Downer, J. (2014). Disowning Fukushima: Managing the credibility of nuclear reliability assessment in the wake of disaster. *Regulation and Governance*, 8(3), 287-309. <https://doi.org/10.1111/rego.12029>

Early version, also known as pre-print

Link to published version (if available):
[10.1111/rego.12029](https://doi.org/10.1111/rego.12029)

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

This is the pre-peer reviewed version of the following article: Downer, J. (2013), Disowning Fukushima: Managing the credibility of nuclear reliability assessment in the wake of disaster. *Regulation & Governance*. doi: 10.1111/rego.12029 , which has been published in final form a t10.1111/rego.12029

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/>

Disowning Fukushima:

Managing the Credibility of Nuclear Reliability Assessment In The Wake Of Disaster.

<Forthcoming in *Regulation & Governance* (2013)>

John Downer ¹

Abstract:

This paper reflects on the credibility of nuclear risk assessment in the wake of 2011 Fukushima meltdowns. In democratic states, policymaking around nuclear energy has long been premised on an understanding that experts can objectively and accurately calculate the probability of catastrophic accidents. Yet the Fukushima disaster lends credence to the substantial body of social science research that suggests such calculations are fundamentally unworkable. Nevertheless, the credibility of these assessments appears to have survived the disaster, just as it has resisted the evidence of previous nuclear accidents. This paper looks at why. It argues that public narratives of the Fukushima disaster invariably frame it in ways that allow risk-assessment experts to 'disown' it. It concludes that although these narratives are both rhetorically compelling and highly consequential to the governance of nuclear power, they are not entirely credible.

Eight years involved with the nuclear industry have taught me that when nothing can possibly go wrong and every avenue has been covered, then is the time to buy a house on the next continent.

~Terry Prachett

1. FUKUSHIMA FOOTBALL

*Introduction*²

Nuclear accidents and their attendant dramas – evacuations, denials, health alerts, heroics – make compelling news coverage, and so in March 2011, when a giant earthquake off the coast of Japan instigated a series of reactor failures at the Fukushima Daiichi nuclear plant, the world's media were transfixed. The world watched avidly as the accident quickly escalated in severity, one increment of the

¹ SPAIS. University of Bristol. Email: jrd22@cornell.edu

² The author would like to thank Chick Perrow; Lynn Eden; Peter Ladkin; Benoît Péloupas; Lori Bystrom; Anne I. Harrington; Henry Rothstein, Rebecca Slayton, Columba Peoples and three anonymous reviewers for their time, support and advice at different stages of this paper's overlong gestation. All its failings are, as ever, the author's alone. This research was funded by the Stanton Foundation.

‘International Nuclear Event Scale’ at a time, until ‘Fukushima’ reached taxonomic parity with ‘Chernobyl’ as a new synonym for technological disaster.

Amid the discourse of disaster, it was tempting to reflect on the hundreds of other nuclear plants around the world. And when British and French officials started advising their nationals to evacuate Tokyo (advice, it would later emerge, the Japanese government was itself considering at the highest levels [Quintana 2012; Fackler 2011]), it was easy to wonder if the people of London, Paris or New York might one day find themselves hurrying bleary-eyed children into cars and filing into tense traffic jams, uncertain if they would ever see their homes again.

Nuclear authorities – the IAEA; regulatory agencies; national governments – and their experts quickly moved to assuage such fears, each explaining, in varying ways, why the disaster in Japan should not undermine the credibility of the nuclear industry in general. Given that the same authorities had been promising for decades that Japanese reactors were safe, however, the accident posed an important question: how can the assurances of nuclear experts remain credible in the wake of Fukushima?

This is the question that motivates this paper. The answer, it will suggest, lies in the fact that both technical and popular accounts of the disaster routinely frame it in ways that allow nuclear experts to ‘disown’ it; thereby protecting their authority, the legitimacy of their work, and more broadly, the viability of an ideal on which both depend. Rather than try to provide another account of the accident,³ in other words, this paper will look critically at the accounts themselves: their rhetoric, their reasoning, and their relationship to wider narratives around the governance of critical infrastructures.

A Failure of Foresight

The first thing to note about the Fukushima disaster (henceforth referred to as ‘Fukushima’) is that was *unexpected*. And not just in a specific sense – nobody ever expects a specific nuclear plant to melt down on any given day – but in the much more foundational sense that nobody expects *any* nuclear plants to melt down.

To understand this distinction consider the realm of civil aviation. In this realm we understand that airplanes very occasionally, but nevertheless *routinely*, crash, even though any *specific* crash is always unexpected. We accept such accidents as an inevitable cost of the technology and we anticipate them in our plans and institutions. The same logic does not hold in the nuclear sphere, however. Reactors, in this respect, are more comparable to dams or bridges, in that public decisions about them are predicated on an understanding that the chance of them failing catastrophically is so low as to be negligible. Modern democracies, we might say, are institutionally blind to the possibility of nuclear meltdowns.

This blindness was evinced by the degree to which Japan seemed under-prepared for the disaster. At the plant itself, for instance, procedures and guidelines were woefully insufficient for a meltdown,

³ For accounts of the accident itself see eg: Osnos (2011); Strickland (2011).

forcing operators to respond on an almost entirely *ad hoc* basis. When the lights went out they had to borrow flashlights from nearby homes to study the plant's gauges (Osnos 2011: 50). On-site dosimeters maxed-out at levels that were far below those that could be expected from a catastrophe, unable to display readings any higher. And, as the plant's former safety manager would later testify, emergency plans "...had no mention of using sea-water to cool the core" (quoted in Adelstein & McNeill 2011), an oversight that caused unnecessary, and perhaps critical, delays. The bureaucratic infrastructure beyond the plant evinced similar shortcomings. Official announcements were often ill-considered and characterized by denial, secrecy, and refusal to accept outside help (Perrow 2011b: 50). Important medications were scarce, with byzantine rules hindering their effective distribution (Kubiak 2011). Evacuations were poorly managed, with vital information being withheld in ways that exposed many families to unnecessary danger (Onishi & Fackler 2011).⁴

We should be wary, moreover, of seeing such failures of foresight as a specifically Japanese problem. Very few observers suggest that other nations are substantially better prepared for such a contingency (see, eg: Perrow 2011b: 46-7; Kahn 2011). As Perrow (2007) and others compellingly argue, all societies routinely make choices about nuclear power that seem myopic when considered in relation to a potential disaster. A conspicuous example being the routine 'clustering' of several nuclear reactors in a single facility, a choice that offers significant economic and political benefits but creates conditions where the failure of one unit can propagate to others (Perrow 2007: 136). (The Fukushima plant, for instance, was a cluster of six reactors – something that workers must have rued as the fallout from the first explosion created a deadly radioactive hot-zone around its neighbors, ultimately thwarting efforts to contain the disaster [Osnos 2011: 50; Strickland 2011]).

This is not to say that disaster planning is entirely absent in the nuclear sphere, but rather that it is routinely insincere and insufficient. Clarke and Perrow (1996), for instance, describe some of the evacuation planning undertaken for Shoreham Nuclear Power Station on Long Island, but argue that the assumptions of these plans were so unrealistic that they are more properly understood not as earnest contingency preparation but as elaborate public performances – bureaucratic Kabuki.⁵ When organizations do attempt to think more earnestly about nuclear disaster, meanwhile, the plans they produce routinely lack any institutional authority. The US Nuclear Regulatory Commission (NRC), for instance, developed "Severe Accident Management Guidelines" (SAMGs) for directing reactor operations in the event of "unanticipated accident sequences" (ie: events like Fukushima), but training in these guidelines was voluntary. At the time of the crisis the NRC did not require that operators demonstrate any knowledge of the SAMGs or their application (Lochbaum 2011), and a recent audit of US plants has found them to have been largely neglected (NRC 2011).

⁴ A failure to utilize and distribute forecasts from a computer system known as the 'System for Prediction of Environmental Emergency Dose Information,' for instance, meant that many evacuees spent several days in areas that were known, by some, to be dangerously contaminated (Onishi & Fackler 2011).

⁵ In this way, we might say that the plans for nuclear meltdowns echo those for nuclear wars (eg: Eden 2004; Clarke 1999).

Nuclear Probabilism

This lack of preparedness for nuclear disasters is puzzling, given that with the recent turn to ‘resilience’ in policymaking. Western states now invest heavily in planning and preparing for disasters in many spheres – in respect to flooding and terrorism, for instance (eg: Walker & Cooper 2011; Collier 2008; Duffield 2011). So why the neglect in the nuclear sphere?

The answer lies in an institutionally deep-rooted confidence that contingency planning is unnecessary for nuclear disasters. For, as Beck (1999: 150) puts it, “our risk assessment bureaucracies have found ways to deny systemic hazards”. Policymakers treat the risks of nuclear meltdowns as if they are *objectively calculable* in way that the risks of terrorism are not, and *entirely preventable* in a way that floods are not. Thus they defer to expert assurances that nuclear accidents, unlike floods and acts of terror, will not happen, and so need not be prepared for.

