Downer, J. (2019). On Ignorance and Apocalypse: A Brief Introduction To 'Epistemic Accidents'. In J-C. LeCoze (Ed.), *Safety Science Research: Evolution, Challenges and New Directions* (1 ed.). CRC Press. <https://doi.org/10.4324/9781351190237-5>

Early version, also known as pre-print

Link to published version (if available):
[10.4324/9781351190237-5](https://doi.org/10.4324/9781351190237-5)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the submitted manuscript (SM). The Version of Record (VOR) is available online via Taylor and Francis at <https://www.taylorfrancis.com/books/e/9781351190237/chapters/10.4324/9781351190237-5>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

On Ignorance and Apocalypse: A Brief Introduction To ‘Epistemic Accidents’

For LeCoze, J.C. (ed.) “New Perspective on Safety Research” (forthcoming 2018)
Pre-edited draft. Please do not cite.

John Downer¹

I thought her unsinkable and I based my opinion on the best expert advice available. I do not understand it.
~Philip A. S. Franklin, Vice President of White Star Line, on the sinking of the Titanic.

After the ship has sunk, everyone knows how she might have been saved.
~Italian proverb

A Close Call?

In the late 1990s a German biotech company altered a common bacterium — *Klebsiella planticola* (or *K-planticola*),² which plays a role in decomposition — by adding a gene that would make it turn plant material into alcohol. The newly modified bacterium, which they called ‘SDF20,’ was intended for farmers. It would allow them to take waste they would usually burn and convert it into ethanol they could sell, plus a nutrient-rich ‘sludge’ they could use as fertilizer.

An organism that simultaneously reduced burning and produced two valuable resources seemed like a win-win, but SDF20’s appeal waned sharply after a study conducted at the university of Oregon (Holmes et al. 1999). The Oregon researchers had noticed that the leftover sludge contained live SDF20, and decided to explore the bacterium’s effects on the ecosystem. In a series of experiments they introduced the sludge to soil samples containing wheat plants. Within two weeks, the plants had died. The live SDF20 had interacted with micro-organisms in the soil (its ‘biota’) in ways that were ultimately fatal to the crops it supported.

¹ Global Insecurities Centre, School of Sociology, Politics and International Studies, University of Bristol. email: john.Downer@bristol.ac.uk

² *Klebsiella planticola* has since been reclassified as *Raoultella planticola*.

The study and its implications remained relatively obscure and uncontroversial until 2001, when the lead researcher, Elaine Ingham, invoked it while testifying about the potential hazards of GMOs to the New Zealand Royal Commission on Genetic Engineering. She claimed the US Environmental Protection Agency (EPA) had missed the danger SDF20 posed because its initial tests were conducted with sterile (biota-less) soils, and, as a result, had approved it for field trials, which were only cancelled in the wake of her study. Had the field-trials gone ahead, she suggested, SDF20 might easily have spread worldwide: displacing the ubiquitous natural *K-planticola*, damaging soils and decimating crops around the globe (Porterfield 2016; Walter et al. 2001).

Ingham's testimony quickly became embroiled in the acrimonious politics around GMOs. Environmental advocates invoked it as evidence of a narrowly averted apocalypse; painting dire portraits of a runaway, laboratory-born organism racing across the planet, leaving wastelands and famines in its wake (e.g. Robbins 2002; Brockway 2010). Proponents of GMOs responded robustly. In a rebuttal to Ingham's original testimony, a group of scientists denied that the EPA had formally approved field tests; highlighted an erroneous citation in her paper, and claimed her testimony went beyond the published evidence. Others built on this testimony to discredit Ingham herself: casting the erroneous citation as a falsified source, and a routine departmental review as an investigation of her integrity (Walter et al. 2001; Krebs 2001). Such efforts were well-funded and rhetorically effective. Ingham retracted her claim about formal EPA approval, corrected the erroneous citation, and allowed that her fears went beyond her direct observations. None of these concessions amounted to an admission of wrongdoing or substantial error — scientists routinely speak to matters beyond their direct observations — yet the optics were bad, and the alarm is now remembered by many as an example of environmentalist hysteria.

