

Normal Accidents of Expertise

Stephen Turner. 2010. Normal Accidents of Expertise. The definitive version of this paper has been published in *Minerva* 48:239-258 by Springer, LCC, all rights reserved. It is available at: <http://www.springer.com> DOI 10.1007/s11024-010-9153-z

Stephen Turner
Department of Philosophy
University of South Florida

The organizational theorist Charles Perrow invented the term “normal accidents” ([1984] 1999) to apply to the kinds of failures that are inevitable in complex systems, such as chemical plants or nuclear plants, which are “coupled” in a tight way, meaning that when something went wrong in one part of the system it would have consequences for other parts of the system (Perrow [1984] 1999, 72, 89-100). Some of the consequences in such systems are predictable linear consequences of failures of components in the system. Others, which will be my main concern here, will involve multiple failures of components with unpredictable consequences. These consequences can include catastrophic outcomes. Loosely coupled systems, such as an assembly process that is made up of several component assembly processes that operate independently but need to converge at the final assembly of a product, also fail, but they allow for more time to correct for failures and to invent responses. Tightly coupled systems are more sensitive: responses need to be immediate and routinized, which means that failures must be anticipated in advance. The failures that can be anticipated in advance are those for which redundant systems can be created or standard operating procedures for correction can be applied in the face of failure. Many failures in a coupled system, however, result from unanticipated interactions between elements of the system in the face of the failure of multiple elements of the system, which produces problems that cannot be solved by standard procedures. When tightly coupled

systems fail in this way, the results can be catastrophic. Loosely coupled systems are more forgiving.

Producing expert knowledge is different from producing nuclear power or chemicals, but there are similarities. Error and error correction are a major element of what systems producing expert knowledge do, and they involve a complex division of labor which is also a division of knowledge: technical or expert knowledge of highly specialized kinds have to be brought together to make the system work. In some respects, however, producing expert knowledge and producing chemicals are especially similar. In big science there is always complex instrumentation which must be made to work, and which must be understood. Errors involving the functioning and design of instruments are endemic to big science. The technologies are often unique, cutting edge, and difficult to test in the conditions where they are designed to be used. Unanticipated interaction effects have the potential of producing catastrophic results.

There have been two major failures of expert communities in the last decade: the failure of the economics community to grasp the implications of the rise of various financial instruments in relation to mortgages and the alterations of mortgage policies, and the failure of the climate science community to anticipate the drop in the rate of temperature change over the last decade. These failures have been much commented on, and the discussion is still in the phase of lack of consensus— including a lack of consensus over whether there was a failure, or whether the failure was major or minor. Each involved complex interactions. It is appropriate to ask whether the model of normal accidents applies to these cases. In what follows, I will do no such thing, though I will return to the issues with climate science in the conclusion. It is clear, however, that much of the political discussion of experts and expert consensus is blind to the general problem of catastrophic expert failure, which will be my concern here. Of course, this issue relates to the

problem of skepticism about experts, as it should. But the point of the paper is to identify a specific problem that arises with systems producing expert knowledge.

The Question of Symmetry: Can Science Studies Talk about Failure?

Talk of “failure” and “error” in science studies invites inverted commas: according to the methodological principle of symmetry, failure and error are socially constructed in the same way as “true” scientific facts are. Is this consideration relevant here? The field has itself seen a sea change in the discussion of such topics, especially in relation to risk. Major figures such as Harry Collins and Bruno Latour have modified their views in order to avoid being seen as skeptics, particularly about global warming. Latour says that

I intended to emancipate the public from prematurely naturalized objectified facts. Was I foolishly mistaken? Have things changed so fast? In which case the danger would no longer be coming from an excessive confidence in ideological arguments posturing as matters of fact-- as we have learned to combat so efficiently in the past-- but from an excessive distrust of good matters of fact disguised as bad ideological biases! While we spent years trying to detect the real prejudices hidden behind the appearance of objective statements, do we now have to reveal the real objective and incontrovertible facts hidden behind the illusion of prejudices? (2004, 227)

Collins has claimed that we are in a third wave of science studies in which the reality of expertise is accepted. Michael Lynch suggests that the source of the problem, the symmetry thesis, itself

be treated as problematic: He revises Bloor's famous formulation, to suggest a discovery program that "would be impartial with respect to truth or falsity, rationality or irrationality, success or failure [*and symmetry or asymmetry*]. Both sides of these dichotomies will require explanation" (Lynch 2010, 11-12). These are all efforts to problematize the problematizing device of symmetry itself. But is this necessary? And does it bear on the application of Perrow's argument to expertise?

Answering this requires a return to some fundamentals. The term symmetry derives from a distinction in explanatory form. Causal explanations are symmetric; teleological explanations are not. In a teleological explanation, there is a favored endpoint that does not require further explanation. The classic example is the Aristotelean theory of motion: an object seeks its appropriate final place (Taylor 1964, pt. 1, chap. 1, sec. 4, pp. 21-25). Once it is there, there is no more motion, except for disturbances of the natural order. Objects "fall" in order to get to their place in the natural order. The order itself is explained by some higher principle. In the case of science, this kind of explanation involved truth or correctness and error. The history of science traditionally was written as a mixed history of truth and error. The truths, or at least the fact that scientists correctly believed in true facts, didn't require explanation—they had, so to speak, arrived at their natural end. Error, however, did require explanation.¹ In writings like Laudan's *Progress and its Problems* (1978), social explanations were given the task of explaining error, but not truth. Dispensing with symmetry meant dispensing with the favored status of correctness. But it also meant a change in the object of explanation. The teleological view ran together belief and truth: truth became a sufficient explanation for belief. The symmetric view was that truth was irrelevant to explanation of belief: beliefs were natural facts of their own, and in the real natural world of belief, explanations of true and false beliefs worked in the same way. Truth in a

final sense was a separate matter.

One advantage gained by this convention of explaining belief rather than truth, and explaining symmetrically rather than using truth to explain belief, is to avoid retrospective bias—to recreate the situation in real time—as the scientists faced the issues in question. But this is symmetry about causes or explainers, not about the thing to be explained. The thing that was being explained was the naturalistic fact that people believed something to be true. Science studies scholars elaborated a complex vocabulary to describe taken-to-be-true facts without calling them true. In actor network theory, the way was to show that elements of the network had enrolled more additional elements in their network than other things. In Andrew Pickering's writings on practice, terms like 'hypothesis' and 'test', along with 'failure', were replaced by the idea of the dialectic of resistance and accommodation (Pickering 1995, 22; Turner 1999, 152-5). But the point of the explanations, terminological gimmicks aside, was precisely to account for the fact that some things were taken to be fact. So the objects of explanation of the asymmetric accounts and the symmetric accounts were the same: what science arrived at and took to be true. They were merely described differently. It is an additional, and problematic, thesis, that descriptions using "true" and "fact" without quotations can never be justified, and even whether science studies itself can avoid appealing to these notions. This was the source of a large literature on reflexivity in the 1980s, and a response to the fascination in the same decade among philosophers of science with "scientific realism," which was the thesis that the beliefs of science were indeed true, and that this was the explanation of the success of science.