Various scholars have noted this institutional disavowal of nuclear disaster (eg: Rip 1986; Clarke 2005), noting that it emerged in the 1960s and 1970s as a response to specific political contingencies. For reasons relating to the Cold War and nuclear weapons, they argue, the West (particularly the US) was committed to nuclear power at the highest levels (Kuznick 2011). Yet official nuclear regulatory studies were threatening the political viability of this commitment by projecting that the potential consequences of an accident could be severe. To circumvent this problem, scholars argue, regulatory agencies such as the US NRC (then the Atomic Energy Commission) shifted the focus of their studies away from assessments of the *consequences* of a disaster (their primary focus until then) and towards assessments of *likelihood*. Rather than justifying nuclear power by exploring the implications of accidents, in other words, authorities sought to draw a distinction between “credible” accidents, which were worthy of political concern, and “hypothetical” accidents, which were not (Rip 1986 4-9; also Clarke 2005).⁶ A position that was neatly illustrated by a former Safety Director of British Nuclear Fuels Limited, when he declared that: “once an accident becomes inconceivable, as far as I am concerned it is impossible” (in Donoghue 1977: 14-15).

At the heart of this transition was a turn to probabilistic reliability analysis: the creation and adoption of a series of calculative audit tools that offered regulators a means of ‘objectively’ establishing that nuclear accidents were too improbable to merit serious discussion (Rip 1986: 7-9; also: Fuller 1976: 149-186). The NRC’s 1975 “WASH-1400” study of reactor safety is usually regarded as the key document in this revolution. The study, which leant heavily on the 1967 ‘Rasmussen report’ on reactor accidents, was hailed as “a landmark achievement in risk analysis” (Rip 1986: 7). Like the Rasmussen report, it arrived at its accident probabilities by using models that combined the failure data of generic reactor parts and sub-systems to mathematically derive the failure probability of the wider system (Rip 1986: 8).

⁶ As Sagan (2004: 944) has observed in relation to nuclear weapons “organizations can [marginalize] low-probability events by transforming them into assumptions of impossibility.”

Significant elements of the scientific and engineering community were deeply skeptical of this turn to probabilism, criticizing both the Rasmussen report and the WASH study that built on it. A 1977 Congressional review panel of safety experts, for instance, “...criticized [the Rasmussen report’s] inadequate data base and questionable methodological and statistical procedures...” arguing that they “...led to uncertainty bands which greatly understated the actual uncertainty” (Rip 1986: 8). Brookhaven National Laboratory, which had overseen previous WASH studies and had been the preferred contractor for WASH-1400, took a similar line. It rejected the Rasmussen study’s probabilistic approach to reactor assessments, arguing that “...it would be ‘charlatanism’ to make predictions about probability, and that only ‘fringe members of the statistical community’ would attempt to do this with the meagre data available” (Rip 1986: 7; see also: Fuller 1976: 141-8).

Brookhaven’s recalcitrance on this issue eventually led it be replaced as contractor for WASH-1400, but doubts remained. In 1978, the NRC sponsored an *ad hoc* ‘Risk Assessment Review Group’ (RARG), which found in its report that the new safety studies were failing to “...sufficiently emphasize the uncertainties involved” in their findings, with the effect that probabilistic assessment was being “...misused as a vehicle to judge the acceptability of reactor risks” (RARG 1978: ix-x). The same year, a member of the NRC ‘Advisory Committee on Reactor Safety’ echoed this view, writing of WASH-1400 that “...although the job was done in a workman-like way, many of the underlying facts, which must be known to accurately predict the course of an accident, are lacking...” with the effect that “...the quantitative estimates of the probability of the various accident chains must be viewed with some reservations” (Okrent 1978: 17). Under pressure from such reports, even the NRC distanced itself from WASH-1400, conceding, according to Rip (1986: 8), that the report “did not regard the numerical estimate of the overall risk of reactor accident reliable.”

For all the recognition of WASH-1400’s imperfections, however, its probabilistic approach to risk became entrenched in the mechanics of nuclear regulation (Rip 1986: 8). Subsequent NRC safety analyses, such as CRAC-II and NUREG-1150, updated WASH-1400’s analytical models, but all adhered to the same underlying premise that the probability of accidents was a fundamentally *calculable* property. (And all similarly determined that serious accidents were incredibly unlikely).

Today, probabilism lies at the heart of all nuclear regulatory oversight, with almost every decision being driven by Probabilistic Risk Assessment (PRA), and any qualifications and caveats happening ‘offstage’ (insofar as they happen at all). In a recent declaration to a UK regulator, for instance, *Areva*, a prominent French nuclear manufacturer, invoked probabilistic calculations to assert that the likelihood of a “core damage incident” in its new ‘EPR’ reactor were of the order of one incident (per reactor) every 1.6 million years, and to conclude that the probability of a core-melt is “infinitesimal” (in Ramana 2011). Similar assessments performed on earlier generation reactors like Fukushima have found the risks to be higher, but still nowhere of an order that would justify significant public concern. An oft-cited figure that framed nuclear discussions before 2011, for instance, put the likelihood of a severe accident

in an operational reactor at one in a hundred-thousand years (see eg: Glaser 2011). Such calculations legitimate almost all contemporary public pronouncements and policy discourse about nuclear power (explaining, for instance, why the costs of accidents are routinely ignored when weighing the costs of different energy options [eg: OECD 2010]).

The Ideal of Objectivity

Reservations about probabilistic numbers have certainly not disappeared from engineering discourse, even if they have declined in their prominence and prevalence. An overview of a 2010 joint NASA/NRC meeting to discuss the role of risk assessments, for example, suggests the discussion involved a nuanced understanding of the limitations of probabilistic assessments (Youngblood et al. 2010); it even goes so far as to emphasize the line: “All models are wrong, but some models are useful” (Youngblood et al. 2010). On the same page, however, the same document highlights the perils of communicating such complexities: “Overly detailed models may erode decisionmaker confidence,” it cautions (ibid.). And this latter concern is more illustrative of modern nuclear discourse. For although nuclear engineers have continued to discuss the nuances of probabilism, such discussions have all but disappeared from the dialogues between those experts and the publics or policymakers they serve.

Despite the common refrain of nuclear regulators that probabilistic assessments are intended to enhance the “transparency and objectivity of decisionmaking” (NRC 2002), there has been an increasing disconnect between the public and private discourse of professional engineering circles. Assessment experts may be aware of the limitations of their calculations, in other words, but on an institutional level, at least, they actively occlude such limitations from non-specialists. In keeping with the foundational role that probabilistic assessments now play in nuclear regulatory decisionmaking, regulators, together with the industry, routinely promulgate the notion that experts can take a nuclear power plant – an enormously complex system with sophisticated socio-organizational dimensions and a stochastic operating environment – and mathematically deduce its reliability to seven decimal places (the level implied by some projected accident frequencies).

Elsewhere (eg: Downer 2011b), I have referred to this idea – that complex technological properties like risk and reliability should be wholly, objectively and quantitatively knowable through formal rules and unbending algorithms – as the “Ideal of Mechanical Objectivity,” a phrase I borrow, slightly loosely, from Porter (1995). It is an ideal that characterizes most modern technological discourse, but it is especially prominent in the context of nuclear power (Wynne 1982). It is easy to imagine why. As US authorities determined in the 1960s, the potential gravity of nuclear accidents begs for declarative assurances and calculative certainty.

On a day-to-day basis this Ideal of Mechanical Objectivity is reflected in the subtext of expert pronouncements about nuclear risk (“*The math is the math*”; “*It’s not an opinion, it’s a calculation*”). It is implicitly evident in the substantial body of social science research that treats nuclear risk as an

established property, to be contrasted, for instance, with peoples' perceptions of that risk (eg: Slovic 2012; Weart 1988). An unguarded illustration of its prevalence in nuclear regulatory discourse can be found the NRC's (2004) "Guidelines for External Risk Communication" – a brightly illustrated and uncommonly revelatory document that outlines a straightforwardly calculative understanding of risk, which it encourages regulators to defend in their public pronouncements. "Avoid making statements such as 'I cannot guarantee...' or 'There are no guarantees in life...'" it advises, as "...statements like these contribute to public outrage because they reinforce feelings of helplessness and lack of individual control" (NRC 2004: 38). (It is worth noting, moreover, that the report – authored by the NRC's *Division of Systems Analysis and Regulatory Effectiveness* – does not suggest that policymakers are less in need of such protection. Quite the opposite. In a sidebar reserved for exemplary quotes and aphorisms it says: "When the media publishes the NRC's talking points and messages and people refer to them for decisionmaking, that's success" [NRC 2004: 48].)

Sociological Misgivings

Sociologists of knowledge have long recognized that this 'mechanical' notion of technology assessment is idealized, and have worked to highlight its shortcomings. Most such critiques highlight how assessment calculations quietly exclude other considerations by shaping what Foucault would call the 'conditions of discourse'. Writers such as Wynne (1983); Rip (1986); and Otway & von Winterfeldt (1982), for instance, criticize technology assessment for its 'narrowness of vision,' arguing that it tends to impoverish definitions of what is at stake leading to ever more refined answers to what are often the wrong questions. When Wynne (1983: 27) speaks about the "myth of expert objectivity," for example, he is pointing to the fact that there are varying interpretations of 'rationality' by which we might want to judge a technological system. Otway & von Winterfeldt (1982: 133) make a similar point, arguing that a group's attitudes to nuclear power might legitimately hinge on such factors as "...whether the technology leads to a reliance on technical experts; whether it will increase economic growth; [...] and whether it will increase the power of big business." All of which, they point out, concern social and economic matters that are quite separate from risk in the sense invoked by formal risk assessments.

Most critics in this tradition follow a similar path, emphasizing the ways that formal risk calculations limit discussions about social priorities, rather than exploring the inherent limitations of the calculations themselves.⁷ The social science literature does, however, contain several well-articulated arguments for why complex systems like nuclear plants *necessarily* harbor risks that escape the calculus of formal assessments.