As the dust settles on the controversy, however, it should be clear that Ingham's fears were far from ungrounded. Her research strongly indicated that SDF20 might have endangered plant life had it spread. And given the ubiquity of natural *K-planticola*, it was reasonable to imagine

that it might have spread uncontrollably if released into the wild, as other organisms have in the past. It might have been true, moreover, that the EPA had not issued formal approvals for field tests prior to her study, yet it is easy to imagine that it could have done so, and would have done so in time. At the time of the incident the agency's standard GMO assessment practices paid little account to the wider ecosystem; framed, as they were, by assumptions better suited to chemicals than to living organisms. Its initial lab tests with sterile soil would not have highlighted dangers arising from SDF20's interaction with soil biota. Nor would other tests mandated by the Toxic Substances Control Act or the Federal Insecticide, Fungicide, Rodenticide Act. Reasonable people can disagree about how close the world came to a disaster with SDF20, but it would be foolish to dismiss Ingham's concerns as baseless alarmism.

Somewhat troublingly, in other words, it is at least *plausible* that the US Environmental Protection Agency, in conducting its routine safety assessments, might have unwittingly unleashed the biological apocalypse. Perhaps even more troublingly, it is difficult to imagine who would have been at fault in this scenario. The scientists involved would have been following proper protocols; performing all the mandated tests prior to testing SDF20 in the wild. And while it is true that those tests, with their sterile soil, were unrepresentative of the real world in consequential ways, it is also true that the true representativeness of any test (much like that of any experiment, simulation or model) is always an inherently imperfect and unknowable property (see e.g. Pinch 1993; Mackenzie 1996; Downer 2007).

The Limits of Interrogation

The value of tests — whether they be of a bacterium, a reactor, or even a jet engine — lies in their ability to *represent* the 'real world' while simultaneously *differing* from it in key respects. In some cases this means stripping away facets of the world that might otherwise occlude the phenomena being investigated. Tests designed to compare the tensile strength of metals under pressure, for example, are maximally reductive. They seek to isolate those two properties —

tensile strength and pressure — by strictly controlling every other significant variable (moisture; purity; temperature; manufacturing imperfections, etc) that might skew the result and keep it from being comparable across samples. In other cases the phenomenon under investigation relates directly to performance in the real world, and here tests must be maximally representative. In these circumstances testers seek to reproduce every facet of the world that might affect the result, while simultaneously retaining the differences that give the test value. Most ‘safety’ tests fit into this category. They seek to recreate the world in every way except those that create actual hazards. (So flight tests are conducted without passengers, for example). In every case, however, the quality of a test is a function of testers’ ability to identify and control significant variables: either to include or exclude them, depending on the context. A test for measuring the strength of a metal is meaningless if it doesn’t control the metal’s exposure to the world; and a test for measuring the safety of a GMO is meaningless if it doesn’t control the world’s exposure to GMOs.

Identifying and controlling significant variables is key to ensuring a test’s validity, therefore, but how are experts to ensure they have identified and controlled every significant variable? Philosophers sometimes call this the ‘problem of relevance.’ It is implicated in every technological test, scientific experiment, theoretical model, and computer simulation. And it has no satisfying solution.

The essence of the problem is that the number of factors that might be significant to a test are potentially infinite, while the number of factors that testers can control is inherently finite.³ Hence, every test, by its nature, might be unrepresentative in a significant way that the tester has not considered. This has complex ramifications; not least because it implies that there is a fundamental circularity to our interrogations. We use tests (experiments, models, etc.) to examine the validity of our theories about the world (whether they be about the safety of a GMO, or the strength of a metal). Yet every test also embodies its own theories about the world, in the sense that it reflects assumptions about the factors it must represent and how

³ Hence the term ‘finitism,’ which describes the philosophy that builds on this insight.

well it has represented them. But if testing theories always means relying on other theories (i.e. about the representativeness of the test itself), then there can be no real certainty. For every finding can be questioned by disputing the bases on which it was observed. A test that ‘proves’ the safety of a bacterium, in other words, could always be an imperfect test.⁴