In the discussion of examples that follows, I will dispense with, or at least bracket, the convention of symmetry as well as abstract concerns about premature naturalization of facts. The cases I will discuss here are both cases in which what was taken to be true changed from a

situation of confusion to one in which the relevant facts came to be taken for granted truths. The language of failure and error was employed, unproblematically, by the people involved in the state of confusion: the confusion was motivated by what they took to be failures, errors, contradictions, and the like. Trevor Pinch, in his study of one of these cases, uses the sanitizing term “perceived contradiction” (1986, 156). In this case the contradiction was perceived by everyone, and was there, so nothing is added by noting that it was “perceived.” The dispute was about the meaning of the contradiction. There was no disagreement or confusion about the factual character of the events themselves or the fact that they were failures, in the banal sense of unexpected and unwelcome outcomes. The same, of course, cannot be said of discussions of climate science. In the case of climate science, there are people who deny that any failures of prediction have occurred. Yet in this case as well there is a more or less clear acceptance by virtually everyone that there has been an unexpected and unwelcome outcome, and active discussion of the meaning of this outcome. For this reason, I will suggest, the situation parallels the early moments of the cases I will describe in more detail.

Perrow’s oxymoronic term “normal accidents” captures the core idea of explanatory symmetry without ever appealing to the term itself, or entangling him in the issues of reflexivity about facts. His point is to show that the *same* organized technical processes that produce success in a complex production technology also, as part of their “normal” operations, produce “accidents,” namely, unintended outcomes that are unanticipated and also not corrected for by normal safety measures, such as redundant systems, pressure valves, shutoff mechanisms, circuit breakers, and so on. The effect of this is to redefine “normal” in relation to “accidents”: to describe a class of accidents that arise from the “normal” organization of complex productive processes. The accidents he wants to “normalize” or treat symmetrically are accidents beyond

those which the safety mechanisms for the discrete elements of the process ordinarily correct.

The point of this paper is to extend this reasoning to expert systems: to show that it is “normal” for unintended and unanticipated outcomes to occur in expert systems that are not correctable by the usual safety mechanisms-- meaning in the case of science such things as peer review, replication, and findings from related work in the network of science (cf. Polanyi [1946] 1964), and also to show that the problems are especially acute for tightly coupled expert systems. The cases I will discuss here-- the failure of the Large Hadron Collider and the Homestake experiment-- do not involve problematic asymmetries. The results of the Homestake experiment were unambiguously a “problem” for the science involved. No one seriously denied that something was wrong, somewhere: the predictions had failed. There was no favored solution: instead there was spirited and contentious disagreement. As Pinch showed, the disagreements about the meaning and importance of the failure were connected to the social locations of the participants and critics, especially their disciplinary background (1986, 172-5). But this point turned out to explain very little about the outcome of the controversy. The winners won for reasons familiar to STS scholars: the main findings, which the experimenter obdurately held to, were eventually linked up to other experiments, and the claims of his critics, a few of which were plausible at the time that the main finding was a lonely anomaly, could not accommodate these experiments. One of the interpretations proposed showed how the results could be accommodated by changing the theory in a way that accorded with some earlier speculation in physics. It turned out to be correct.

In short, the critics of the experiment lost, whether one chooses to describe this loss “sociologically,” philosophically, as was done at the time by Dudley Shapere in a prescient paper, “The Concept of Observation in Science and Philosophy” (1982), or in terms of the way

the scientists themselves eventually came to understand the situation. In the cases of failure I will discuss here, the term serves equally for its STS sense and its sense in the language of the scientists themselves. Indeed, if these languages were not largely translatable, it would be impossible for STS to explain science. In using the terms in this way one need not violate explanatory symmetry. Truth is not used as an explanation, and using the terms does not amount to acting as an omniscient observer illegitimately taking sides or producing a Whig history in which the victory of the winners is construed as the inevitable outcome of the historical process. Symmetry is a convention that avoids these mistakes. But it is a convention that makes sense only in terms of the purposes it serves, and it serves these purposes in some contexts and not others. In the cases to be discussed here, artificially problematizing the language of error and failure, does nothing to change the explanation.² Symmetry about causes, in the form of Perrow's showing that accidents are the result of normal processes, does change the explanation, and it is this kind of symmetry that will be preserved in what follows.

CERN: Normal Accident?

The CERN Large Hadron Collider is well known and no purpose would be served by describing it here. The system itself is a vast collection of electro-magnets— 1734 large magnets and 7724 smaller corrector magnets. The magnets direct and accelerate proton beams placed in a tunnel cooled by superfluid helium, to produce collisions whose results are captured by detectors. The experiments require a massive amount of energy, and energy levels within the experiment itself are an order of magnitude higher than in previous accelerators. After its launch on September 10, 2008, it suffered a major failure on 19 September 2008 which has been extensively documented

(Bajko et al 2009; CERN 2008; Chalmers and Henderson 2008; Hill 2008; Rossi 2010). It was shut down, a large effort was put into diagnosing the failure, and it was repaired, with the systems that failed replaced or cleaned and fixed, and redesigned. The failure occurred during powering tests of the main dipole circuit.

The symptom of the failure was a release of helium that heated the super-cooled tunnel: a “quench.” The initial response to the accident followed the pattern described by Perrow: confusion and a radical underestimation of the seriousness of the problem: the wrong outcomes are “not only unexpected, but are *incomprehensible* for some critical period of time” (Perrow [1984] 1999, 9). Although the leak of helium and the loss of helium were reported in the logbook, the fact that the fire brigade was called was entered but later deleted. “Sources” were upbeat:

A CERN source who had initially feared that the incident could delay the LHC’s operations by several weeks said hopes were rising during the afternoon that the damage was limited. “They just managed to bring back the compressors again, so who knows,” the source said. “These guys make miracles and could bring back the system in a week.” (Chalmers and Henderson 2008)

There were no miracles; further inspection found considerable and unexpected damage. Exactly the kind of confusion described by Perrow set in. An extensive investigation was launched, and preliminary report was made (CERN 2008), then a follow up report (Bajko et al 2009).

