



Downer, J., & Ramana, M. V. (2020). Empires built on sand: On the fundamental implausibility of reactor safety assessments and the implications for nuclear regulation. *Regulation and Governance*, 0(0), 1-22. Advance online publication. <https://doi.org/10.1111/rego.12300>

Peer reviewed version

Link to published version (if available):
[10.1111/rego.12300](https://doi.org/10.1111/rego.12300)

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

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Wiley at <https://onlinelibrary.wiley.com/doi/full/10.1111/rego.12300> . 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/>

Empires built on sand: On the fundamental implausibility of reactor safety assessments and the implications for nuclear regulation

John Downer¹ and M.V. Ramana²

Forthcoming in: *Regulation & Governance* (2020).

Uncorrected proof.

Abstract:

This paper explores the nature of expert knowledge-claims made about catastrophic reactor accidents and the processes through which they are produced. Using the contested approval of the AP1000 reactor by the US Nuclear Regulatory Commission (NRC) as a case study and drawing on insights from the Science and Technology Studies (STS) literature, it finds that the epistemological foundations of safety assessments are counterintuitively distinct from most engineering endeavors. As a result, it argues, those assessments (and thus their authority) are widely misconstrued by publics and policymakers. This misconstrual, it concludes, has far-reaching implications for nuclear policy, and it outlines how scholars, policymakers, and others might build on a revised understanding of expert reactor assessments to differently frame, and address, a range of questions pertaining to the risks and governance of atomic energy.

Keywords:

expertise; governance; nuclear regulation; nuclear safety; risk assessment.

¹ Global Insecurities Centre, School of Sociology, Politics and International Studies, University of Bristol, Bristol, UK

² Liu Institute for Global Issues, School of Public Policy and Global Affairs, University of British Columbia, Vancouver, British Columbia, Canada

1. Introduction

In April 2011, about a month after the catastrophe at Fukushima Daiichi, the operator of Vermont Yankee nuclear plant — a US reactor of similar design to the lost Fukushima reactors — sued the Vermont State government in federal court (Clayton, 2011). The suit indirectly pitted the state against the Nuclear Regulatory Commission (NRC), over the authority of the latter’s reactor safety assessments. Vermont Yankee was aging, and the state wanted to force its closure on safety grounds when its original license expired the following year. The operator, determined to keep the plant, pointed out that the NRC had recently renewed the reactor’s safety license for a further 20 years, and contested the state’s right to question the regulator’s conclusions.

The case wended through the system until August 2013, when the Court of Appeals eventually sided with the operator. The NRC was the sole authority on reactor safety, the court ruled, and the state could not overrule its determinations (NEI, 2013). The decision ended up being moot, as the operator eventually closed the plant on economic grounds (Wald, 2013), but it raises important questions about the role of nuclear regulators and the authority of their assessments.

Like nuclear regulators around the world, the NRC routinely makes official pronouncements on the safety of the reactors under its jurisdiction. Such pronouncements are based on many factors, but at their heart are Probabilistic Risk Assessments (PRAs), which — insofar as the regulator approves a reactor for US operation — always categorically dismiss the possibility of Fukushima-scale accidents. They justify this conclusion, in essence, by establishing that a reactor’s ‘reliability’ (defined narrowly here as its *mean-time-to-catastrophic-failure*, as opposed to, say, the downtime it requires for maintenance) to be so high — over a million reactor-years between failures is now common³ — that serious accidents can be considered improbable to the point of being implausible. Thus, the NRC considers Fukushima-level disasters in reactors under its jurisdiction to be “hypothetical” rather than “credible” events: theoretically possible, in the same way that giant meteorite strikes or extraterrestrial invasions are theoretically

³ In a recent declaration to a UK regulator, for instance, *Areva*, a prominent French nuclear manufacturer, asserted 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 (in Ramana 2011).

possible, but no more worthy of serious consideration or disbursement (Fuller, 1975, pp. 149–186; Rip, 1986, pp. 7–9; Wellock 2017).⁴