By far the most prominent approach to the latter argument is Normal Accident Theory (NAT) (Perrow 1999 [1984]), with its simple but deceptively profound insight that accidents caused by very improbable

⁷ It is true that the literature on contemporary organizations offers many discussions of complex but ultimately invalid calculations being exploited for instrumental ends (e.g.: Lampland, 2010), but this discourse rarely delves into the 'hard' calculations of scientists and engineers.

confluences of events (which no risk calculation could ever anticipate) are ‘probable’ in systems where there are many opportunities for them to occur (ie: that the organizational logic of extremely complex systems allows for billions of potential ‘billion-to-one’ accidents, and so it is only to be expected that we would see them from time to time.) There is another approach to this argument, however, that draws on the sociology of knowledge and Science and Technology Studies (STS) literatures. It points to the irreducible epistemological ambiguity of all technological knowledge (eg: MacKenzie 1996; Wynne 1988), and argues that this ambiguity, although manageable in most spheres, becomes crippling in assessments of complex, safety-critical systems that deal with very small probabilities (Downer 2011b).

These sociological arguments about why nuclear risk calculations *must* be insufficient are important, especially given that expert authorities are inherently unwilling to undermine their own credibility and self-interest in a sustained way (Wynne 2011: xxii). At the same time, however, they can obscure a more interesting sociological question about why Fukushima (and, indeed, the disaster-punctuated history of nuclear power more broadly) doesn’t speak for itself.

Enduring Credibility

Such misgivings carry little weight in wider public discourse, however. Power (2011: 28) argues that crises damage the credibility of expertise, but the history of nuclear risk assessment seems to challenge this finding. For nuclear experts remain credible even though, as Ramana (2011) highlights, nuclear disasters happen far more often than formal calculations predict. Indeed, the fallibility of nuclear risk assessments is evinced by a litany of accidents and near accidents,⁸ some of which – Windscale; Three Mile Island; Chernobyl – are widely recognized and remembered, but none of which seem to have effectively undermined the pervasive belief that nuclear risks can be calculated with objective certainty.⁹ Windscale, we believe, was an exception; Chernobyl was an exception; Three Mile Island (TMI) was an exception; Fukushima seems to be following the same trajectory.

Admittedly, Fukushima has led to policy reversals in some spheres, most notably in Japan and Germany, but in most other contexts it has done little to undermine public faith in nuclear risk calculations. Even before the close of 2011 there were strong indications that the long-anticipated resurgence of nuclear power would survive the disaster, and by March 2012, the UN was predicting that the global use of nuclear energy could increase by as much as 100 percent in the next two decades, and projecting

⁸ Any list of nuclear accidents will be inherently contentious, but Wikipedia (as of Dec. 2012) points to twenty-four “Nuclear power plant accidents and incidents with multiple fatalities and/or more than US\$100 million in property damage.”

⁹ Proofs of low probability, it should be said, are not absolute proofs of safety. So it is true, theoretically, that Fukushima’s risk calculations could have been correct and Japan could have just been unlucky. Given the levels of probability outlined above, however, Japan would have had to have been incredibly unlucky for this to be the case. It is theoretically possible that a perfectly balanced coin, flipped a thousand times, will land heads-up on every flip. But if this actually happened, then any *a priori* calculations proving the coin’s ‘perfect balance’ would have to be exceptionally convincing to be credible, whereas nuclear risk calculations are contested at best. (Even if Fukushima’s assessments were correct, in other words, the accident should undermine our Bayesian confidence in them.) This is not to mention the other nuclear accidents, of course, or the fact that, as we will see below, few experts now contest that there were fundamental errors in Fukushima’s risk assessments.

further growth of between 35 percent and 100 percent by 2030 (In Dahl 2012). Such predictions were supported by a series of decisions that implied both a continued institutional confidence in nuclear power and a large degree of (at least passive) consent from the public. Before the accident, a total of 547 reactors were either proposed, planned or under construction throughout the world; by early 2012, this number had increased to 558 (Holloway 2012). In February 2012, Britain and France signed a formal agreement that paved the way for a new generation of reactors in both countries (Press Association 2012; BBC 2012). That same month, in the US, the NRC approved licenses to build two new reactors in Georgia, the first such approvals since TMI in 1979 (Abernethy 2012). The following month it approved two more in South Carolina (Wingfield & Johnsson 2012).

At the risk of trivializing a serious issue, the relationship between public discourse and nuclear risk assessment might be compared with a recurrent scene in Charles Schulz's pensive comic strip, *Peanuts*. In the first such scene, Charlie asks Lucy to hold a football for him to kick. Lucy obliges only to whip the ball away at the last second, causing him to fall. The scene's subsequent iterations always involve Lucy convincing Charlie that 'this time it will be different' and that he should trust her to hold the football, only for her to pull the same stunt again and again. Charlie is fooled in this way for years, always giving in to Lucy's reasonable explanations and never learning the larger lesson that Lucy is not to be trusted.

Western societies, we might say, have a similar relationship with nuclear expertise. The nuclear establishment seems to enjoy a 'Lucy-esque' ability to justify why its past failings should not count against the current and future credibility of its assurances. We keep being asked to take a punt on calculations that then let us down, but after each failure our confidence rebounds and we bet again on 'calculative certainty.' Proponents of nuclear reliability calculations, in other words, have routinely found ways to 'escape' from being tarred by evidence of their apparent failures – ways to maintain their credibility, even as the consequences of past accidents loom large in the public imagination and heavy on the public purse. These arguments are so convincing that even Fukushima, with all its visceral undeniability, cannot compel publics and policymakers to recognize that reliability calculations are less credible than they appear.

The remainder of this essay will examine the shape of these arguments, and the extent to which they are credible.

2. RATIONALIZING THE MELTDOWN

The question of when it should be appropriate to reject a belief has long been of interest to philosophers of science. Karl Popper (1959), for instance, is widely (if slightly inaccurately) remembered for arguing that all knowledge-claims are essentially theories that should be discarded if they are disproven ('falsified') by experience. A strict Popperian, then, might consider the Fukushima meltdowns

to be proof that nuclear reliability assessments are inaccurate, and balk at our refusal to acknowledge the falsification of the claim that such assessments are trustworthy.

As historians and philosophers of knowledge have long acknowledged, however, an unwillingness to bow before seemingly disconfirming evidence is not unusual in science, and is not always desirable in principle (eg: Stanford 2009; Duhem 1954; Quine 1960). They point out that no ‘disproof’ can ever be conclusive. An amateur scientist who tests the boiling point of water and finds it to be 98 degrees Celsius, for instance, is far more likely to have discovered a faulty thermometer than to have ‘falsified’ the claim that water boils at 100 degrees. In other words, the purity of the water and the workings of the thermometer can themselves be treated as theories, and thus can be contested in lieu of the boiling point.¹⁰ So it is that scientists routinely hold onto theories even in the face of contrary evidence. Newton did not abandon his theories of gravity simply because he lacked a good explanation for birds.

It is arguably misleading, therefore, to see Fukushima as *self-evident* proof that nuclear reliability assessment cannot work or that nuclear authorities are in denial. For in much the same vein as Popper’s critics above, the discourse around Fukushima not only denies that the accident discredits the practice of reliability assessment, it also offers rationales for *why* the accident should not disprove the claims of expert reliability calculations more broadly.

These broad rationales are woven into diverse and highly overlapping narratives. To understand their underlying logics, however, it helps to parse them into four core arguments:

- i. That the assessments did not *actually* fail (what I will call the *interpretive defense*).
- ii. That the failure of one assessment is not relevant to the credibility of others (what I will call the *relevance defense*).
- iii. That the assessments were sound, but people did not obey the rules (what I will call the *compliance defense*).
- iv. That the assessments were flawed, but now they are fixed (what I will call the *redemption defense*).

I will sketch each of these in turn, before discussing their merits in a later section.

i. The Interpretive Defense

One way of insulating reliability assessments from a nuclear accident is simply to *deny there has been a failure*. Occasionally such denials are straightforwardly untrue, such as when Soviet authorities denied

¹⁰ Philosophers, following Duhem, refer to this as the ‘underdetermination of theory by evidence’, or sometimes as ‘confirmation holism’.

the accident at Chernobyl, even as they were organizing mass evacuations (Perrow 2011b: 50). Such denials are rarely so ostensibly mendacious, however, as most exploit the fact that the definition of ‘failure’ is open to legitimate interpretation.

It is rarely straightforward to say when, or if, a nuclear plant has definitively ‘failed.’ In part, this is because all reactors are designed to tolerate some degree of malfunction – not every blown fuse or broken valve at reactor constitutes a ‘failure’ of the plant or a nuclear ‘accident,’ for example – and reliability calculations reflect this. On a slightly deeper level, moreover, critical technologies are often designed to fail ‘safely’. Reactors have containment structures designed to keep failures from becoming major public hazards, for instance, and reliability analyses account for this as well. If an unanticipated reactor event is relatively contained, therefore, then observers can plausibly say that the plant – as a *system* – has functioned as designed, and that there was no real ‘accident’ that should challenge the validity of its reliability assessments.

The spectacular failure of Fukushima, we might think, should have been unambiguous in these regards. Yet even as the significance and extent of the disaster became undeniable, experts were able to further parse the definition of ‘failure’ to deny that their assessments had failed. This was visible, for instance, in the many ‘beyond design basis’ arguments.