Dilemmas like those outlined above are at the crux of modern epistemology. For hundreds of years, great philosophical minds wrestled with the ‘scientific method’: dissecting the logics of induction and deduction in an effort to formulate an ineluctable route to definitive facts. None succeeded. Instead, the great 20th Century breakthroughs came from those who embraced the impossibility of the task. In different ways, thinkers such as Kripke (1982); Kuhn (1996 [1962]); Feyerabend (1975); Quine (1951); and Bloor (1976), many of them building on Wittgenstein (2001 [1953]), demonstrated the impossibility of the ‘perfect proof.’ Collectively, they taught scholars to embrace fundamental indeterminacy; encouraging them recognize that evidence is always contingent on prior assumptions, and that facts are always open to revision.

In their wake, a generation of STS (Science and Technology Studies) scholars have built on these ‘finitist’ insights by exploring their material ramifications in contemporary science and engineering. Sociologists such as Collins (1985), Latour (1987), and Mackenzie (1990) have illustrated how the ‘problem of relevance’ creates meaningful uncertainties in contemporary scientific and technological knowledge-claims. Others have explored the implications of this in a wide range of contexts, from the courtroom (e.g. Lynch 1998) to the halls of government (e.g. Jasanoff 1990). Few, however, have closely and systematically examined how an epistemological understanding of uncertainty might contribute to our understanding of accidents.

This neglect is unfortunate, for the inherent uncertainties of tests and models would seem to have direct bearing on the question of why systems fail. Tests are essential to both achieving and measuring technological safety. They underpin the foundational premises (such as

⁴ Formulated slightly differently, in terms of the ‘underdetermination’ of theory by evidence, this is a foundational dilemma in the philosophy of knowledge, sometimes referred to as ‘the Duhem-Quine problem’ or ‘confirmation holism’ (see Quine 1951).

measurements of tensile strength) on which experts design technologies to perform safely; and they simultaneously underpin expert assessments of those technologies' safety performance. If tests of the design are unrepresentative, then assumptions about its behavior could be wrong in ways that have catastrophic potential. And if tests of that behavior are similarly unrepresentative (perhaps for the same reason), then that catastrophic potential might be invisible. Insofar as the representativeness of a test cannot be known for certain, therefore, then a system might fail because an assumption implicit in its design proves to be erroneous, even though there were logically consistent grounds for experts to hold that assumption before (although not after) the event. I have proposed to call such failures 'Epistemic Accidents' (Downer 2011).

Epistemic Accidents

We need not look far to find evidence that epistemological dilemmas play a significant role in accidents. Accident investigations routinely identify causes (or causal factors) in the form of logically-held but nevertheless erroneous assumptions about the design of a system or its assessment. Take, for example, the SDF20 scare outlined above, where the EPA's safety tests were framed by a misleading assumption about the (ir)relevance of soil biota. A more detailed example can be found in Downer (2011), which explores the fate of *Aloha Airlines Flight 243*: a Boeing 737 that suffered a massive mid-flight fuselage failure over Hawaii in 1988. This failure stemmed from a misunderstanding of how aluminum fatigues in highly specific conditions (relating to imperfections in an airplane's manufacturing process, a specific rivet configuration, and the saltwater-infused operating environment of the Pacific islands). The airplane's designers had misunderstood this fatigue behavior because they had never experimented with these exact conditions, and the safety tests they performed for the airplane's regulators did not reveal the misunderstanding for the same reason (Downer 2011).

If we look further afield to some of the more iconic accidents of the last few decades, we again see evidence of epistemological finitism wreaking havoc with experts best-laid plans. A good case could be made, for instance, that the 2011 Fukushima accident was an Epistemic Accident, at least in part. In the most fundamental sense, the reactors failed because the plant's flood defenses were overwhelmed by the tsunami that struck them. Those defenses performed as designed, however, the problem was that their design was inadequate, and it was inadequate because it was premised on an erroneous belief about the maximum possible tsunami. This erroneous belief about tsunamis, in turn, stemmed an erroneous belief about the maximum size of potential earthquakes off the coast — a belief, which, nevertheless reflected the best seismology available at the time the plant was designed.⁵ And the error was not revealed in testing because the same earthquake assumptions that informed the design also informed the models through which regulators assessed the flood defenses. At the time the plant was constructed, no engineering analysis could have identified the circumstances of its demise.

