1), the characteristic width of the citation peak is approximately two years, which, I think, is remarkably short [20]. Once again, this seems encouraging: the community as a whole can apparently establish very quickly whether an extraordinary claim made by prominent scientists is wrong or fraudulent. Naturally, we can only establish this for things that we now know to be wrong. Is it possible that science as a whole or certain subjects within it carry a baggage of wrong facts or misconceptions, which linger on and perhaps are no longer even questioned? After all, early science is full of them — from luminiferous aether to phlogiston. It is extremely difficult to answer this second question with any certainty, and the absence of prominent cases of this kind, in which longstanding theories or solid data sets were completely debunked, is in a way a bit worrisome. The persistence of misconceptions seems more likely in research fields with limited possibilities of experimental design, since they have to rely on a relatively small number of crucial datasets — medicine and cosmology are among these subjects, albeit for completely different reasons. A very interesting case in this respect is that of 'dark energy': although this concept is widely accepted and is now part of the standard cosmological model known as Λ-CDM, a minority of prominent scientists is open to the possibility that it might be a misconception born of statistically insignificant observations [21].

V. THE EVOLUTION OF SCIENTIFIC CONSENSUS: COLLAPSE AND REBIRTH

The examples I have provided thus far sketch a rather peculiar portrait of scientific consensus: though rather robust against the spread of errors and misconceptions, previous consensus can collapse very rapidly when new facts or a new theory emerge, and it may take the scientific community a very long time to reform around new shared views. Here are some interesting questions:

i) Are consensus collapses predictable?

ii) Are there the equivalent of 'fault lines' that could be monitored, similar to what is done to predict earthquakes?

iii) Is the strength of scientific consensus even measurable?

Although I am not aware of any serious attempt to predict the collapse of scientific consensus, one could surmise that there could be several more-or-less typical 'failure modes'. One of them — falsification — was proposed by the philosopher of science Karl Popper in his classic 1959 book 'The Logic of Scientific Discovery' [22]. Falsification occurs when the current consensus theory or belief makes statements of the kind "There is no such thing as X" or "Y is not possible", and X or Y are subsequently found to be the case. A spectacular example of belief falsification occurred during the Bovine Spongiform Encephalopaty (BSE) crisis. BSE, also known as 'Mad Cow Disease', was first recognised in 1986. By the end of the decade, there was solid scientific consensus around the following statements: a) BSE had developed between 1970 and 1980, due to cattle eating feed containing the remains of sheep that had contracted a disease known as scrapie. b) BSE could not be transmitted to humans and therefore c) eating the meat of affected cattle was completely safe. One could have perhaps identified early on the weak link in this logic — the disease had already crossed the species barrier at least once, why could it not do it again? In fact, these beliefs were shaken in 1990, when a cat started to show signs illness, and eventually shattered completely when several dairy farmers fell ill with what is now known as variant Creutzfeldt-Jakob Disease (vCJD) — a direct consequence of exposure to the BSE 'prion' agent [23]. As of 2012, about 170 cases of have been recorded in the United Kingdom, and 50 cases in the rest of the world. A much larger but unknown number of people have certainly been exposed to the prion, and, given the very long incubation period, may yet die of the disease (there is no treatment for vCJD). I believe that this crisis is paradigmatic, because the sudden collapse of previous beliefs was followed by a period of uncertainty about the true consequence of the crisis for human health — a situation that is to some extent still persisting. Epidemiological studies would suggest a very small impact with an ultimate life toll not exceeding a few hundreds, but there is no consensus on epidemiological modelling of vCJD.
Regarding my last point iii), attempts have been made to measure the strength of consensus by analysing networks of bibliographic citations [24]. The idea is that, at times of controversy, the community would be divided in opposing factions, each with their own set of 'classic' reference papers and