The basic facts of the event as described in the reports and the subsequent explanations by the chief engineer for the project (Rossi 2010) were these. The cryostatic vacuum enclosure in

which the experiments were conducted had a safety relief device which allowed for the release of helium into it in case of failure, but it was not designed for the pressures that developed, which were “considered ‘beyond design’” (CERN 2008, 1), and it did not release helium at the rate needed to avoid damage to the system. In short, this was an anomalous event, not anticipated in the design. The mystery was the cause of the event. The proximate cause settled on by the investigators, with the caveat that, as the parts were destroyed in the accident, any conclusion must remain “speculative” was a defective bus connector between electromagnets. The components had been extensively tested, in accordance with normal procedures for preventing accidents, though it happens that the particular section that failed was the only one of eight that had not been tested. They had not, however, been tested with one another in the system itself. They were vulnerable to unanticipated interaction effects, and in this case one of the most mundane: an arc or short resulting from excessive resistance in one of the components.

The “normal” safety devices worked as designed. The power tripped off when resistance was detected, and the relevant systems— the quench or helium detector and the energy discharge systems that rapidly powered down the system by inserting resistors-- all worked as expected. In a second, however, an arc developed, blowing a hole in the helium enclosure, letting the helium into the beamline— several tonnes worth. This was an unanticipated interaction which involved first two failures, which led to a third. The beam vacuum degraded and the insulation began to degrade in neighboring subsectors. When the pressure reached the level of atmospheric pressure the spring-loaded relief discs opened, and helium was released from the tunnel. But the release valves, not designed for this rate of escape, could not prevent the internal pressure from rising with such force that it pulled magnets from their supports and filled the tunnel with soot. This sequence was not evident at the time. As Lyn Evans, the project’s leader, told *The Times*, “It was

only by going back through the data that we see a signal that was considered insignificant at the time. A helium valve was opening and that's because the fault was drawing more and more helium to keep things cool, until it couldn't any more. The safety system was not good enough to cope" (*The Times* 2009).

The electrical fault that led to the arc was presumably the result of improperly soldered connections, which in turn was the result of a decision to use a solder with silver and tin content, which had a higher melting point and didn't flow as easily as lead, resulting in bad connections. The decision to use this solder was based on a desire to protect the workers from lead. Given that the instrument required 1200 tonnes of superconducting wire for the magnets alone, it is no surprise that there would be a short somewhere when the system was powered up. As the engineer in charge, Lucio Rossi, admitted, it occurred in a "relatively low-tech system," namely "an interconnection between magnets." But he argued that the nature of the problem, however, reflected a complex set of "faults in design" that "even if comparatively trivial, are intimately connected with superconductivity and its subtleties" (Rossi 2010, 1). What interacted, in this case, was a set of faulty designs— the solder and the magnet connection, parts that worked as designed but failed in the unanticipated conditions, signals that were missed, valves that were not designed for the event as it occurred, and so forth. This was not a simple case of a part failing to work and its backup failing to work.

Other physicists familiar with such systems anticipated that failures of this general kind would occur. Their initial responses, unlike the CERN sources, were that the problems were no surprise. "It's what happens when you start up a big, superconducting machine," Judy Jackson, spokeswoman for the Fermi National Accelerator Laboratory outside Chicago, told the Associated Press. "Our impression is that what happened last Friday at CERN caught them more

by surprise than it did us” (Hill 2008). The problem was not that there was a lack of normal measures, such as backup systems. “The thing about the LHC was it has not just caused a quench, but there are systems that are supposed to prevent it from melting and dumping helium,” said Michael Harrison, who worked on Fermilab's Tevatron collider and designed and built the United States' other superconducting collider at Brookhaven on Long Island. “So it was obviously something else that went on as well” (Hill 2008). The redesign and fix for the problem reflected the fact that multiple systems failed. New systems were designed to solve these problems (Henderson 2009).

What is Knowledge Risk?

We can identify three key places in this narrative of the accident where things went wrong or were at risk of going wrong because of errors or limitations of our own understanding. These are, to coin a phrase, knowledge risks: risks we incur in the course of acting on imperfect knowledge. The first involved the low-tech system, the bus, which failed because of the solder used in the connections. A decision to use a silver compound solder was made for safety reasons. The probability of an electric fault was increased by this decision. The increase was far greater than realized at the time of the decision. A second involved the failure to recognize the significance of the signal that the helium valve was opening. The third was the generic risk identified by Rossi: that strange things happen in relation to superconductivity. Knowledge risks are ubiquitous: there were hundreds, perhaps hundreds of thousands, of decisions made in the course of the construction of the collider that involved similar kinds of risks. Each of them might have gone wrong— failed to produce the expected results— for reasons that were not known to, or understood

by, the decision-maker.

Science is a means of reducing knowledge risk, though only one means among many. Experience, tacit and explicit, also reduces knowledge risk: one prefers an experienced surgeon to the one with the most recent textbook and little experience. But science occupies a special role in relation to knowledge risks because it is a system for producing the knowledge that reduces these risks. Like any system, it has its own risks, and its own knowledge risks. It can be characterized in Perrow's terms, as loosely or tightly coupled. Classical science was a very loosely coupled system. There were many errors made by scientists, but there was plenty of time to correct errors, relatively few serious consequences of error, especially in pure science, and there were good mechanisms for correcting the errors. Much of the literature on science in the period after the Second World War, which looked backwards to science as it was before the war, was concerned with these corrective mechanisms. The fact that scientists had a career incentive to refute the conjectures of others (Popper [1963] 1965) and to replicate the work of others (Zuckerman and Merton 1971; Ziman 1978) plus the network-like, interdependent structure of science, which meant that errors would be detected as a result of the use of findings in adjacent areas (Polanyi [1946] 1964), were taken to be key to the error reduction process.

The science system and the institutions of science supported these processes. Science was a 'liberal' institution, in which authority was always exercised indirectly. Scientific societies did not proclaim orthodoxies nor (except for a small number of incidents in which national societies intervened in controversies) attempt to settle questions of scientific truth directly. Consensus was allowed to develop on its own. And in some cases it took centuries to happen. Practices of peer review were developed which facilitated the forming of consensus by excluding crazy and unscientific ideas from journals. But peer review had a light hand— it was understood to be a

means of screening out obvious error rather than affirming one true answer.

The idea was that truth would emerge without coercion in a community of independent scientists forming their own ideas about what was true. The consensus that emerged was all the stronger for being uncoerced. And it was open to correction by the same means as it was formed— the unforced assent of competent practitioners. There were standards— standards of competence, for example. The club was exclusive, or at least excluded people. Academic positions were hard to come by. But accredited scientists were free to form their own opinions, pursue their own hunches, and report their results if they were not obviously erroneous in the expectation that truth would emerge from the process of conjecture and refutation that was supported by the competition between scientists for reputation and recognition and from competition between research institutions to employ the most meritorious and visible scientists and associate themselves with their discoveries.