A similar case might be made for the 1986 *Challenger* disaster. In a far-reaching (1996) book on the tragedy, Diane Vaughn lays much of the blame for the disaster on a tendency among NASA engineers to normalize 'deviances' in the behavior of a key component (the "o-rings" sealing the booster rockets) over time. Yet her account is rich enough to support a rival interpretation: specifically, that the meaning of 'deviance' in a complex technical system — where it would be normal for performance to imperfectly match predictions in myriad ways — could not have been transparent to engineers in advance of the failure.

Accidents are complex phenomena, and any specific example of an Epistemic Accident will be contestable, yet the principle stands on its own logic. If we accept the premise that some accidents result from erroneous beliefs about the functioning of a system. And we further

⁵ There is a case to be made that Fukushima's design did not reflect the best seismology available at the time of the accident (Nöggerath et al. 2011). Insofar as we expect operators to reassess, redesign and rebuild nuclear plants in light of (contested) changes in the upstream science on which those plants are premised, therefore, then Fukushima's claim to being an epistemic accident is arguably diminished.

accept that even the most rigorous and well-founded beliefs about the functioning of systems necessarily contain uncertainties. Then it follows that Epistemic Accidents must exist.

New Perspectives: NAT & EAT

What is the value of understanding accidents in these terms? There are several ways to address this question. One straightforward but significant answer is that it shows why some failures are inherently *unpredictable* and therefore *unavoidable*. To recognize the existence of Epistemic Accidents is to acknowledge that no amount of organizational restructuring, application of intelligence, or hard work will ever ‘solve’ the problem of accidents.

Recognizing this limitation is not the same as saying that the prevalence accidents cannot be greatly reduced, or that attempts to reduce them are not valuable and important, but it does have far-reaching implications in certain contexts. Most notably, perhaps, in the context of debates around critical systems, where exceptional hazards are often offset by extreme claims to reliability. Safety debates in most technological spheres accept that entirely failure-free operations is an unreasonable aspiration. In spheres where failures would be most consequential, however, this common-sense understanding has long been suspended. Reactors, perhaps most notably, would probably not be tenable politically unless the possibility of catastrophic accidents were excluded entirely from decisionmaking processes (Downer 2016; Rip 1986).⁶ Here — where accidents can incur costs (direct and indirect) of over a trillion dollars, together with unknown but potentially vast long-term health risks — politics embrace an ideal of technological mastery that is highly at odds with basic epistemology.

The argument that accidents are inherently unavoidable is neither new nor unique, however. Most notably, it has asserted by Perrow (1999 [1984]), and other proponents of his ‘Normal Accident Theory’ (NAT) (e.g. Sagan 1993). The argument outlined above does not seek to

⁶ Or so it is believed by most nuclear authorities. A similar argument might be made for nuclear deterrence networks, although the case here is more complicated (see Sagan 1993).

undermine Perrow's thesis in any way, but rather to compliment it. Epistemic Accidents are not same as Normal Accidents, and, examined closely, the two have meaningfully different ramifications. Indeed, the act of exploring these differences serves as a useful way of examining the implications of Epistemic Accidents. It is to this, then, that we now turn.

NAT has many facets and deserves more unpacking than can be afforded here (see e.g. Le-Coze 2015; Downer 2015; 2011). At its core, however, is a simple but profound insight: that accidents caused by highly improbable 'billion-to-one' confluences of otherwise minor events (which no expert could anticipate in advance) are statistically probable in systems where there are a great many opportunities for them to occur. (Perrow, for instance, points to the extraordinary chain of events that led to the meltdown at Three Mile Island.) If Epistemic Accidents are an emergent property of epistemological limitations, we might say, then Normal Accidents are an emergent property of structure and probability.