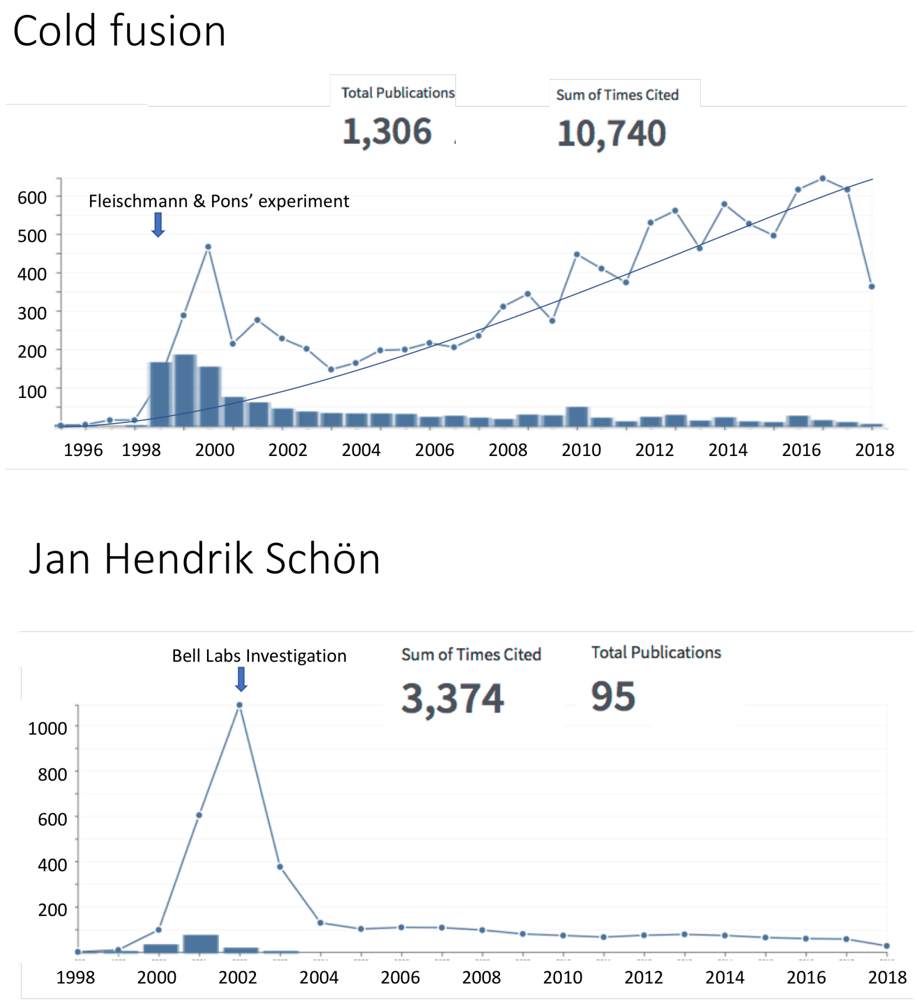

FIG. 1. (Number of publications (blocks) and related citations (dots and continuous line) for papers containing the keyword 'cold fusion' **top**) and authored by Jan Hendrik Schön (**bottom**). The increasing background of citations to 'cold fusion' papers relates to a different field of nuclear physics.

with a tendency to cite each other rather than their opponents. The degree of 'epistemic rivalry' could be measured through an index, known as network modularity score, which can be extracted from patterns of citations. The idea is extremely interesting, and indeed the signatures of known controversies (such as: "Is smoking harmful?") in the modularity score are clear, though rather weak. However, the biggest problem with this method is that it cannot distinguish between positive consensus ("We all agree that X is true") and negative consensus ("We all agree that information is insufficient to establish whether Y is true or false").
One might therefore ask a further question, perhaps the most important one:

iv) How can non-specialists establish what the consensus is on a given topic?

As a preamble to the full answer, one should observe that *true consensus can only exist outside active research topics*. This is clear if one considers that *scientists regard as their mission to challenge established views*, so, paradoxically, *consensus in active areas of research is somethings nobody fully believes*. Nevertheless, the following are clearly two distinct questions, and Alex Müller would have answered them very differently:

a: What do you believe to be true?

b: What do you think the majority of your colleagues believes to be true?

This observation is the basis of my answer, which is very simple, if somewhat naïve. If you want to establish the consensus on a given topic, ask scientists in that sub-field and proximal sub-fields *what the consensus is* (not what they believe). I am persuaded that, in most cases, the answers would be remarkably consistent.

VI. SCIENCE AND GOVERNANCE: A PRAGMATIC VIEW

The view of science I have outlined poses a great challenge to those who are concerned with political and economic governance: if even healthy, non-pathological science is bound to undergo very rapid changes, is it still possible to state, with Bertrand Russell that "Science is at no moment quite right, but it is seldom quite wrong..." [25]? A global perspective of rational governance is even more challenging in the light of theories of social constructivism, according to which "First-world science is one science among many" [26].