The ‘design basis’ of a nuclear plant is the set of assumptions that frame the work of the experts who design and assess that plant. So if the design basis states that the worst flood a nuclear plant will be subjected to is ten meters, and the plant fails because it is flooded to a depth of twenty meters, then both the engineers who built the plant and the assessors who certified it can deny they have failed because the flooding was ‘beyond design basis’. There were no errors in their calculations, they can say, because the calculations never claimed that the plant could survive such events; it is not their fault if Fukushima was struck by what insurance contracts refer to as an ‘Act of God’. As the American Nuclear Society (2011) put it, for instance: Fukushima “...could actually be considered a ‘success’ given the scale of this natural disaster that had not been considered in the original design.” A view that was echoed by Sir David King, the former chief scientific advisor to the UK Government, who, shortly after the accident, reassured the press that all the nuclear plants affected by the tsunami had “acted as they were meant to, including Fukushima” and that the tsunami was “an extremely unlikely event” that overwhelmed defenses designed for much lower levels of flooding (in Harvey 2011); and which appears in even purer form in this passage in the *New American*:

“...the Fukushima “disaster” will become the rallying cry against nuclear power. Few will remember that the plant stayed generally intact despite being hit by an earthquake with more than six times the energy the plant was designed to withstand, plus a tsunami estimated at 49 feet that swept away backup generators 33 feet above sea level” (Hiserodt 2011).

ii. The Relevance Defense.

A second means of logically protecting the credibility of nuclear risk-assessment from a conspicuous failure is to isolate the failed assessment and then establish that it is *unrepresentative* of assessment practices as whole. If accounts can show that Fukushima's (or, more widely, Japan's) risk calculations were *significantly different from other such calculations*, in other words, then *the failings of those calculations can be theirs alone*.

Claims about relevance and representativeness, in varying guises, are extremely common in the discourse around Fukushima, (as they are in the discourse around most technological accidents). Most are framed around the representativeness of the technology itself. A report for the American Academy of Arts and Sciences, for instance, explained that the plant's reactors (GE Mark-1s) did "...not reflect the safety improvements of more recently designed reactors" (Marvel & May 2011: 3), while Guardian columnist George Monbiot (2011b), echoed the voice of many journalists when he described the facility as "a crappy old plant with inadequate safety features". Logically, of course, such claims about the plant's design do little to redeem the practices by which it was assessed and approved, but they often imply that assessment practices evolved with the technology, with the result that the practices that governed the Mark-1's approval are no longer characteristic of assessment practices more generally. There are echoes of this argument, for instance, in an internal UK government email sent soon after the accident (and later published in the Guardian) where a civil servant writes: "We need to [...] show that events in Japan, whilst looking dramatic, are all part of the safety processes of this 1960's reactor."¹¹

Another very common form of the relevance defense rests on claims about the representativeness of Fukushima's specific assessment regime and regulatory environment. These claims often draw distinctions between what happened "over there" (in Japan) and what goes on "over here" (in Europe/US) (Gusterson 2011; Greene 2012). The *New York Times*, for instance, reported at length on Japanese regulatory shortcomings, running articles with titles such as "Culture of Complicity Tied to Stricken Nuclear Plant," (Onishi & Belson 2011); and "Japan ignored or long hid nuclear risks" (Onishi & Fackler 2011b). Many official reports have followed suit, with even the official report of the Japanese Diet finding that the accident's "fundamental causes" lay in the "ingrained conventions of Japanese culture" (NAIIC 2012: 9).

iii. The Compliance Defense.

As the headlines above suggest, most (although not all)¹² of the arguments about Japanese exceptionalism attribute the disaster to an unusual degree of malfeasance or incompetence on behalf of its operators and overseers. The report to the Japanese Diet, for instance, strongly recommended that Fukushima be understood as a 'man-made' disaster (NAIIC 2012), while an earlier report by the Carnegie

¹¹ See: <http://www.guardian.co.uk/environment/interactive/2011/jun/30/email-nuclear-uk-government-fukushima?intcmp=239> (accessed Jun 2012).

¹² Some, for instance, attribute it to Japan's unusual seismic activity.

Endowment for International Peace came to much the same conclusion (Acton & Hibbs 2012a). The press, meanwhile, highlighted a spectrum of purported malfeasances (eg: Ichida et al. 2011; Broder 2011), reporting, for instance, that TEPCO (the operator) had covered-up a series of regulatory breaches for years, including data about cracks in critical circulation pipes that could have been instrumental in the catastrophe (eg: Adelstein & McNeill 2011).

Highlighting human failure in such ways offers a third narrative through which accounts of Fukushima can redeem the credibility of risk assessments. This is because we (as a society) routinely allow technology assessments to assume, within certain bounds, that the people operating or regulating critical systems will strictly adhere to rules, or, at minimum, will only violate those rules in predictable and circumscribed ways (Downer 2011a: 276-7).¹³ In other words, all technology risk assessments embody implicit and explicit caveats, such as: “...given proper maintenance,” or “...if handled correctly,” and because we accept these caveats it is possible to exculpate the assessments after a disaster by highlighting operator-error, noncompliance or malfeasance – to claim, essentially, that the *‘calculations would have been sound if people had only obeyed the rules.’*

Accounts of the errors and wrongdoings around Fukushima overwhelmingly invoke this argument. They ‘isolate’ the accident by blaming it on the human failings, which they then construe as local deviances from the norm rather than as fundamental problems with the norm itself. The implication, again, being that the disaster is not (or need not be) ‘relevant’ to other reactors or risk assessments. The Carnegie report, for instance, carries the title “Why Fukushima Was preventable” and, after documenting a series of organizational errors, concludes that: “In the final analysis, the [...] accident does not reveal a previously unknown fatal flaw associated with nuclear power” (Acton & Hibbs 2012a: 2). (A conclusion echoed by the NAIIC report).

This defense leans on more than Japanese exceptionalism, however, with some accounts of Fukushima point to more widespread errors and malfeasances. Accounts that noted the plant’s lack of guidance for operations under extreme crisis, for instance, could hardly ignore that such guidance was deficient in the US and elsewhere (eg: Lochbaum 2011). On these occasions a different logic invariably serves to redeem the credibility of nuclear reliability assessments: the argument that such problems are ‘fixable.’ The Carnegie Endowment report, for instance, suggested that accidents might be prevented in the future by “...periodically reevaluating plant safety in light of [...] evolving best practices” (Acton & Hibbs 2012: 2).

In fact, this ‘promise of perfectibility’ is prominent enough to merit being highlighted as a final logic by which to defend assessments from disaster. To wit:

¹³ It is understood, for instance, that pilots get tired and sometimes respond inaccurately when under pressure, and this is why rules stipulate (and risk assessments assume) that there be two in a cockpit to share loads and check each other’s work.

iv. The Redemption Defense.

"If nuclear power is to have a future in this country..." declared a New York Times editorial after the accident, "...Americans have to have confidence that regulators and the industry are learning the lessons of Fukushima and taking all steps necessary to ensure safety" (NYT 2011). This sentiment neatly captures a final approach to redeeming risk assessments, which is simply to concede the existence of errors in the plant's risk calculations (in their assumptions about tsunamis, for instance) and then to argue that experts have *identified the errors, altered the assessments, and remedied the problem*. In this way even far-reaching assessment failures can be made irrelevant to the credibility of future assessments.

Such arguments are implicit in almost all the discourse around Fukushima. In the immediate aftermath of the accident, for example, various nations and organizations announced plans to review their risk assessment practices for errors. The EU, for instance, announced it would reassess all its nuclear plants (Willsher 2011). As did the NRC, which, after its own swift review, publicly found that its assessments "do not adequately weigh the risk" of threats to emergency generators (Wald 2011). Of course, with each devil uncovered in the details of the assessment process there came an exorcism. The NRC began reframing its assessments to better account for the newly apparent threats of flooding. The IAEA announced a 'five-point plan' to strengthen reactor oversight around the world (Amano 2011).

It is not clear that many of these reviews led to substantial corrections. When the NRC subsequently approved the licenses for two new reactors in Georgia, for instance, it did so with no binding commitment to implement changes in federal requirements arising from the NRC's post-Fukushima work, (a factor that led then NRC Chairman, Gregory Jaczko, to vote against the approval) (Abernethy 2012). It is extremely difficult, moreover, to imagine that such reviews would ever conclude that the assessment process was fundamentally unrealizable. This was not their purpose.

3. REDEMPTION CROSS-EXAMINED

Here, then, are four narratives of redemption, encompassing arguments about definition, relevance, compliance and correction. Four tropes through which accounts of Fukushima can implicitly and explicitly reconcile the unexpectedness of the accident with the Mechanical Ideal of Objectivity and the idea that quantitative risk assessments, despite their failure in this instance, should remain credible as a basis for nuclear policy. But how credible are these narratives? This question will be the focus of the next part of this paper.

Without doubt, the narratives outlined above can be rhetorically compelling, and they undoubtedly contain a degree of truth. Failure *is* an ambiguous property. Assessment practices *do* differ in ways that

limit the broader relevance of specific shortcomings. People *can* undermine assessment calculations if they break prescribed rules, (and clearly some *did* break rules in this instance). Risk calculations undoubtedly *will* improve with the lessons of experience. And so, in Fukushima's wake, it is vital that we endeavor to more accurately weigh evidence of harm; examine ways that assessment regimes differ; highlight errors and wrongdoings; and hone assessments by identifying errors. To return briefly to the philosophy of science: it is undeniably *good* and *necessary* that nuclear regulators, like scientists, do not always abandon their theories and practices at the first signs of falsifying evidence.