This practice of using reliability-based probability to bound catastrophic reactor accidents out of formal consideration is an unusual approach to risk management, which, in most contexts assesses the *consequences* of a hazard as well as its *likelihood* (Miller 2003: 165). It was developed in the mid-1970s, by the US nuclear regulator (then the Atomic Energy Commission [AEC]), after it became clear that public examination and discussion of major reactor accidents might threaten the political viability of atomic energy (Wellock 2017: 699). Since then, it has become the internationally accepted way of framing the questions of reactor safety (Balogh 1991; Wellock 2017: 709).

This is to say that the contention that reactors are knowably designed to levels of reliability that make catastrophic accidents too unlikely to warrant public consideration has become a foundational premise of the practices, rationalities and logics through which all states manage atomic energy: the ‘governmentality’ of reactors worldwide. It shapes reactor siting decisions; as is evident, for instance, in the location of Indian Point nuclear plant, thirty-five miles north of Manhattan. It shapes civil emergency planning and preparedness; as can be seen, for instance, in the fact that the US mandates evacuation radii, medication stockpiles and fallout-tracking models to cope with minor radiation leaks but not for more serious incidents (Clarke & Perrow 1996; Lochbaum et al. 2014, p.63). And, in myriad ways, it shapes the discourse and scholarship around reactors. When leading environmentalists — from journalists like George Monbiot (2006) to climate scientists like James Hansen (2009) — entreat world leaders to embrace atomic energy in the fight against climate change, for example, then their reasoning draws uncritically (albeit usually implicitly) on formal assurances about reactor safety. When economists explore the costs of atomic energy without pricing-in potential catastrophes (e.g. Ansolabehere et al. 2003), they are doing the same. So too are the many academic studies — by historians (e.g. Weart 1988), psychologists (e.g. Slovic et al. 1982), and social theorists (e.g. Douglas and Wildavsky 1983) — that try to reconcile the ‘objective facts’ of reactor safety (derived from PRA

⁴ This practice emerged in the late 1960s and early 1970s, when US regulators began to worry that studies exploring the implications of meltdowns were threatening the political viability of atomic energy (Rip, 1986, pp. 7–9).

assessments) with the seemingly incommensurate levels of public concern aroused by atomic energy.

A key premise of choices like those outlined above is that official pronouncements about reactor safety (and thus the reliability assessments on which they depend) are highly credible and authoritative: ‘established facts’ rather than ‘qualified opinions.’ The combined economic, social and environmental costs of catastrophic reactor disasters are extreme to the point where they are widely considered to be politically intolerable (Wellock 2017; Downer 2016).⁵ Decisions to ignore such costs when framing emergency plans, making siting choices, performing economic analyses and constructing social theories would be difficult to justify if the validity of reactor reliability assessments — or the idea that their failure-probability is an inherently knowable property of reactors — was in dispute, or seen as legitimately disputable (Watson 1994).

So it is that a deep and abiding institutional faith in the credibility and authority of reactor reliability assessments is etched into all nuclear policies, infrastructures and discourse. The depth of this faith is evident in its resilience. Historically, catastrophic reactor accidents have occurred far more frequently (by orders-of-magnitude)⁶ than official assessments have promised (Cochran, 2011; Ramana, 2011; Raju, 2016; Wheatley et al. 2017), but this has not meaningfully undermined those assessment’s institutional standing. Instead, states have interpreted catastrophes like Fukushima and Chernobyl as unrepresentative ‘one-offs’ with limited bearing on the wider industry’s safety or the ability of experts to assess it (Downer, 2014; Hamblin 2012). Even amid the continuing repercussions of these outsized dramas, the preponderance of discourse, analysis and scholarship around atomic energy remains committed to the idea that reliability is a ‘known’ (and knowable) property of reactors. Declarative claims about reactor failure probabilities still occupy the high ground in public discussions of nuclear energy: legitimating policy and shaping public opinion (Perrow 1999; Ramana & Kumar 2014).