In this brief note, I clearly do not have the space for an extensive discussion of these topics. I shall, therefore, adopt a pragmatic view, which is only concerned with the functional value of science, completely disregarding its epistemic value. Even if it is nothing else, science is indubitably a *tool* to predict the future in a more or less approximate way, and to transform it through technology [27]. This point of view has an interesting implication: inasmuch as it is a tool, science could be *shaped by function*, and therefore be, to some extent, culturally non-specific; in the same way, hammers or swords, though somewhat different across cultures, do not cease to be recognisable as such. This is in fact the universal perception of science practitioners, when they meet colleagues from other cultures [28]: unquestionably, they do things in a slightly different way, but differences tend to disappear when we 'go down to business and talk science'. Moreover, this pragmatic view of science can still function against extreme skepticism, and to some extent represents a protection against challenges to authority (it is a lot easier to challenge the authority of a philosopher than that of a dentist...)

Let us therefore be pragmatic and define science as a tool to predict the future, no doubt an imperfect one but arguably the best we have. How are those concerned with governance to make the best out of this tool? Here is a set of equally pragmatic rules:

- *Attempt to establish what the consensus is and how strong it is with the same ethical rigour that you expect of your scientists.*

The strongest temptation for decision makers is to endorse minority views that suit their constituency. This is extremely easy because, as I have just stated, nobody fully believes in the consensus view. Moreover, when collapses in consensus occur, there is a very high statistical probability that somebody had predicted the 'right' answer, so there are always plenty of past examples of successful 'challenges to the establishment' to draw upon. These days, many people are inclined to believe in conspiracy theories, whereby consensus is swayed by lobbies through money ('Big Pharma', allegedly suppressing evidence against vaccines) or influence ('The Liberal Elite' allegedly peddling anthropogenic global warming). Although external influences should never be underestimated, history teaches that they were rarely successful in suppressing evidence for very long [29].

- *Always ascertain what the consensus* actually *says*.

This is particularly difficult when consensus is based on statistics rather than experimentation, as in the case of epidemiology. Here is a topical example: the overwhelming scientific consensus is that common vaccines such as MMR are reasonably safe, but what does this mean? It is in fact entirely possible that for a very small subsection of the population, as yet unidentified, some vaccines could be significantly more harmful than for the vast majority of others. However, in the absence of an identified correlation, refusing vaccination is plain irrational. Here, the consensus has no obvious fault lines, in that identifying the at-risk subgroup (if it even exists) would still result in the majority of the population being vaccinated.

- *Try to establish possible fault lines in the received wisdom, and always ask 'what ifs...'*

Once again, we can gain useful insight from the unfolding of the BSE crisis: in the UK, a ban on the use of high-risk offal for human consumption — the brain, spinal cord and spleen — was issued in 1989, before any evidence of human transmission. The unknown civil servants responsible for the ban have undoubtedly saved hundreds, if not thousands, of lives.

VII. ACTIONABLE EVIDENCE

I conclude this brief essay by proposing the concept of *actionable evidence*. I ask the following questions:

a: When is the scientific consensus sufficiently strong to demand action by policy-makers?

b: Is it conceivable that policy-makers may adopt a *less prudential* course of action than the one recommended by current scientific consensus?

c: When should policy-makers adopt a *more prudential* course of action than the one recommended by current consensus?

I should say from the outset that my answer to question b) is a resolute 'no': fortunately, adopting policies that are openly against prudent scientific consensus is politically unacceptable in most countries, and policy-makers usually prefer to deny the existence of consensus by fabricating non-existing controversies. This is precisely where the methodology I have introduced could be most useful: if a rigorous process to ascertain consensus was to be widely accepted and adopted (e.g., by the press), the fabrication of controversies for the purpose of political gain would be at the very least more complicated. In answering a), I propose that, when the action recommended by scientific consensus is also the most prudent, it should be taken as representing *actionable evidence* and should be adopted without hesitation. A case in point are actions recommended to counter anthropogenic global warming — a field where there is near-universal consensus without obvious fault lines, and where urgent action is demanded [30]. The answer to c) is more nuanced: I already mentioned the decision to ban human consumption of offals during the BSE crisis. This was a more prudential approach than recommended by scientific consensus, and it clearly had a short-term economic cost, but was undoubtedly correct. Risk perception is also a factor here: strict regulations around very-low-level radiation exposure and slight radioactive contamination are not supported by scientific consensus, but any relaxation would be politically unacceptable because of a heightened perception of radiation risk. Ultimately, we must conclude that question c) does not have an exact answer: like science, politics contains an element of art and intuition. When due diligence is done, when the scientific community and the electorate have been consulted, policy makers should form their own judgement and take responsibility for it.