Perrow's thesis offers an elegant explanation for why some accidents (albeit only a small minority by his accounting)⁷ are fundamentally unavoidable. The kinds of events that come together to make them are too innocuous in themselves to serve as meaningful warning signs; and the specific sequences that make those events catastrophic are too improbable and variable to ever register as a threat in advance. At the same time, however, there are accidents that would not register as Normal Accidents, but which can nevertheless be understood as unavoidable from the perspective of epistemology. Herein, therefore, lies one specific argument for thinking epistemologically as well as probabilistically. To wit:

i. Epistemic Accidents broaden the spectrum of potentially unavoidable failures.

By Perrow's calculus, accidents that arise from a single point-of-failure are not 'Normal,' and, as such, ought to be foreseeable and avoidable (Perrow 1999: 70-71).⁸ Yet Epistemic Accidents

⁷ Perrow (who more often refers to such accidents as 'system accidents') uses the word 'normal' to connote inevitability rather than commonness.

⁸ He calls them 'Component Failure Accidents'.

can arise in this way: from a single understanding that proves to be erroneous (about the relevance of sterile soil, for instance, or the fatigue properties of aluminum), rather than from an unanticipated confluence of otherwise anticipated events. This does not imply that all non-Normal accidents are unavoidable, of course, but it does imply that there are unavoidable failures that NAT would not recognize.

Once an understanding has been identified as erroneous, however, then there is no reason for it to be unavoidable again. And herein lies a second distinctive feature of Epistemic Accidents relative to Normal Accidents...

ii. Epistemic Accidents imply a distinct relationship to institutional learning and technological achievement.

One important property of Normal Accidents is that they rarely challenge common understandings of the component-level events that cause them. The specific circumstances that come together to instigate a Normal Accident — failed warning lights, stuck valves, human error, and so on — are usually unremarkable in themselves. Normal Accidents are only remarkable by virtue of the unique manner in which these circumstances combine. The investigators of TMI were not surprised to find that valves sometimes stick and warning lights sometimes fail, they were only surprised to find such an improbable combination of such failures.

This odd combination of accident-level uniqueness and component-level banality is important because it leads to a second property of Normal Accidents, this one highly consequential: the fact that they are not very *edifying*. The banality of the events that line-up to create a Normal Accident means that they offer few insights to their investigators, who are unlikely to learn anything new about valves, warning lights or human behavior. Learning from Normal Accidents is difficult, in other words, because they are ‘one-of-a-kind’ events. (For although it is logical to anticipate a fatal ‘one-in-a-billion’ coincidence in systems that allow for billions of

such coincidences, it is not logical to expect the same exact coincidence twice. The exact circumstances that led to TMI might never reoccur, even if the same plant ran for another ten thousand years.) Acting to prevent the recurrence of a sequence that might never reoccur is not useful, therefore, as the next fatal coincidence is vanishingly unlikely to involve the exact same sequence. For these reasons, Normal Accidents usually only teach one important lesson, and it always the same lesson: that unpredictable confluences of seemingly trivial events can instigate disasters.

Epistemic Accidents are very different in these respects. Unlike Normal Accidents, they are caused by specific, consequential events, rather than by random, one-in-a-billion coincidences of events. And, almost by definition, they challenge conventional understandings of those events. They arise from misconstruals, and if those misconstruals are not remedied then they are liable to reoccur. EPA testing regimens that continued to mandate sterile soil, for example, would pose a hazard to soil biota every time the agency tested a new bacterium.

For the same reasons, however, Epistemic Accidents can be edifying. The EPA's close encounter with SDF20 offered insights that would allow it to change its tests. And the fact that other bacteria could pose similar risks in the future meant that those changes could contribute meaningfully to the safety of its work.⁹ In other words, the fact that Epistemic Accidents reveal shortcomings in our knowledge — ways in which our tests and models are unrepresentative — means that experts can leverage hindsight to improve systems and practices over time. If a future bacterium ever decimates the world's soils because regulators failed to test its effects on soil biota, the resulting catastrophe would not be an Epistemic Accident.