We can think of knowledge risk on the model of errors of hypothesis testing: the error of accepting false hypotheses on the basis of the data, and the error of failing to accept true ones. Both types of error were filtered out by this system, but slowly. The error of accepting false hypotheses was filtered in several ways, first by the system of publication, which relied on peer review of one kind or another to exclude obviously wrong assertions or obviously bogus findings, by the system of granting degrees and society memberships only to competent practitioners, and by the fact that there were career incentives, in the form of recognition and reputation, that rewarded scientists for disproving widely held scientific ideas. The error of failing to accept true hypotheses was corrected for by the need for scientists to produce novel findings: a true hypothesis was likely to be a fruitful one, and scientists could not afford to ignore this fruit. Rivalry, indeed, meant that there were strong incentives to exploit new findings in

one's own scientific work, and in exploiting the findings one would either confirm them or find reasons to reject them.

What this loosely coupled science was good at was exploiting opportunities. The decisions on what research to do were made by scientists running labs, typically with a small group of students and technicians. Equipment was important, but relatively cheap and easy to obtain or make. It was possible for industries to sponsor labs which were given relatively large amounts of freedom to work in an area of science, such as textile chemistry lab run by Michael Polanyi in the 1930s, which allowed the research of the lab to benefit industry without being involved directly in the development of new products. What the German chemical industry paid for in this case was world class expertise in physical chemistry as it related to a general area, and this is what they got.

There were other loosely coupled systems which had the same outputs: knowledge that enabled industrial production. Bell Labs was a model of a patient industrial organization. It required a strong sense of the goals of the organization and of the technology that could be improved by new and inventive science. Expectations were clear— scientists were paid to produce results that could be used. But the uses need not be immediate, and experience had shown that research in given areas of physics produced results that could be put to use along with many other results that either would never be used or which might remain curiosities for some time. These kinds of projects filtered out error in their own ways. In the case of Bell Labs, its scientists published and were an active part of the organized community of science itself. But the fact that they were expected to do research that eventually led to practical improvements in telecommunications meant that their results were often also tested in practice by other members of the Bell team.

This kind of loosely coupled structure contrasts with such tightly coupled projects as the Manhattan Project, where there was a high level of technical complexity, clear goals, many knowledge risks, short time frames— in this case the pressure of war. This project succeeded. But the German and Japanese projects, which had the same goal, failed. In the German case, the failure was a result of mistaken ideas about the physics. And it is worth recalling that Oppenheimer opposed the development of the H-bomb on mistaken grounds related to his own beliefs about the relevant physics (Thorpe 2006, 210), which were especially important because the knowledge in question was subject to secrecy restrictions that prevented it from being filtered for error in the classical way. This is suggestive: tightly coupled systems of knowledge production generate special kinds of knowledge risks, risks that arise when the filtering mechanisms change, that are analogous to the risks Perrow identifies for tightly coupled production systems.

Very Tightly Coupled Expert Systems

In the 1960s a famous experiment, known as the Homestake experiment, was designed with the aim of empirically testing some basic physical ideas about solar neutrinos. The experiment began in 1970. The best physics of the time predicted that solar neutrinos would pass through the earth at a particular rate. But no one had ever tested this by seeing what happened inside the earth. The Homestake experiment was an attempt to do so. The experiment was conducted at the bottom of an abandoned gold mine in South Dakota. The design was itself complex. Solar neutrinos are difficult to detect. The method was to build a huge vat and fill it with perchloroethylene, a chemical used, as it happens, in dry cleaning. The chlorine atoms were the means of detection:

when a solar neutrino collided with one, it produced a radioactive isotope of argon, which could be extracted by bubbling helium through the vat, and used to estimate the number of neutrinos that had been captured.

It was widely believed at the time that this would be a crucial experiment— a test not only of ideas about solar neutrinos, but of the image of the world provided by particle physics. The results, however, were incomprehensible. The vat did detect the argon by-product of the collisions. But it detected it at a rate far below what was predicted. Something was right: neutrinos were detected. But something had also gone seriously wrong. But what? The experiment depended on a reasonably well understood technology of collection, with an understood physical chemistry. The estimates were based on this stable knowledge— and the assumption that the knowledge risks were known and controlled. The model of the sun that was being tested was also well supported.

The Homestake experiment, like the CERN accident, involved a complex instrument, and elements which depended on a variety of disciplines or subdisciplines. The experiment thus involves both an epistemic problem and a social problem of organizing the decision-making involved in the construction of the experiment and its interpretation. The “social” problem is also epistemic: a problem of the aggregation of knowledge.³ The knowledge relevant for the experiment was distributed among many people, with specialist knowledge and perhaps tacit knowledge that was not possessed by other members of the team working on the experiment. Moreover, and not surprisingly, critics of the experiment from different disciplines had different interpretations of what had gone wrong. Some even claimed, at first, that the discrepancy was not as significant as it appeared.

What brought these scientists together was the experiment, not common membership in a discipline. They needed to manage the fact that the experiment as a whole depended on a division of labor that was also a division of expertise. This was a relationship that differed radically from the relationship between competitors in a discipline. They were not, accordingly, motivated to refute one another's core scientific ideas, much less the work involved in the experiment— they were neither competent to do this, nor had any reason to, until the experiment failed. The enterprise they were engaged in was a co-operative one. It required some level of mutual understanding and trust. On the other hand, they were strongly motivated, both internally and externally, to defend the aspect of the experiment that was part of their expertise and for which they were responsible. The disciplinary knowledge that went into the experiment, from multiple disciplines, was entrenched. Raymond A. Davis, the experimenter, later recalled it this way: “My opinion in the early years was that something was wrong with the standard solar model; many physicists thought there was something wrong with my experiment” (2002). Indeed they did. They focused, as Perrow's model of normal accidents suggests they would, on the aspects of the experiment that routinely go wrong in experiments: possible errors in calculations and errors in the experiment itself. But, as Davis later said, they were unable to find the error. “Years of measurements produced consistent answers and many tests showed that there were no problems with experimental procedures” (2002). The experiment continued, with the same results— detecting only close to a third of the predicted results. Data collection ended almost a quarter century after it began, with no solution.