---

[1] Following Alan D. Sokal, this Enlightenment idea can be summarised as follows: that there exists an external world, whose properties are independent of any individual human being and indeed of humanity as a whole; that these properties are encoded in"eternal" physical laws; and that human beings can obtain reliable, albeit imperfect and tentative, knowledge of these laws by hewing to the "objective" procedures and epistemological strictures [sic] prescribed by the (so-called) scientific method. A scientific realist, Alan D. Sokal was responsible of a famous hoax, in which he pretends to dismantle this view by quoting several social constructivist philosophers and imitating their style and language. This nonsensical paper was published in the journal *Social Text* [31], together with serious contributions by his competitors, marking the culmination of the so-called "science wars" of the late 1990s'. Sokal's paper is humorous and well worth reading.

[2] T. S. Kuhn, *The structure of scientific revolutions*, 4th ed. (University of Chicago Press, Chicago, London, 2012 [1962]).

[3] As Kuhn observed, a book is more likely to detract from one's scientific reputation than to enhance it.

[4] An editor of the prestigious journal *Nature* recently recounted the curious interaction between Einstein and one of his predecessors, dating back to the late 1930s'. Having submitted a paper to *Physical Review* on gravitational waves with his assistant Nathan Rosen, Einstein received a letter explaining that the journal would be happy to publish it, after he had considered the revisions (entirely correct, as it turned out) suggested by an external reviewer. Einstein's response was along these lines "Dear Sir, we (Mr. Rosen and I) had sent you our manuscript for publication and had not authorised you to show it to specialists before it is printed. I see no reason to address the in any case erroneous comments of your anonymous expert. On the basis of this incident I prefer to publish the paper elsewhere".

[5] I am generalising a bit here, since not all communities operate in this mode, but what I wrote corresponds to the *modus operandi* of a large section of the scientific community. In the UK, for example, lab-based physics and engineering receive 29% of the grant funds, followed by medical research (22%) and 'big science' research (15 %, largely for particle physics and astronomy, but also including the large research infrastructures that I use).

[6] It is named after Richard Burdon Haldane, who in 1904 and from 1909 to 1918 chaired committees and commissions that recommended this policy.

[7] In this and other respects, the ERC is, I believe, one of the best science funding programmes worldwide.

[8] The time it would take for a superconducting current to decay in a closed loop is comparable to the age of the universe. By contrast, currents in a loop made of the best 'normal' metals (e.g., very pure silver) would decay in about 100 *micro*seconds.

[9] P. F. Dahl, Historical Studies in the Physical Sciences **15**, 1 (1984).

[10] And yes, Teller is Edward Teller, the 'father of the hydrogen bomb'.

[11] J. G. Bednorz and K. A. Müller, "NOBEL LECTURE - Perovskite-type Oxides - the new Ap-