This property of Epistemic Accidents is important because it offers a unique perspective on innovation and safety. Elsewhere (Downer 2017), for example, I have argued that one of the ways that modern civil aviation has achieved such remarkable levels of safety is by leveraging its vast well of service experience. Over tens of millions of flights, its experts have slowly honed

⁹ For the record, the author is not actually aware if the EPA *did* update its test regimen in the wake of SDF20. He assumes it did, but should probably look this up before this paper goes near a press.

their understandings by: a) fastidiously deconstructing accidents for epistemological insights about their designs and assessment practices; and b) adhering a common, very specific, jetliner paradigm that ensured that any hard-earned insights continued to be relevant.¹⁰ This is to say that decades of service with very similar airplanes has allowed aviation engineers to slowly weed-out the significant uncertainties in both their designs and the tests through which they verify those designs.

Understanding technological achievement this way highlights the limits of technological ambition and it helps outline effective routes to achieving ultra-high levels of safety in complex systems. At the same time, however, it emphasizes the costs of such achievements. By directly linking the impressive safety of modern civil aviation to the industry's long, painful, but ultimately instructive history, it highlights the difficulty of replicating that success without that history. The modern governmentality of many critical technologies is premised on an understanding that experts can achieve and assess ultra-high levels of performance on the basis of tests and models alone. This is deeply misleading. The fact that we can build safe jetliners does not mean that we can build safe reactors, and when dealing with new hazardous technologies — be they bombs or bacteria — no amount of scrutiny or oversight will ever fully shield us from accidental catastrophe. It is important that we recognize this.

¹⁰ A practice I have called 'innovative restraint'.

Works Cited:

[AAIB] Air Accidents Investigation Branch (1990) "Report on the accident to Boeing 737-400, G-OBME, near Kegworth, Leicestershire on 8 January 1989" Report No: 4/1990. 25 August 1990 HMSO. London.

Bloor, D. (1976) *Knowledge and Social Imagery*, Routledge & Kegan Paul, London.

Brockway, R. (2010) "How a Biotech Company Almost Killed the World (With Booze) [cracked.com](http://www.cracked.com/article_18503_how-biotech-company-almost-killed-world-with-booze.html) April 3, 2010 online: http://www.cracked.com/article_18503_how-biotech-company-almost-killed-world-with-booze.html (accessed 1 Jun. 2017).

Collins, H. (1985) *Changing Order*. SAGE. London.

Downer, J. (2007) "When the Chick Hits the Fan: Representativeness and Reproducibility in Technological Testing," in *Social Studies of Science*. 37 (1): 7-26

Downer, J. (2011) "'737-Cabriole': The Limits of Knowledge and the Sociology of Inevitable Failure" *American Journal of Sociology*. 117 (3): 725-762

Downer, J. (2015) "The unknowable ceilings of safety: Three ways that nuclear accidents escape the calculus of risk assessments." in Taebi, B. & Roeser, S. (eds.) (2015) In in Taebi, B. & Roeser, S. (eds.) (2015) *The Ethics of Nuclear Energy: Risk, Justice and Democracy in the Post-Fukushima Era*. Cambridge University Press: 35-52

Downer, J. (2016) "Resilience in Retrospect: Interpreting Fukushima's disappearing consequences." in Herwig, A. & Simoncini, M. *Law and the Management of Disasters: The Challenge of Resilience*. Routledge. London: 42-60

Downer, J. (2017) "The Aviation Paradox: Why we can 'Know' Jetliners but not Reactors" in *Minerva*. Forthcoming.

Feyerabend P. (1975) *Against Method: Outline of an Anarchistic Theory of Knowledge*. Redwood Burn; London.

Holmes, M. T., Ingham, E. R. Doyle, J. D. and Hendricks, C. W. (1999). "Effects of *Klebsiella Planticola* SDF20 on Soil Biota and Wheat Growth in Sandy Soil." *Applied Soil Ecology* 11(1):67-78.