At the same time, however, Karl Popper, champion of 'falsifiability', undeniably had a point when he argued that we sometimes have to let ideas be disproven by experience. Insofar as scientists have improved their understanding of the world over time (and clearly they have), it is because they have often abandoned cherished paradigms because of disconfirming evidence. Evidence might be infinitely interpretable, as Popper's critics suggest, but there has to come a time when it becomes too credible to ignore, even if not everyone agrees.

With this in mind, therefore, it is worth looking critically at the 'redemption narratives' outlined above. The following section of this paper will briefly weigh the merits of each in turn.

i. Interpretation Reexamined

When considered critically, the 'interpretive defense' – emphasizing the plant's 'design basis' and the ambiguous definition of failure – is difficult to defend. For while it is undoubtedly useful to parse the meaning of failure in many contexts, it is hard to deny that Fukushima's meltdowns represent an assessment shortfall; especially given that this view has since become orthodoxy among the regulatory community itself (see eg: NAIIC 2012).

The logic of the 'beyond-design-basis' argument, for instance, has merits in some circumstances but falls down when used to justify assessment failures. Nobody denies that the earthquake and tsunami were beyond Fukushima's design basis. But although this exculpates many of the individuals involved in the plant's design and construction (they could only build the plant to the specifications they were given), it does very little to redeem the credibility of the plant's regulatory assessments. The real art of risk calculation does not just involve *applying* formulae correctly, but *framing* those formulae correctly. The design basis of a nuclear plant, almost by definition, should be an accurate representation of the conditions the plant will encounter in its lifetime. To say that an event was 'beyond design basis,' therefore, is essentially to say that an integral element of the assessment was *wrong*. It is difficult, therefore, to see how a failure to anticipate the possible severity of a natural hazard does not constitute

an assessment failure. Extreme natural hazards, after all, are precisely the kind of occurrences that expert risk assessors try to envision in their calculations.¹⁴

ii. Relevance reexamined.

'Relevance claims' feature prominently in the discourse around all nuclear disasters. After Three Mile Island, for instance, authorities worked hard to assure the public that it was an exceptional event with little bearing on nuclear power more generally: a "teething accident" as Socolow (2011) would later characterize it. Indeed, just three years before Chernobyl, the head of the IAEA's safety division confidently concluded that its design made a TMI-type meltdown "practically impossible," when discussing its reactor type (RMBK) in the *IAEA Bulletin* (Semenov 1983: 51); a view that was repeatedly echoed by the soviet authorities. This, of course, was soon disproven. But this did not stop Western observers from making much the same claim again about Chernobyl. After the Soviet disaster Western observers drew an invisible line around the Iron Curtain and claimed it was an exception that said very little about their own reactor designs and assessment practices – "a *Soviet* accident, not a nuclear accident" as many officials would later refer to it. It became routine for Western experts to exclude Chernobyl from their assessment calculations, as could be seen, for instance, when the Indian Nuclear Commission assured visitors to its website that "The statistical risk from living next to an intelligently-designed (*i.e., not RBMK*) nuclear power plant is equivalent to driving 125 feet"[emphasis added].¹⁵

This history of misleading assurances is the appropriate lens through which to view claims about Fukushima's exceptionalism, the logic of which quickly begins to wilt under close scrutiny. Take for instance, the claims of Japanese exceptionalism. Fukushima was a Japanese facility, but its GE Mark-1 reactors were designed in the United States and are still employed in 23 US plants (Marvel & May 2011: 3). This means the design must have been specifically assessed and approved by US authorities on various occasions, as well as by authorities in a range of countries that constructed the same reactors at different times (these include Switzerland, Spain, The Netherlands, Italy and India).

It is not illogical, of course, to argue that assessments that deemed a design safe for the US might have found it unsafe for Japan; especially considering Japan's distinctive seismic activity. Yet the US manufacturer played a leading role in the initial stages of Fukushima's approval, no doubt using the same assumptions and assessment models. The seismic variables would have differed, but the controversies that the accident highlighted about Japanese earthquake models echo ongoing debates and uncertainties in around US seismic standards (eg: McCann 2011). There is every reason, in other

¹⁴ It is worth noting, moreover, that the events in Japan, although unusual, were far from unthinkable. The 9.0 earthquake was large, but not to a degree that should have made it obviously beyond consideration. The area, at the intersection of three fault lines, was long expected to be due a seismic event, and the quake, when it came, was only the fourth largest of the last century. Experts had repeatedly tried to highlight the danger (Perrow 2011b: 47), yet, as the head of Kinki University's Atomic Energy Research Institute told the Wall Street Journal: "The earthquake and tsunami [...] both exceeded [the accepted] engineering assumptions by a long shot" (in Shirouzu & Smith 2011).

¹⁵ Indian Nuclear Commission online. [Accessed 2010; since removed].

words, to imagine that Fukushima's assessment calculations would have been acceptable elsewhere. At the time of the accident, we should remember, Japan had a first-class reputation for managing complex engineering infrastructures. As the title of one op-ed in the *Washington Post* put it, "If the competent and technologically brilliant Japanese can't build a completely safe reactor, who can?" (Applebaum 2011).

Claims that Fukushima is exceptional because assessment standards have changed are similarly unconvincing. It is certainly true that assessment practices have evolved, such that modern assessments differ significantly from those that governed the Mark-1, but it is also true that successive regulatory regimes had periodically reviewed and reassessed the Mark-1 reactors using updated calculative criteria, and declared them to be safe. In this regard as well, therefore, the accident's relevance to wider assessment practices is difficult to dispute.

iii. Compliance reexamined

Claims of exceptional malpractice and error merit more serious consideration. Such arguments also have a long history; error and malpractice being common themes of almost all accident investigations.¹⁶ The formal reports that identify human failures at Fukushima echo those of previous disasters such as the 'Kemeny Commission' investigation into Three Mile Island, which pointed to operator error as the cause. Thereby establishing a longstanding caricature of TMI as, in Socolow's (2011) words, the result of "appalling deficiencies in worker training" (see also: Wynne 1983: 23; Otway & Misenta 1980). Soviet investigations into Chernobyl, meanwhile, similarly identified human error as the primary cause, with the plant director and five other operators being sentenced to long prison terms (Schmid 2011: 20).

It is debatable in both cases whether the specific charges were wholly warranted, as we will see in a later section below, yet errors and malfeasance undeniably do contribute to accidents in complex organizations and there is compelling evidence that they contributed, at least in part, to the Fukushima disaster.

Even if we can blame a nuclear accident on errors or misbehavior, however, this does not necessarily redeem the assessments that claimed the accident was not credible. This is because such redemption would imply that perfect assessments of (and control over) nuclear risks are possible if everyone acted rationally and followed the rules, and, further, it would imply that it is reasonable for assessments to expect everyone to act rationally and follow the rules. Both of these premises are implausible.

Take, first, the idea is that it is realistic for assessments to assume that people will always follow the rules. Accounts of Fukushima, like those of Chernobyl and TMI, invariably treat its apparent human

¹⁶ Many scholars attribute the emphasis on human error in the discourse around accidents as much to institutional expediency as to human deficiency. Perrow (1983; 1999), for instance, argues that institutions try to limit their liability for accidents by leveraging ambiguities in the definition of 'operator error'.

failings as aberrant, but the kinds of misbehavior they document are arguably quite unexceptional. Over the last thirty years a large and sophisticated body of academic literature has emerged exploring how human-error, noncompliance and malfeasance relate to accidents. Vaughan (1996), for instance, speaks of the “normalization of deviance”; Rasmussen (1997) of “migration to the boundary”; and Snook (2000) of “practical drift.” This discourse – of which Silbey (2009) offers a useful overview – sometimes conveys a diffuse impression, often implicit and unintentional, that human behavior is a manageable problem: “something akin to noise in the system” as Silbey (2009: 342) puts it. Yet it is a problem that has proven stubbornly resilient to sociologists’ proscriptions. As Wynne (1983: 23) notes, “There are countless cases where technologies have failed [...] because somewhere in the social labyrinth of their enactment, people have not acted according to the designers’ unrealistic assumptions and faiths.” Indeed, it is difficult to imagine any institution in history in which every rule was followed all the time.¹⁷ In the nuclear industry alone, for instance, there exists an extensive literature documenting regulatory failings in routine operations (see eg: Perrow 2007; Perin 2005). To a seasoned observer, therefore, it should be entirely unsurprising that the media spotlight generated by Fukushima unearthed evidence of malpractice in nuclear plants around the world (eg: Donn 2011).¹⁸ The simple, resilient truth is that people sometimes disobey rules, and when nuclear risk assessments fail to capture this in their calculations then it is their own failing, akin to miscalculating the frequency of earthquakes.

Let us now consider the more intuitive, premise implied by the ‘compliance defense’: that perfectly compliant operators (if they existed, contrary to my argument above) *could*, in principle, always follow the rules. This idea seems unproblematic, yet a growing body of literature suggests it is unrealistic. This is because the idea of ‘perfect rule compliance’ implies ‘perfect rules’ – the idea that complex institutions can realistically offer unambiguous proscriptions for every contingency – whereas successive studies suggest that even the most expansive stipulations need interpretation and cannot relieve workers of having to make some decisions in uncertain conditions (eg: Schmid 2011; Mackenzie 2003). Close accounts of technological work routinely find that work to be *necessarily* and *unavoidably* ‘messier’ in practice than it appears on paper (eg: Langewiesche 1998; Downer 2007; Wynne 1988), with the effect that ‘error’ and ‘non-compliance’ are often ambiguous concepts, which, like ‘failure,’ are open to legitimate interpretation. “Regulations will always be imperfect,” as Perrow (2011b: 46) puts it, “they cannot cover every exigency.” Wynne, for example, speaks of TMI in terms of “...an unresolvable contradiction between the demands of routine operation and occasional abnormal conditions” (1983: 23-7). He argues that the operators responded rationally to the ambiguous rules and information available to them. “To pretend that more competent operators would have avoided the emergency,” he

¹⁷ Christian tradition maintains that the very first two humans on Earth disobeyed their only rule, commanded unto them by God Himself.