Trevor Pinch's history of this controversy recounts the stage of confusion, and it is characteristic of the kinds of accidents Perrow describes. Pinch quotes a summary of the 1972 Irvine Solar Nuetriono conference: “the critical problem is to determine whether the discrepancy

is due to faulty astronomy, faulty physics, or faulty chemistry” (1986, 157). The four domains he outlines made a different proposal: that the argon was trapped in the tank by some unknown process, causing the detector to fail to collect all of the argon produced by the interactions with neutrinos (a problem of radiochemistry), that the nuclear cross-sections were based on extrapolations from high to low energy situations that may have been unwarranted (a problem of nuclear physics), the possibility that the neutrinos oscillate between three phases, so that only one phase is detected (a problem of neutrino physics), and the possibility that the flow of solar neutrinos is not steady because of mixing of outputs from the sun, which might produce long term anomalies (a problem of astrophysics) (Pinch 1986, 157-9).

Pinch’s interpretation of the controversy revolved around such things as Davis’s credibility and suggested that this was more important than the results themselves. His credibility was based largely on the fact that Davis took the criticisms seriously, and carefully tested and retested the processes himself in response to criticisms and suggestions (Pinch 1986, 175). But the crux of Pinch’s interpretation involved the notion of interpretive flexibility, and the idea that there were resources in the theoretical background to the experiment that could be exploited to account for the results without getting rid of anything basic in the physics or the solar neutrino model itself. But his respondents put limits on this flexibility, by rejecting some interpretations as *ad hoc* because they were *post facto* interpretations designed to explain away the experiment rather than being well-grounded in actual physics. Interestingly, though Pinch dismissed the uses of *ad hoc* as inconsistent, and just another sign of interpretive flexibility, one of the hypotheses that was not treated as *ad hoc* involved the issue of oscillations (Pinch 1986, 189), and indeed, as Pinch notes, “if we assume the evidential context of the experiment is neutrino physics then ideas such as neutrino oscillations seem to follow rather naturally (Pinch 1986, 194-5).

When Pinch wrote it made some sense to treat the “consensus” that there was a serious discrepancy as a “social accomplishment,” to treat interpretive flexibility as an irreducible fact of science, and so forth. In retrospect, however, his account, and the disagreements and speculations he described, were part of the stage of confusion itself. The problem with the experiment was explained only later, in a 2001 paper (Schlattel). The paper reported results with a new kind of detector, which was able to detect neutrino oscillation. The paper explained that only one of the three “flavors” that the neutrinos oscillated between could be detected by the methods in the experiment. This was the third that the results had consistently reported. The experiment itself had been sound. It lived up to the expectation that its results would be of fundamental significance, and Davis’s hunch was correct. The significance turned out to be that the standard model of particle physics needed to change: neutrinos had been thought to be without mass and to be of the same type when they left the sun and when they arrived on earth. It was eventually solved by recognizing that they have mass and therefore can change type or oscillate, and did so in the course of traveling from the sun.

The stubbornness of this problem was a result of the interaction between two errors— the error of thinking that the detector would capture all of the neutrinos and the mistaken belief, deeply grounded in the standard model of massless neutrinos, that they did not oscillate. The two elements—the theory underlying the collection of data and the model that was being indirectly tested but was presupposed by the measurement procedures—combined to produce the incomprehensible results. Neither could be directly tested by the instrument Davis used. There was nothing “wrong” with the calculations, the chemistry, and the procedures themselves. It is normally preferable to have a theory of the experiment that is independent of the theory that is being tested, but this is not always possible. In this case, the solution required a technological

deus ex machina.

The story of the Homestake experiment is comforting, if we look only at the beginning and the end. If we ignore the long period of confusion, the outcome conforms to the classical model. The problem was taken up by other labs. Finally a detector was invented that could detect the oscillating neutrino. There were career incentives for solving this problem, conjectures and refutations, and an ending which affirmed the significance of the original experiment— a Nobel Prize for Davis, the experimenter. The self-correcting machinery of science worked, though it took a long time to do so. If we look at the three decades of incomprehension, however, the picture is less comforting. The consensus of the scientific community was wrong in its initial reaction, and wrong for decades. As Davis later recalled, “Many distinguished physicists suggested explanations for the low solar neutrino flux that now seem fanciful” (2002). But these explanations were, for many years, treated as sufficient grounds to dismiss the experiment. The reaction of the ‘scientific community’ is telling. The Homestake experiment was one of the highest profile and most anticipated experiments in the history of science. Davis was no Mendel toiling in obscurity. In the absence of a routine explanation for the results, the scientific community chose more or less to ignore them, and to accept as plausible explanations that were largely groundless, but conserved the standard model of the sun and the particle physics behind it. Rather than treating the results as a major anomaly that needed to be addressed and accounted for, the community lost interest in the results.

The correction process prevailed. But the circumstances were unusual. Despite Davis’s suspicions of the standard model, the hope of the funding agency was that the result would confirm the standard model. Yet when it did not, the experiment was not shut down. The anomalous results were important enough to continue to secure funding for the project for

decades, perhaps in the hope that the experiment would correct itself. Davis, by his own account, was strongly motivated by his suspicion that the standard model was wrong. But if this had been the stated motive of the research, it would likely not have been funded. International competition was important as well. Scientists outside the US continued to work on the problem, and had different funding sources with less commitment to the consensus point of view and more willingness to take a chance on a topic that the experiment had opened up to new thinking. This was exceptional. It is difficult to get funding for a twenty-five year long experimental run with a complex and expensive design. Nor is funding readily available for refutations of the standard model. In this case, the funders probably were motivated by a desire to confirm the standard model, and to continue funding in order to account for the anomaly. In other cases where costs are high, such as time on expensive telescope systems, it has been claimed that only projects that fit with standard models get access. We may think of this loop, in which funding decisions restrict to the standard models the kinds of science that is done, which in turn justifies the funding decisions, as a way of tightening a loosely coupled system. But the effect of tighter coupling, as we have seen, is to produce a particular kind of knowledge risk— risk of failures that fall outside of those that can be corrected by standard operating procedures or technical solutions.

Policy-Oriented Science as a Tightly Coupled System

One of the central elements of tightly coupled systems is time: there are short lags between events in the system. As a consequence, there is little time to make corrections other than those provided by routines, back-up systems, and standard operating procedures for error correction.

The reason the classical model of science worked so well in correcting errors was that it had time—time for conjectures to be refuted, for alternative ideas to develop and get support from new findings. The classic model also supplied a motivation to do these things, and sufficient independence between sources of support to allow scientists the freedom to follow their hunches about what was wrong.