proch to High Tc Superconductivity," (1987).
[12] The absolute world record for a superconducting critical temperature is 203K (–70 Celsius), which is 'room temperature' in winter in some parts of the world. It was reached in 2015 by a German team by applying a pressure of 1.5 million atmospheres to humble H_2S, better known for its rotten eggs smell. Ironically, by all evidence H_2S is a conventional BCS superconductor.
[13] One recent example of the latter in solid-state physics has been the 'topology revolution' of the early 2000s', which led to revisiting both theoretically and experimentally many materials and phenomena previously considered rather boring.
[14] An interesting recent example in physics has been the controversy over solar neutrinos — a problem that hang around for about 40 years [32] and was eventually resolved at the turn of the millennium, leading to the 2015 Nobel Prize. Initially framed as a controversy between physicists and observational astronomers, the terms of its resolution would have surprised both camps.
[15] Bibliometric indicators rely on the number of citations received by a given paper in a given time, which is now readily available through bibliographic search engines. Journal impact factors measure the average number of citations received by each paper published in that journal, and are used as proxies for its prestige. Personal impact factors and derived metrics do the same for the quality of a scientist. Both are imperfect and prone to distorting scientific output.
[16] First published on 4 November 1869 with the mission to inform scientists of key advances outside their own field, *Nature* is arguably the most prestigious of all scientific journals. Its parent academic publishing company, Springer Nature, reported revenues of EUR 1.64 billion in 2017.
[17] See, for example, "How journals like Nature, Cell and Science are damaging science", by Randy Schekman, The Guardian, 9 December 2013.
[18] The recent paper on gravitational wave detection by the LIGO and VIRGO collaborations (B. P. Abbott, *et al.*, Phys. Rev. Lett. **116**, 61102 (2016)) — perhaps the most important physics paper of the last few years — was not published in *Nature*. As already mentioned, some close-knit communities now function entirely by *preprints*. Others propose to replace the editorial model, usually by some kind of on-line voting scheme.
[19] Founded in 1925 but tracing its origins to the end of the nineteenth century, Bell Laboratories was the premier industrial laboratory worldwide, and was responsible for breakthroughs such as the transistor and the discovery of the cosmic microwave background radiation. Eight Nobel Prizes have been awarded for work completed at Bell Labs. Once owned by the all-powerful American Telephone and Telegraph Company (AT&T), in 1996 Bell Labs were taken over by Lucent, a much smaller telecom company formed after AT&T's divesture of its technology business unit, and is now owned by Nokia. After the Schön scandal, its interest in basic research has waned, but its influence has not, since former Bell alumni are in key science decision-making positions worldwide.
[20] Sometimes, errors are discovered even faster, before they had the time to hit the peer-reviewed literature (but, sadly, not the general press). A well-known case is that of neutrinos allegedly travelling faster than the speed of light. In 2011, the OPERA experiment — a collaboration between CERN and the Gran Sasso Laboratories in Italy — mistakenly observed neutrinos travelling 0.002% faster than the speed of light, and disseminated this result through preprints and press releases. During peer review in *Journal of High Energy Physics*, the error was eventually traced to a loose fibre optics cable, and many subsequent experiments confirmed that neutrinos do indeed travel at the speed of light. This consensus was reached very quickly, but not before a true firestorm of interest had hit the press and the general public.
[21] J. Colin, R. Mohayaee, M. Rameez, and S. Sarkar, (2018), arXiv:1808.04597.
[22] K. R. Popper, *The logic of scientific discovery* (Routledge Classics, Abingdon (UK), New York (USA), 2002 [1959]).
[23] A prion is a mis-folded protein, which can cause other identical proteins to mis-fold in a chain reaction. Further human exposure to the prion was curtailed by massive herds slaughter at the cost of several billion pounds in the UK alone.
[24] U. Shwed and P. S. Bearman, Source: American Sociological Review **75**, 817 (2010).
[25] B. Russell, *My philosophical development* (1959).
[26] Paul Feyerabend, introduction to the Chinese edition of 'Against method', 1988.
[27] The Italian philosopher E. Severino espouses an even more radical idea: modern science would have renounced any claim to the truth, replacing the 'Will to knowledge' of the ancient Greek philosophers with Nietzsche's 'Will to power'.
[28] Originating in western Europe and, later, in the US, 'western-style' science is widespread across the continents, with a strong representations in Asia, South America and to a lesser extent Africa.
[29] An interesting counter-example is presented by U. Shwed and P. S. Bearman [24]. By using citations network analysis, they show that controversy regarding the dangers of tobacco smoke was provisionally settled and reopened in different formulations over a period of 25 years (1960–1985), fuelled in part by research funded by tobacco companies.
[30] It would be too long to summarise the state of the climate change 'controversy' here, but in essence there is no evidence at all of a controversy among scientists, because it was agreed long ago to asses risk through statistical analysis of different computer modelling predictions. The 'controversy' is likely to have been fabricated by politicians by hand-picking outlier predictions.
[31] A. D. Sokal, Social Text **46/47**, 217 (1996).
[32] J. N. Bahcall and R. Davis, Science **191**, 264 (1976).