¹⁸ It is also important to note, moreover, that nuclear risk assessments cannot discount the possibility that someone will *deliberately* do something self-destructive and pathological. The nuclear industry recognizes this possibility. It studies ‘insider threats’ and the resilience of containment structures to the impact of errant airliners. But it fails to fully incorporate such factors into the assessments that frame policy and public discourse.

writes, “is to imply that complete knowledge of the reactor system can exist and be exploited by competent operators” – something he deems implausible (Wynne 1983: 23).

The vagaries of human performance and the ambiguities of rules are probably quite marginal to risk calculations in most organizational circumstances; which no doubt explains why there is a tradition of neglecting such issues in engineering assessments. At the levels of reliability expected of nuclear reactors, however, – where meantimes-between-failure are expected to be thousands, if not hundreds-of-thousands of years – even highly unlikely compliance issues will unavoidably have a much larger bearing on the validity of an assessment than they would in normal circumstances. In the context of nuclear power, therefore, such issues become enormously significant. As I have elaborated elsewhere (Downer 2011b), engineers need to question their longstanding premises when operating with ultra-high reliabilities.

iv. Redemption reexamined.

Recognizing the irreducible vagaries of human behavior (and the rules governing it) offers a critical perspective on the final “we-found-and-fixed-the-error” redemption narrative. Not least, because it shows that some shortcomings can never be corrected.

A principle shortcoming of the “found-and-fixed” narrative lies in its unrealistic promise of ‘completeness’. Few critics would argue that nuclear assessment practices cannot be improved by learning from past failures, or that improving them is not a worthwhile endeavor,¹⁹ but ‘improved’ is not the same as ‘perfect’. The lessons of Fukushima will lead to better risk assessments, in other words, but there is no way of knowing how *much* better. Even if experts can offer compelling reasons to believe that the errors revealed by the accident have been resolved, therefore, they can offer no incontrovertible reasons to for us to believe that the assessments are now error-free. (A goal that this paper has outlined several in-principle reasons to believe is fundamentally unrealizable.)

Perhaps more purely than any of the other arguments outlined in this paper, therefore, the promise of perfectibility appeals for us to follow Charlie Brown in subordinating experience to hope. It invites us to disregard the history of nuclear regulation, and imagine that this time it will be different. “The perfectly safe reactor is always just around the corner” says Gusterson (2011). We might say the same about the perfectly accurate risk assessment.

Nuclear experts routinely deflect questions about the accuracy of their risk assessments by emphasizing their use of error-bars and generous margins, designed to safely accommodate calculative uncertainties. Such tools are misleading, however, as they are framed by fallible understandings of where, exactly, the uncertainties lie, and offer little defense against challenges that come from ‘outside the box.’ Appeals to

¹⁹ Indeed, the history of nuclear power, like that of all complex systems (Perrow 2011b: 51-2; Downer 2012b; Petroski 2008), is one of continuous learning from accidents and near accidents.

generous error margins misconstrue the uncertainties of nuclear risk assessments by assuming that any errors will be at the margins of their calculations. (Such that, for example, a failure probability estimated to be one-in-ten-million might, at worst, be one-in-nine-million instead.) But accidents in critical systems sometimes happen for reasons that nobody even anticipated, rather than because engineers miscalculated a known variable (Downer 2012b), and when this happens risk projections can be off by orders of magnitude. Even assessments with the most generous error-margins can be upturned by ‘unknown unknowns’: unrecognized events and failure-modes that exist entirely outside the calculations. Nuclear experts are hastily rethinking their tsunami calculations in the wake of Fukushima, for instance, but who is to say that a freak tsunami will be the next under-anticipated event to challenge a nuclear plant. The lead author of a 2011 NRC report on flood preparedness (Perkins et al. 2011), for instance, has claimed that the Oconee plant in South Carolina is almost certain to melt-down if the upstream Jocassee dam were to fail, and, further, that such an event – at Oconee or elsewhere – is hundreds of times more probable than the tsunami that hit Japan (Global Research 2012).²⁰

It is not unreasonable, finally, to maintain that adequately (albeit not perfectly) accurate risk assessments might be achievable if experts were to steadily learn from their failures.²¹ For proponents of this view, however, Fukushima still raises damming questions about the timescales involved, and how willing we are to tolerate the learning process. How many accidents would be required before we could be confident that our assessments were accurate? “The Japanese radiation victims and the dead plant workers will be glad to know that in their disaster lies our salvation” writes Perrow (2011a). He is being arch, of course, but he has a point. Arguing that we should not reject nuclear risk assessments on the basis of a few ‘teething problems’ begs the question of when disasters should truly be definitive. If the Cold War had ended in a thermonuclear exchange, would the ensuing holocaust have invalidated deterrence theory, or merely suggested it needed refinement?

The only fact that Fukushima demonstrates absolutely unambiguously is that devastating oversights can exist in what authoritative experts ardently claim to be rigorous, objective and conservative risk calculations.

4. CONCLUSION

The disaster-punctuated history of nuclear power ought to speak for itself about the limitations of risk assessments, but our narratives obfuscate that history by rationalizing it away. For experience can only

²⁰ He has also formally alleged that the NRC censored his report by “...intentionally mischaracteriz[ing] relevant and noteworthy safety information as sensitive, security information in an effort to conceal the information from the public.” (Global Research 2012).

²¹ The reliability of modern jet-liners, for instance, is built on lessons gleaned from billions of flights that have elicited thousands of accidents. But this was only made possible by an extraordinarily high tolerance for risk in the early decades of aviation, (which would be unrealistic in the nuclear sphere), and a commitment to what I call ‘Innovative restraint’ (which the nuclear industry cannot match.) (Downer 2011a, 2011b).

‘show’ if we are willing to ‘see,’ and the lessons of Fukushima, like those of the accidents that preceded it, will always be opaque to us if our narratives consistently interpret it as exceptional. So it is that even as the dramas of Fukushima linger, and in some ways intensify, the Ideal of Mechanical Objectivity survives with its misleading impression that expert calculations can objectively and precisely reveal the ‘truth’ of nuclear risks. This has critical policy implications.

At a 2012 post-Fukushima enquiry, Japan’s Trade Minister, Yukio Edano, testified that the Japanese nuclear regulator had rejected global standards for disaster response out of fear that implementing them “would undermine public trust” (in Kubota 2012). His disclosure illustrates the truth of Power’s (2011: 29) observation that “a society which seems to manage risk via the intensification of auditing and monitoring, in fact makes itself more vulnerable by damaging the institutional conditions for encountering fundamental surprise.”²²

Irreducible uncertainties exist in even the most rigorous calculations, and we need to recognize this. For what is the purpose of a risk calculation of dubious reliability? It is not useful to say that there is a *one-in-a-million* chance of a meltdown, but an *unknowable* (but nevertheless *meaningful*) chance that that figure is completely wrong because it assumes an erroneous design basis or an implausible model of human behavior. Such a claim offers no compelling reason for policymakers to defer to expert assurances that nuclear accidents will not happen; no justification for them to avoid considering the possibility of accidents when evaluating the costs and benefits of proposed nuclear plants; and no grounds for shirking the high costs of disaster preparation. In other words, it reintroduces all the problems that nuclear regulators tried to escape from with their initial turn to probabilism in the 1960s.

The pervasive idealization of formal risk assessments, which so many narratives of Fukushima reaffirm, both narrows the democratic discussions around nuclear power, and perverts the processes through which it is governed. The false surety it projects allows the cost/benefit projections that frame nuclear decisionmaking to silently discount the evidence of past accidents in the tacit understanding that disasters are somehow *aberrant* and *avoidable* rather than *endemic*. The result, as outlined above, being a deep-rooted institutional reluctance to adequately plan for worst-case scenarios, and an institutional inability to consider whether, as Perrow (2011b: 52) succinctly puts it: “some complex systems with catastrophic potential are just too dangerous to exist, not because we do not want to make them safe, but because, as so much experience has shown, we simply cannot.”

We would be better positioned to govern the atomic age if we could institutionalize the idea that nuclear risk assessments are *contestable judgments* more than they are *objective truths*. Academic and policy debates that treat nuclear risk assessments as ‘established facts’ are likely to come to very different conclusions than those that are willing to grasp the nettle of uncertainty. Uncertainty demands

²² The idealized notion of risk assessment does not even need to *convince* in order to *constrain*, as it shapes the institutional structures that even skeptics must work within. Bureaucracies inevitably limit their scope of action when they formally and authoritatively ‘deny’ technological hazards, for it is difficult – legally, bureaucratically and rhetorically – for policymakers to justify, or regulators to require, the consideration of risks that have been officially declared to be negligible.

‘possibilist’ thinking (Clarke 2005). It favors the ‘precautionary principle’ (Collingridge & Reeve 1986).²³ It might not aid us in making nuclear power safer, but it would more adequately frame the question of whether the costs of trying are too high to bear. A question that must have weighed heavily on Japan’s prime minister in 2011, as he grappled with the grim calculus of catastrophe, and – harboring, he would later confess, “apocalyptic visions of a deserted Tokyo” (in Osnos 2011: 50) – ordered the formation of ‘suicide squads’ to combat the unfolding disaster.