The anomalous finding was, from the point of view of the field, an accident. The accident was detected and recognized as significant in part because the experiment was conducted in a tightly coupled area of science, in which precise predictions were not only possible but normal. Policy-oriented science, and one cannot help but think of climate science as an example of this, works very differently. It is a feature of this kind of science that it is time-limited, usually by the demands of the decisions that need to be made or by the desires of the participating scientists to have an impact on policy. It is also commonly the case that policy-oriented science is done by an artificially created community of scientists who share particular policy concerns, who make funding decisions about one another's work, and are bound up with one another in many other ways. Thus policy-oriented science truncates the processes of science in order to produce a needed consensus in a shortened time frame, in a situation where there are known to be gaps between the evidence and the decisions, policies, and practices they are supposed to guide.

The issues with these systems can best be understood by examining the institutions involved, which operate in ways that are different than “normal” science itself, or regulatory science. The institutions are familiar: in the United Kingdom, there are Royal Commissions. In the United States, the National Academy of Sciences was created in 1863 to perform this service, and produces a steady stream of reports in response to governmental requests as well as on its own initiative. The United States National Institutes of Health has a “Consensus Development

Program” which holds major conferences with the key researchers in an area on nearly a monthly basis during the academic year on such topics as “Vaginal Birth after Caesarean” and “Enhancing the Quality of Colorectal Cancer Screening.” The slogan of one of the programs for outreach sponsored by this program is “Mind the Gap: Bridging the gap between evidence and practice” The website shows a white-coated medic stepping into a blank space— perhaps a street, perhaps the void (NIH 2009). The phrase and the image capture the issues nicely: the generic problem to which they are a response is the fact that the evidence is not enough to guide practice unequivocally, and that some sort of fraught step, even a leap, is needed to get from the evidence to practice. The reason for these institutions is this: the science alone— the “evidence,” in this phrase— is insufficient, conflicting, or simply not in a form that can guide practice. But practice is necessary, ongoing, and consequential. So the process of bridging the gap needs to be helped along.

In one sense, these institutions can be treated as marginal to science itself. They are, in the famous phrase of Leigh Star and James Griesemer (1989), boundary organizations, engaged in producing boundary objects, in the form of such things as commission reports that are usable for decision-makers or policy. But if we consider science as a knowledge production system that can be more or less tightly coupled, the role of these institutions becomes more relevant. Every major science funding system is concerned, increasingly, with the impact of research. Most of these systems now involve a) large scale funding initiatives targeted at the production of knowledge relevant to politically important knowledge domains, and b) specific requirements for describing and identifying impacts, including public impacts, in the text of proposals themselves. These are coupling devices: they assure that the research is coupled more closely to the rest of the science system, as well as to public goals. The public goals in question are the same kinds of

goals, and usually the exact same goals, that these institutions— consensus conferences and commissions— are concerned with. Coupling research to impact is not voluntary, though the specific impacts are selected by the applicant. The effect of these practices, and many other related practices, is to create a more tightly coupled knowledge production system than the classic model. But to the extent that the system becomes more tightly coupled it becomes susceptible to risks, including knowledge risks, which are characteristic of tightly coupled systems.

New Kinds of Consensus

To be a policy research community is to promise some sort of expert knowledge about the policy problem, rather than mere advocacy of a policy supported by data. But to get to this goal in a timely manner, one which does not wait on the long process of consensus formation required by the classical, loosely coupled model, requires truncating the process of consensus formation. This is precisely the purpose of consensus seeking institutions, such as consensus conferences and commissions, and the goal of targeted funding systems. But there is a price to be paid in terms of knowledge risk.

Climate science is the field that today is the most visible example of a policy-oriented science field that is tightly coupled. The history of the field provided by Spencer Weart (2003), which is a defense of the scientific credentials of the field and its legitimacy as science, bears this out. Weart describes the slow development of the subject under the conditions of loosely coupled science. In this phase, money to pursue key ideas, such as the study of the accumulation of CO₂ in the atmosphere, was scarce, and the interested scientists were able to do research on these

topics only more or less by accident: the generous military funding of science in the 1950s, for example, provided some key elements (Weart 2003, 31). Geophysical research generally is “platform” science: it requires big instruments which in turn required extensive lobbying by researchers to acquire (Weart 2003, 40). Only when public fears were aroused (Weart 2003, 401) and the conservation movement picked up on the problem of global warming (Weart 2003, 44) did the money begin to flow freely.

Climate modeling is at the core of the study of global warming. And like CERN and the Homestake experiment, the materials that go into models come from a variety of scientific fields. The climate prediction outputs of these models were, not surprisingly, sensitive to assumptions, parameters, small variations inputs, and so forth. Weart recounts a long history of failure to make even plausible predictions, such as models that predicted that a doubling of CO₂ would produce a rise in global temperatures of 10c or higher, as well as the fiasco of the predicted new ice age. Weart concedes that these “rudimentary” models were wildly wrong, though they were presented as quite serious policy relevant predictions at the time, by scientists such as Stephen Schneider, who are still active and still aggressively promoting the policy relevance of their current views (Weart 2003, 83). Weart’s openly Whiggish perspective on these models was that “most research begins with flawed theories, which prompt people to make better ones” (Weart 2003, 84). He treats the fact that “a simple calculation (whatever problems it might have in detail) could produce a catastrophic outcome” (albeit on paper) as a “disturbing discovery” (Weart 2003, 84).

Weart makes the same points about the intertwined social and epistemic complexity of the science: communication between the relevant fields was a problem (Weart 2003, 125, 149). Thus aggregating climate-relevant knowledge was a problem, a problem which was not matched by the controlling social structure of a discipline, as was the case with normal science, but done

under a different structure. Climate itself, it was clear from the outset “was a staggeringly intricate complex of interactions and feedbacks among many global forces” and that even relatively simple phenomena, such as the doldrums in the tropical seas, defied explanation (Weart 2003, 111). He concedes that this field does not operate as a normal scientific discipline does:

the social structure is not cohesive. A community in one specialty cannot check the work of researchers in another branch of science, but must accept their word for what is valid. The study of climate change is an extreme example. Researchers cannot isolate meteorology from solar physics, pollution studies from computer science, oceanography from glacier-ice chemistry, and so forth. The range of journals they cite ... is remarkably broad. This sprawl is inevitable . . . But the complexity imposes difficulty on those who try to reach solid conclusions about climate change. (Weart 2003, 195)

He acknowledges that the models rely on *ad hoc* devices. They fit the actual climate, even roughly, “only because they had been laboriously ‘tuned’ to match it, by adjusting a variety of arbitrary parameters” (Weart 2003, 177). He also acknowledges that generating money for the field was intertwined with politics throughout: self-interest, emotional commitment, and the political cause of promoting action on the problem coincided (Weart 2003, 151).