Jasanoff, S. (1990) *The Fifth Branch: Science Advisors as Policymakers*, Harvard University Press.

Krebs, A.V. (2001) Commentary: Searching for a Fair Resolution Concerning Controversial Story on Possible Effects of *Klebsiella P* on the Environment. in The Agribusiness Examiner. Issue #119 online: <http://gmwatch.org/en/news/archive/2001/8951-full-story-of-the-dr-elaine-ingham-controversy-over-klebsiella-p> (accessed 1 Jun. 2017).

Kripke, S. (1982) *Wittgenstein on Rules and Private Language*. Harvard University Press.

Kuhn, T. (1996 [1962]) *The Structure of Scientific Revolutions* (3rd Edition); The University of Chicago Press; Chicago.

Latour, B. (1987) *Science in Action: how to follow scientists and engineers through society*. Harvard University Press: Cambridge, (Mass.)

Le-Coze, J. C. (2015). "1984-2014. Normal Accidents. Was Charles Perrow Right for the Wrong Reasons?" *Journal of Contingencies and Crisis Management* 23 (4): 275-86.

- Lynch, M. (1998) 'The Discursive Production of Uncertainty: The OJ Simpson "Dream Team" and the Sociology of Knowledge Machine' in *Social Studies of Science* 28 (5/6): 829-869]
- MacKenzie, D. (1990) *Inventing Accuracy: A Historical Sociology of Nuclear Weapon Guidance*. MIT Press, Cambridge (mass).
- MacKenzie, D. (1996) 'How Do We Know the Properties of Artifacts? Applying the Sociology of Knowledge to Technology', in R. Fox (ed.), *Technological Change* (Amsterdam: Harwood): 249-51.
- Nöggerath, J.; Geller, R. & Gusiakov, V. (2011) "Fukushima: The myth of safety, the reality of geoscience" *Bulletin of the Atomic Scientists*. 67(5): 37-46
- Perrow, C., (1999 [1984]). *Normal Accidents: Living With High-Risk Technologies (2nd edition)*. Princeton. New York
- Pinch, T. (1993) "'Testing – One, Two, Three ... Testing!': Toward a Sociology of Testing', *Science, Technology & Human Values* 18(1): 25-41.
- Porterfield, A. (2016) Did you hear about the GMO that almost destroyed all life? Genetic Literacy Project. Sept. 6. online: <https://geneticliteracyproject.org/2016/09/06/hear-gmo-almost-destroyed-life/> (accessed Jun 1, 2017)
- Quine, W. v. O. (1951) "Two Dogmas of Empiricism" *The Philosophical Review*, 60: 20-43.
- Rip, A. (1986) "The Mutual Dependence of Risk Research and Political Context." *Science & Technology Studies* 4 (3/4), 3-15.
- Robbins, J. (2002) "A Biological Apocalypse Averted" *Earth Island Journal*. Winter 2002. online: http://www.earthisland.org/journal/index.php/eij/article/a_biological_apocalypse_averted/ (accessed 3 July 2017).
- Sagan, S. (1993) *The Limits of Safety*. Princeton: Princeton University Press.
- Servick, K. (2016) "Mouse microbes may make scientific studies harder to replicate" Aug. 16, 2016. *Science AAAS* online: <http://www.sciencemag.org/news/2016/08/mouse-microbes-may-make-scientific-studies-harder-replicate> (accessed 2 June 2017)
- Vaughan, D. (1996) *The Challenger Launch Decision*. Chicago: Univ. Chic. Press.
- Walter, C.; Berridge, M. & Tribe, D. (2001) "Genetically engineered *Klebsiella planticola*: A threat to terrestrial plant life? Rebuttal of written and verbal evidence presented by Dr Elaine R. Ingham." *Presentation to the Royal Commission on Genetic Engineering, 1 February 2001*. Press Release: New Zealand Life Sciences Network. online: <http://www.scoop.co.nz/stories/print.html?path=SC0102/S00070.htm> (accessed 1 Jun. 2017)
- Wittgenstein, L. (2001) [1953]. *Philosophical Investigations*. Blackwell Publishing.