Works Cited:

- Acton, J. & Hibbs, M. (2012) “Why Fukushima Was preventable” *Carnegie Endowment for International Peace*. March 2012. Washington DC. Online < <http://carnegieendowment.org/files/fukushima.pdf> >
- Adelstein, J. & McNeill, D. (2011) "Meltdown: What Really Happened at Fukushima?" in *The Atlantic Wire*. July 02, 2011. < <http://www.theatlanticwire.com/global/2011/07/meltdown-what-really-happened-fukushima/39541/> >
- Amano, Y. (2011) “Statement to IAEA Ministerial Conference on Nuclear Safety” 20 June, 2011. Vienna, Austria. Available online < <http://www.iaea.org/newscenter/statements/2011/amsp2011n013.html> >.
- American Nuclear Society (2011) “Japanese Earthquake/Tsunami; Problems with Nuclear Reactors.” March 3, 2011. Online: < http://209-20-84-91.slicehost.net/assets/2011/3/13/ANS_Japan_Backgrounder.pdf >
- Applebaum, A. (2011) “If the Japanese can’t build a safe reactor, who can?” in *The Washington Post*. March 14, 2011
- Beck, U. (1999) *World Risk Society*. Polity Press, Malden.
- Broder, J. (2011) “U.S. Nuclear Industry Faces New Uncertainty” in *The New York Times*. March 13, 2011. Online: < http://www.nytimes.com/2011/03/14/science/earth/14politics.html?_r=1&hp >
- Clarke, L. & Perrow C. (1996) “Prosaic organizational failure” in *American Behavioral Scientist*; 39 (8): 1040-1056
- Clarke, L. (1999) *Mission Improbable: Using Fantasy Documents to Tame Disaster*. University of Chicago Press; Chicago.
- Clarke, L. (2005) *Worst Cases: Terror and Catastrophe in the Popular Imagination*. University of Chicago Press; Chicago.
- Collier, S. (2008) “Enacting Catastrophe: Preparedness, Insurance, Budgetary Rationalization” in *Economy and Society*. 37 (2).
- Collingridge, D. & Reeve, C. (1986) *Science Speaks to Power*. St. Martin's, New York.
- Dahl, F. (2012) “IAEA: ‘significant’ nuclear growth despite Fukushima” *Reuters*. Friday Mar 16, 2012. Online: < <http://www.reuters.com/article/2012/03/16/us-nuclear-energy-growth-idUSBRE82F0MU20120316> >
- Donn, J. (2011) “US Nuke Regulators Weaken Safety Rules.” *Associated Press* 21 June 2011. Online < <http://readersupportednews.org/news-section2/312-16/6336-us-nuke-regulators-weaken-safety-rules> >
- Donoghue, K. (1977) *Windscale Inquiry Oral Transcripts*. Department of the Environment, London. 1977, Day 24: 14-15.
- Downer, J. (2007) ‘When the Chick Hits the Fan: Representativeness and Reproducibility in Technological Testing’, *Social Studies of Science*. 37 (1): 7–26.
- Downer, J. (2011a) “On Audits and Airplanes: Redundancy and Reliability-Assessment in High Technologies.” in *Accounting, Organizations and Society*. 36 (4): 269-283
- Downer, J. (2011b) “‘737-Cabriolet’: The Limits of Knowledge and the Sociology of Inevitable Failure” in *The American Journal of Sociology*. 117 (3): 725-762
- Duffield, M. (2011) “Environmental Terror: Uncertainty, Resilience and the Bunker.” *SPAIS Working Paper*. No. 06-11. University of Bristol.

²³ There are a few indicators that the deference to idealized mechanical objectivity in the nuclear sphere may be slowly starting to change. As noted above, for instance, nuclear discourse in Germany and Japan is much more equivocal than that of the US or UK; and even the (as of writing) newly-appointed NRC Chairperson, Allison Macfarlane, has recently advocated “Less reliance on probabilistic performance assessments” in assessments involving geological systems (Macfarlane 2012: 165).

- Duhem, P. (1954) *The Aim and Structure of Physical Theory*. Princeton, New Jersey, Princeton University Press.
- Eden, L. (2004) *Whole World On Fire: Organizations, Knowledge, & Nuclear Devastation*. Cornell University Press, Ithaca, (NY).
- Fackler, M. (2011) "Nuclear Crisis Set Off Fears Over Tokyo, Report Says" in *New York Times*. February 27, 2012. Online: < http://www.nytimes.com/2012/02/28/world/asia/japan-considered-tokyo-evacuation-during-the-nuclear-crisis-report-says.html?_r=1 >
- Fuller, J. (1976) *We Almost Lost Detroit*. Ballantine Books; New York.
- Glaser, A. (2011) "After the nuclear renaissance: The age of discovery" in *Bulletin of the Atomic Scientists*. 17 March, 2011. Online: < <http://thebulletin.org/node/8629> >
- Global Research (2012) "Coverup: Risk of Nuclear Melt-Down in U.S. Higher than it was at Fukushima" Center for Research on Globalization. Blog. October 20, 2012. Online: < <http://www.globalresearch.ca/coverup-risk-of-nuclear-melt-down-in-u-s-higher-than-it-was-at-fukushima/5309002> >
- Greene, G. (2012) "Science with a Skew: The Nuclear Power Industry After Chernobyl and Fukushima" in *Japan Focus*. Online: < <http://japanfocus.org/-Gayle-Greene/3672?rand=1325868300&type=print&print=1> >
- Gusterson, H. (2011) "The Lessons of Fukushima" in *Bulletin of the Atomic Scientists*. 16 March 2011, Online: < <http://www.thebulletin.org/node/8625> >
- Harvey, F. (2011) "Nuclear is the safest form of power, says top UK scientist." in *The Guardian*. Tuesday, 29 March 2011. Online: < <http://www.guardian.co.uk/environment/2011/mar/29/nuclear-power-safe-sir-david-king> >
- Holloway, J. (2012) "Despite Fukushima disaster, global nuclear power expansion continues." In *Ars Technica*. Mar. 13. Online: < <http://arstechnica.com/science/news/2012/03/despite-fukushima-disaster-global-nuclear-power-expansion-continues.ars> >
- Ichida, T.; Noguchi, Y.; Itabashi, H. and Fujimori, K. (2011) "TEPCO orchestrated 'personal' donations to LDP" in *ASAHI.com*. October 09. Online: <http://www.asahi.com/english/TKY201110080284.html>
- Kahn, L. (2011) "Is the United States prepared for a nuclear reactor accident?" April. 2011. in *Bulletin of the Atomic Scientists*. Online: <http://thebulletin.org/node/8705>
- Kemeny, J. (1979). "The Need for Change: The Legacy of TMI." *Report of The President's Commission on the Accident at Three Mile Island*. October 1979. Washington, D.C.
- Kubiak, W. D. (2011) "Fukushima's Cesium Spew - Deadly Catch-22s in Japan Disaster Relief" in *Truthout*. Monday 27 June 2011 < <http://www.truth-out.org/fukushimas-cesium-spew-eludes-prussian-blues-deadly-catch-22s-japan-disaster-relief/1308930096> >
- Kubota, Y. (2012) "Japan atomic regulator blocked crisis plan revamp - documents" *Reuters*. Fri Mar 16, 2012. Online: < <http://www.reuters.com/article/2012/03/16/japan-nuclear-idUSL4E8EF3R720120316> >
- Kuznick, P. (2011) "Japan's nuclear history in perspective: Eisenhower and atoms for war and peace" in *Bulletin of the Atomic Scientists*. 13 April. Online: < <http://thebulletin.org/node/8727> >
- Lampland, M. (2010) "False numbers as formalizing practices" in *Social Studies of Science* 40 (3): 377-404
- Langewiesche, W. (1998) *Inside the Sky: A Meditation on Flight*. Pantheon. New York
- McCann, M.W. (2011) "Seismic Risk of a Co-Located Portfolio of Dams – Effects of Correlation and Uncertainty" Paper presented at the 3rd International Week on Risk Analysis, Dam Safety, Dam Security, and Critical Infrastructure Management, Polytechnic University of Valencia, Valencia, Spain, October 17-18, 2011. Online: < http://www.stanford.edu/group/npdp/cgi-bin/drplnpdp/sites/default/files/reports/seismic_risk_of_a_co-located_portfolio_of_dams-effects_of_correlation_and_uncertainty.pdf >
- Macfarlane, A. (2012) "Fukushima Lessons: The Disconnect Between Geology and Nuclear Engineering." In *Elements* 8 (3). : 165-166. Online: < http://www.geomin.unibo.it/struttura/Pag_pers/Braga/macfarlane_elements2012.pdf >
- 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: 249-51
- MacKenzie, D. (2003) "A philosophical investigation into Enron" in *London Review of Books*. Monday May 26.
- Marvel, K. and May, M. (2011) *Game Changers for Nuclear Energy*. American Academy of Arts and Sciences. Cambridge, (MA).