These abnormalities (compared to the classical model of disciplinary science) are, for Weart, justifiable on the grounds of time: “in the case of climate, waiting for a sure answer would mean waiting forever. When we are faced with a new disease or an armed invasion, we do not put off decisions until more research is done: we act using the best guidelines available” (Weart 2003, 198-9). This is the language of the NIH consensus conferences, and of

commissions. In the absence of ordinary disciplinary controls, we must accept the claims of climate science on other grounds. His own judgment of the correctness of the dominant views in climate science rests on his “feeling for where scientific claims are reliable and when they are shaky” (Weart 2003, 199), his belief that

The few who contest [the] facts are either ignorant or so committed to their viewpoint that they will seize on any excuse to deny the danger. (Weart 2003, 199)

Accordingly, he closes the book with his own policy recommendations, and the following final sentence:

The spirit of fact-gathering, rational discussion, tolerance of dissent, and negotiation of an evolving consensus, which has characterized the climate science community, can serve well as a model. (Weart 2003, 201)

In short, it is these features of the climate science community that need to be replicated in the domain of policy, and at the same time provide an alternative ground for trust in the conclusions, distinct from the control of disciplinary “social structure.”

This characterization of the climate science community has now been shown to be false in every respect as a result of the release of e-mails from the University of East Anglia’s Climatic Research Unit (an institution Weart praises). It is now evident that facts were mishandled and suppressed, data withheld and lost, peer-review manipulated, a major effort made to discredit those who questioned the “consensus,” and more: this group is now an example

of what Polanyi called the corruption of science. The contrast between these e-mails and the period when this work was done by disciplinary science is striking. Thirty years ago I was given access to the editorial files of the Oceans and Atmospheres series of the *Journal for Geophysical Research*. The Atmosphere series, which is now separate, now advertizes itself as one of the top ten most highly cited research publications on climate change over the past decade. Its sister journal *Geophysical Research Letters* is also among the top ten. At the time, the journal was a model of open discussion. Rejections were rare, although around thirty percent of the submissions were rejected, most of the rejections were because the papers were inappropriate (for example, engineering papers about platforms for science rather than science), or because the English was inadequate. The editor of the journal had a strong sense of serving a community, and because this was platform science, it was a well-defined community: the people who worked in or had access to the satellites and research vessels, and sensors that were required to do the research. Novelty came largely from novel platforms: new satellites or technology. Tolerance was high: a paper could be accepted even if it was speculative, if the editor believed someone in the community might find it useful. Because this was platform dependent science there was a degree of pre-screening: these projects had already gone through evaluations as grants or as applications to use the technology. But the field was well-funded, and there were multiple agencies in different countries supporting research, with no shared political agenda. The goal of the journal was to serve the community, and it erred on the side of inclusion. There were no instances of the kinds of peer-review manipulation discussed in the East Anglia e-mails, and virtually no negative reviews. “Consensus” under these conditions was not something that needed to be asserted or demonstrated. To the extent that consensus was an issue, it was the product of a loosely coupled system that minimized interventions designed to produce

consensus.

Two Kinds of Knowledge Risk

Perrow argues that the people who operate complex systems think in the following way: they make initial choices about how to think about the accident situation in the face of ambiguous data. If the first steps they take are consistent with the choices they have made about how to think about the data, their mental model of the situation, they psychologically commit to the model, and lose the skepticism they might originally have had. New data that are inconsistent with the model do not become reasons for reconsidering the model. Instead, they are seen as mysterious, unilluminating, or incomprehensible. The data are discounted. The similarity between this response and Kuhn's discussion of paradigms and the response of "normal science" to anomalies is striking. And it is a reminder that the problem of knowledge risk is not confined to tightly coupled systems of the kind discussed here. Disciplines, even the open structures defined by the classical model, produce intellectual conformity, which creates distinctive knowledge risks. The negative consequences of this conformity were more important to Kuhn's mentor James Bryant Conant, who noted the limitations of scientists: their tendency to group-think, their proneness to hobby-horses, and their tendency to over-rate the importance of their own disciplinary concerns (Conant 1952, 110-17, 176).

The traditional disciplinary model was fading even as Conant wrote, in the late 1940s and early 1950s. The Homestake experiment itself was an example of what came to be called "Big Science." Science studies scholars have struggled to make sense of the new forms of science that have emerged since then, and especially in the last thirty years: terms like 'technoscience',

‘Mode II’ (Gibbons et al 1994), ‘post-academic’ science (Ziman 1996) have been invented to describe these new forms. Tight-coupling is a feature of this new mode of science. To some extent, the knowledge risks of disciplinary science and the knowledge risks of tightly coupled science overlap. But the differences do matter. Unlike disciplines, the new forms typically involve a complex division of intellectual labor involving multiple specialties and a “social structure,” in Weart’s sense, that serves to aggregate knowledge from these specialties to produce expert conclusions. The products of this kind of science embody this complexity. But complexity itself produces unexpected interaction effects that belong to no discipline. The expertise in question is about things that are not under any disciplinary domain. Consequently, the structures involve mechanisms of consensus production and control that differ from, and involve more overt use of authority, negotiation, and pidgin communication between specialties than in the model of disciplines with paradigms. The problem of aggregating multidisciplinary knowledge is not a problem that can be solved in any disciplinary way. The pressure to produce results leads to the *de facto* elimination or moderation of some of the means of correction provided by the traditional disciplinary model— what Robert Merton called organized skepticism is less effective when the objects of the skepticism are not disciplinarily defined. These new structures have many virtues. They produce results that can guide policy and innovation and produce objects that disciplinary science could not. But there is no epistemic free lunch. Each of these differences creates knowledge risk.

References

Bajko, M. F. Bertinelli, N. Catalan Lasheras, S. Claudet, P. Cruikshank, K. Dahlerup-Petersen,

R. Denz, P. Fessia, C. Garion, J.M. Jimenez, G. Kirby, M. Koratzinos, Ph. Lebrun (chair), S. Le Naour, K.H. Mess, M. Modena, V. Montabonnet, R. Nunes, V. Parma, A. Perin, G. de Rijk (scientific secretary), A. Rijllart, L. Rossi, R. Schmidt, A. Siemko, P. Strubin, L. Tavian, H. Thiesen, J.Ph. Tock, E. Todesco, R. Veness, A. Verweij, L. Walckiers, R. van Weelderen, R. Wolf CERN, Geneva, Switzerland; S. Feher, R. Flora, P. Limon, and J. Strait, FNAL, Batavia, USA (2009) Report of the Task Force on the Incident of 19 September 2008 at the LHC. *European Organization for Nuclear Research* (31 March).