- Monbiot, G. (2011) "Why Fukushima made me stop worrying and love nuclear power." In *The Guardian*. Monday 21 March 2011 Online: < <http://www.guardian.co.uk/commentisfree/2011/mar/21/pro-nuclear-japan-fukushima> >
- [NAIIC] The National Diet of Japan (2012) *The official report of The Fukushima Nuclear Accident Independent Investigation Commission: Executive summary*. Tokyo. Online: < <http://reliefweb.int/node/508828> >
- (NRC) US Nuclear Regulatory Commission (2002) "Nuclear Safety Research Conference 'Notice of meeting'." *Federal Register* Volume 67, Number 199 (Tuesday, October 15, 2002): 63686. [FR Doc No: 02-26166].
- (NRC) US Nuclear Regulatory Commission (2004) "Effective Risk Communication: The Nuclear Regulatory Commission's Guidelines for External Risk Communication." Report NUREG/BR-0308. January. Washington DC.
- (NRC) US Nuclear Regulatory Commission (2011) "Availability and Readiness Inspection of Severe Accident Management Guidelines (SAMGs)" Summary of Observations Temporary Instruction 2515/184. Accession No. ML11154A117 < <http://pbadupws.nrc.gov/docs/ML1115/ML11154A109.pdf> >
- (NYT) New York Times (2011) "In the Wake of Fukushima: If the nuclear industry wants Americans' trust, it needs to embrace reforms, not block them". in *The New York Times* July 24, 2011: SR11. Available online: < <http://www.nytimes.com/2011/07/24/opinion/sunday/24sun1.html> >
- OECD (2010) "Comparing Nuclear Accident Risks with Those from Other Energy Sources" *Nuclear Energy Agency Report No. 6861*. OECD Publications. Paris. Online: <http://www.oecd-neo.org/ndd/reports/2010/nea6862-comparing-risks.pdf>
- Okrent, D. (1978) "On the History of the Evolution of Light Water Reactor Safety in the United States." Available online: < http://fissilematerials.org/library/1978/06/on_the_history_of_the_evolutio.html >
- Onishi, N. and Fackler, M. (2011) "Japan Held Nuclear Data, Leaving Evacuees in Peril" In *The New York Times*, August 8, 2011. Online: < http://www.nytimes.com/2011/08/09/world/asia/09japan.html?_r=2&pagewanted=print >
- Osnos, E. (2011) "The Fallout. Seven months later: Japan's nuclear predicament." in *The New Yorker*. Oct 17, 2011: 26-61
- Otway, H. and Misenta, R. (1980) "Some human performance paradoxes in nuclear operations", in *Futures*, 12 (5): 340-357.
- Otway, H., and von Winterfeldt D. (1982) "Risk Management and Acceptable Risk Criteria," in Kunreuther, H. (ed.) (1982) *Risk: A Seminar Series*. Laxenburg: International Institute for Applied Systems Analysis: 125-143
- Perin, C. (2005) *Shouldering Risks: The Culture of Control in the Nuclear Power Industry*. Princeton University Press. NJ.
- Perkins, R; Bensi, M; Philip, J. Sancakatar, S. (2011) " Screening Analysis Report for the Proposed Generic Issue on Flooding of Nuclear Power Plant Sites Following Upstream Dam Failures" U.S. Nuclear Regulatory Commission Office, of Nuclear Regulatory Research Division of Risk Analysis. July. Online: < <http://pbadupws.nrc.gov/docs/ML1218/ML12188A239.pdf> >
- Perrow, C. (1999) *Normal Accidents: Living With High-Risk Technologies (2nd edition)*. Princeton. New York
- Perrow, C. (1983) "The Organizational Context of Human Factors Engineering" *Administrative Science Quarterly*; 28 (4): 521-541
- Perrow, C. (2007) *The Next Catastrophe: Reducing Our Vulnerabilities to Natural, Industrial, and Terrorist Disasters*. Princeton University Press. Princeton..
- Perrow, C (2011a) "Fukushima, risk, and probability: Expect the unexpected" in *Bulletin of the Atomic Sciences*. 1 April, 2011. Online < <http://www.thebulletin.org/node/8702> >
- Perrow, C. (2011b) "Fukushima and the inevitability of accidents." In *Bulletin of the Atomic Scientists*. 67 (6) 44-52
- Petroski, H. (2008) *Success through Failure: The Paradox of Design*. Princeton University Press; New Jersey.
- Popper, K. (1959) *The Logic of Scientific Discovery*. Basic Books, New York.
- Porter, T. (1995) *Trust in Numbers: The Pursuit of Objectivity in Scientific and Public Life*. Princeton, N.J.: Princeton University Press.
- Power, M. (2011) "Preparing for Financial Surprise" in *Journal of Contingencies and Crisis Management* 19 (1): 28-31
- Quine, W. V. (1960). *Word and Object*. Cambridge, Mass., MIT Press.
- Quintana, M. (2012) "Fukushima Crisis Concealed: Japanese Government Kept Worst-Case Scenario Under Wraps" in *The Asia Pacific Journal*. Jan. 31, 2012. Online: < <http://japanfocus.org///events/view/129> >

- Ramana, M. V. (2011) "Beyond our imagination: Fukushima and the problem of assessing risk" in *Bulletin of the Atomic Scientists*. 19 April, 2011. Online: <http://thebulletin.org/node/8736>
- [RARG] Risk Assessment Review Group (1978) Risk Assessment Review Group Report to the US Nuclear Regulatory Commission. NUREG/CR-0400. October Available online < <http://www.osti.gov/bridge/servlets/purl/6489792-nii8Km/6489792.pdf> >
- Rasmussen, J. (1997) "Risk Management in a Dynamic Society: A Modeling Problem." in *Safety Science*, 27 (2/3): 183-213.
- Reason, J. (1990) *Human error*. Cambridge University Press; Cambridge
- Rip, A (1986) "The Mutual Dependence of Risk Research and Political Context" in *Science & Technology Studies* 4 (3/4): 3-15
- Sagan, S. (2004) "The problem of redundancy problem: why more nuclear security forces may produce less nuclear security." *Risk Analysis* 24 (4): 935-46.
- Sagan, S. (1993) *The limits of safety*. Princeton NJ: Princeton University Press.
- Semenov, B.A. (1983) "Nuclear power in the Soviet Union" in *IAEA Bulletin*, 25 (2): 47-59
- Schmid, S. (2011) "When safe enough is not good enough: Organizing safety at Chernobyl." in *Bulletin of the Atomic Scientists*, 67 (2): 19-29.
- Shirouzu, N. And Smith, R. (2011) "U.S. Sounds Alarm on Radiation" in *The Wall Street Journal*. March 17. Online: <http://online.wsj.com/article/SB10001424052748703899704576203850892933990.html>
- Silbey, S. (2009) "Taming Prometheus: Talk About Safety and Culture" in *Annual Review of Sociology*. 35: 341-69
- Slovic, P. (2012) "The perception gap: Radiation and risk" in *Bulletin of the Atomic Scientists*. 68/67 (3). online: <<http://bos.sagepub.com/content/68/3/67>>
- Snook, S. (2000) *Friendly fire*. Princeton, NJ: Princeton University Press.
- Socolow, R. (2011) "Reflections on Fukushima: A time to mourn, to learn, and to teach." In *Bulletin of the Atomic Scientists*. 21 March, 2011. Available online: < <http://www.thebulletin.org/web-edition/op-eds/reflections-fukushima-time-to-mourn-to-learn-and-to-teach> >
- Stanford, K. (2009) *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford; Oxford University Press.
- Strickland, E. (2011) "24 Hours at Fukushima: A blow-by-blow account of the worst nuclear accident since Chernobyl" In *IEEE Spectrum special report: Fukushima and the Future of Nuclear Power*. November, 2011. Online: < <http://spectrum.ieee.org/energy/nuclear/24-hours-at-fukushima/0> >
- Vaughan, D. (1996) *The Challenger Launch Decision*. Chicago: University of Chicago Press.
- Wald, M. (2011) "Nuclear Plant Safety Rules Inadequate, Group Says" in *The New York Times*. June 15, 2011. Available online: < <http://www.nytimes.com/2011/06/16/business/energy-environment/16nrc.html> >.
- Walker, J. and Cooper, M. (2011) "Genealogies of Resilience: From Systems Ecology to the Political Economy of Crisis Adaptation" in *Security Dialogue* 14 (2).
- Weart, S. (1988) *Nuclear Fear: A History of Images*. Harvard University Press. Cambridge (MA).
- Willsher, K. (2011) "Nicolas Sarkozy makes €1bn commitment to nuclear power" in *The Guardian*. Monday 27 June 2011. Available online: < <http://www.guardian.co.uk/world/2011/jun/27/nicolas-sarkozy-france-nuclear-power> >
- Wingfield, B. & Johnsson, J. (2012) "Scana Receives Approval to Build S.C. Reactors" *Bloomberg Businessweek*. March 30, 2012. Online: < <http://www.businessweek.com/news/2012-03-30/scana-receives-nrc-approval-to-build-south-carolina-reactors> >
- Wynne, B. (1988) "Unruly Technology: Practical Rules, Impractical Discourses and Public Understanding," in *Social Studies of Science* 18: 147-67.
- Wynne, B. (1983) "Redefining the Issues of Risk and Public Acceptance: The social viability of technology" in *Futures*. (February): 13-32.
- Wynne, Brian (2011) *Rationality and Ritual: Participation and Exclusion in Nuclear Decision-Making*. London: Earthscan.
- Youngblood, R; Dezfuli, H; & Siu, N. (2010) "PRA in Design: Increasing Confidence in Pre-Operational Assessments of Risks (Results of a Joint NASA/NRC Workshop)" *Proceedings of 10th International Probabilistic Safety Assessment and Management*

Conference; 7-11 Jun. 2010; Seattle, WA; United States. *NASA Technical Report no: HQ-STI-10-050*. Available online: < <http://naca.larc.nasa.gov/search.jsp?R=20100021223&q=N%3D4294129240%2B4294963045%2B4294177099> >