<http://cdsweb.cern.ch/record/1168025/files/LHC-PROJECT-REPORT-1168.pdf>

Chalmers, Matthew and Mark Henderson (2008) CERN Delays Atom-Smashing over Magnet Fault. *The Times* (20 September).

<http://www.timesonline.co.uk/tol/news/science/article4789673.ece> accessed June 30, 2010

CERN (2008) Follow Up of the Incident of 19 September at the LHC. Press release Geneva, 5 December 2008. <http://press.web.cern.ch/press/PressReleases/Releases2008/PR17.08E.html> (accessed June 30, 2010)

Conant, James B. (1952) *Modern Science and Modern Man*. Garden City, NY: Doubleday & Co.

Davis, Raymond A. (2002) Autobiography. *Nobelprize.org*

http://nobelprize.org/nobel_prizes/physics/laureates/2002/davis-autobio.html (accessed June 29, 2010)

Gibbons, Michael; Camille Limoges, Helga Nowotny, Simon Schwartzman, Peter Scott, & Martin Trow (1994) *The New Production of Knowledge: The Dynamics of Science and Research in Contemporary Societies*. London: Sage.

Henderson, Mark (2009) Large Hadron Collider to Run through the year in Higgs Boson Race. *The Times* (5 June). <http://www.timesonline.co.uk/tol/news/science/article6431714.ece>

(Accessed June 30, 2010)

Hill, Josh (2008) CERN Shutdown: Predictable? *The Daily Galaxy* (29 September)

<http://www.physorg.com/news141814025.html> accessed June 29, 2010

Latour, Bruno (2004) Why Has Critique Run out of Steam? From Matters of Fact to Matters of Concern. *Critical Inquiry* 30(2): 225-248.

Laudan, Larry (1978) *Progress and Its Problems: Towards a Theory of Scientific Growth*.

Berkeley: University of California Press.

Lynch, Michael (2010) The Checkered Career of Symmetry in Recent Social Studies of Science 50th Anniversary Celebration, Boston Colloquium for Philosophy of Science, Center for Philosophy & History of Science, Boston University (15 April 2010)

National Institutes of Health (2009) Medicine: Mind the Gap; an NIH Seminar Series.

<http://consensus.nih.gov/mindthegap/> (Accessed June 30, 2009)

Perrow, Charles ([1984] 1999) *Normal Accidents: Living with High Risk Systems*. Princeton: Princeton University Press.

Pickering, Andrew (1995) *The Mangle of Practice*. Chicago: The University of Chicago Press.

Pinch, Trevor (1986) *Confronting Nature: The Sociology of Solar-Neutrino Detection*.

Dordrecht: D. Reidel Publishing Company.

Polanyi, Michael [1946] 1964, *Science, Faith and Society*, Chicago: The University of Chicago Press.

Popper, Karl ([1963] 1965) *Conjectures and Refutations: The Growth of Scientific Knowledge*. New York: Harper & Row.

Rossi, Lucio (2010) Superconductivity: Its role, its success and its setbacks in the Large Hadron Collider of CERN. *Superconductor Science and Technology* 23(3): 1-17.

<http://iopscience.iop.org/0953-2048/23/3/034001> Accessed June 29, 2010

Rouse, Joseph (2002) *How Scientific Practices Matter: Reclaiming Philosophical Naturalism*. Chicago: The University of Chicago Press.

Schlattl, H. (2001) Three-flavor Oscillation Solutions for the Solar Neutrino Problem. *Physical Review D*, 64, 013009 (June) <http://prola.aps.org/abstract/PRD/v64/i1/e013009> (Accessed June 29, 2010)

Schrage, Michael (2007) Science Must Be More Political. *The Financial Times*, September 25, Comment: 11.

Shapere, Dudley (1982) The Concept of Observation in Science and Philosophy. *Philosophy of Science* 49 (4):485-525.

Star, Susan L. and James R. Griesemer (1989) "Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39." *Social Studies of Science* 19 (4): 387-420

Taylor, Charles (1964) *The Explanation of Behavior*. New York: Humanities.

Thorpe, Charles (2006) *Oppenheimer: The Tragic Intellect*. Chicago: The University of Chicago Press.

Times, The (2009) Large Hadron Collider to Run through the Year in Higgs Boson Race. Science News, June 5, online. <http://www.timesonline.co.uk/tol/news/uk/science/article6431714.ece> (Accessed June 29, 2010)

Turner, Stephen. 1999. Practice in Real Time. *Studies in the History and Philosophy of Science* 30: 149-56.

Turner, Stephen. 2005. Expertise and Political Responsibility: The Columbia Shuttle Catastrophe. *Democratization of Expertise? Exploring Novel Forms of Scientific Advice In*

Political decision-Making, edited by Sabine Maasen and Peter Weingart. Dordrecht: Kluwer, pp. 101-21.

Turner, Stephen (2007) Political Epistemology, Expertise, and the Aggregation of Knowledge. *Spontaneous Generations: A Journal for the History and Philosophy of Science* 1(1): 36-47.

Turner, Stephen (2010) *Explaining the Normative*. Oxford: Polity Press.

Turner, Stephen (forthcoming 2010) Webs of Belief, Practices, and Concepts. *Archives européennes de sociologie/European Journal of Sociology*.

Weart, Spencer (2003) *The Discovery of Global Warming*. Cambridge, MA: Harvard University Press.

Ziman, John (1978) *Reliable Knowledge: An Exploration of the Grounds for Belief in Science*. Cambridge, New York: Cambridge University Press.

Zuckerman, H. & Merton, R. K. (1971) "Patterns of Evaluation in Science: Institutionalization, Structure and Function of the Referee System." *Minerva* 9 (January): 66–100.

Notes

¹ Philosophers pointed out a large problem with this account: according to Taylor, intentional explanations were teleological and asymmetric. The intended end was the favored end of an intentional explanation. So there was a question of whether any account that relied, as the "Strong Programme" in the Sociology of Scientific Knowledge did, on ordinary intentional language, could ever be symmetric in the relevant sense (Turner forthcoming 2010). This discussion was framed in the now passé language of the problem of reasons and causes. But it has now reemerged in the form of the concept of "normativity," where it has the same

implications. The term “belief” itself, for example, has been claimed to be a “normative” concept that cannot be understood “causally” (Rouse 2002, 99).

² The issues with this device are discussed in “Webs of Belief, Practices, and Concepts” (Turner forthcoming 2010).

³ I have discussed the problem of aggregation of knowledge and the issues with the social organization of aggregation elsewhere. See especially “Political Epistemology, Expertise, and the Aggregation of Knowledge” (Turner 2005, 2007).