The principle argument of this paper is that official pronouncements about catastrophic reactor accidents — for all their institutional resilience and standing — are much less credible than is

⁵ Such costs, which are systematically downplayed and obfuscated, are widely under-appreciated (see Downer 2016).

⁶ The exact numbers vary depending on interpretive choices — on whether Fukushima counts as one meltdown or three, for example — but are damning even when made with conservative assumptions. Raju (2016) offers an uncommonly sophisticated analysis of the question, and comes to an unambiguous conclusion. “It is clear that the results of [nuclear safety assessments] are untenable in the light of empirical data” he writes (p. 56).

commonly supposed. And that the authority afforded to them, and to the experts who wield them, is unwarranted. In the sections that follow we will argue that the discourse around reactors and the structures through which they are governed are premised on a fundamental misconstrual of the nature and provenance of the assessments on which those pronouncements depend. Section 2, below, will explain the logic of this argument. It draws on insights from the Science and Technology Studies (STS) literature to explore the authority of engineering expertise in general, and to argue that reactor reliability assessments have an exceptional relationship to that authority. On this issue, if on few others, we contend, it is reasonable for STS scholars to take a strong normative position on the (in)credibility of an expert knowledge-claim. Having made this argument in principle, Section 3 then illustrates and evidences it in practice. To this end it explores two controversies that arose during the NRC's assessment of a US reactor design, the 'AP1000,' which was certified as being unprecedentedly safe. Finally, having established and evidenced an STS argument for normativity with regard to reactor safety, Section 4 then invokes that argument to sketch some normative observations. It outlines how scholars, policymakers and others might build on a revised understanding of expert reactor assessments to differently frame, and address, a range of questions pertaining to the risks and governance of atomic energy.

2. The Limits of Certainty

2.1 Engineering and Authority

To understand why the authority vested in reactor reliability assessments is unwarranted, it is important first to recognize the basis of that authority.

On one level, the credibility afforded to reactor assessments — like those of all expert findings — rests on public trust in the integrity and ability of the individuals and institutions that produce them (Miller 2003: 168).⁷ This is straightforward: assertions carry less water when proffered by actors who are widely considered to be corrupt or incompetent.

⁷ This is rational, for where public policy must hinge on inscrutable claims that non-specialists cannot judge directly, it is a practical recourse to judge the experts attesting to those claims instead (Shapin, 1995).

It is not our intention in this paper to imply that reactor experts are corrupt or incompetent, but authority cannot stand on the basis of integrity and ability alone. Stock-market predictions, for example, are always met with a measure of caution, even when made by seemingly able and upstanding economists. A more fundamental basis for the credibility of reactor assessments, therefore, lies in an understanding of their nature and provenance: a belief about their fundamental amenability to expert analysis. And it is on this that our argument is focused.

Common understandings of the nature and provenance of reactor assessments are grounded in the fact that reactor safety is fundamentally conceived to be an *engineering problem*. As we have seen, the NRC's safety assessments have 'reliability calculations' at their core: they examine the probability that reactors will fail catastrophically.⁸ This is significant because modern societies have long afforded engineering assertions (together with those of 'hard' scientists) more credibility, authority, and institutional standing than most other expert assertions (Miller 2003: 188).

There are various logical justifications for this discrimination, which, loosely speaking, we might reduce to two broad rationales: one principled, the other more practical.

The principled rationale for conferring authority on engineering knowledge lies in a common understanding of its epistemology. This is to say that societies see engineering claims as unusually credible because they conceive of engineering as a rigorous and objective discipline, which produces knowledge grounded in calculation and measurement rather than subjective opinion or judgement (Wynne 1988). Engineering has a few peculiarities in this regard, which we will discuss below; for most purposes, however, it shares this image with the hard sciences — such that it often makes sense to group them together as 'technoscience.' Technoscientific knowledge production is widely assumed to be driven by strict methodologies, which, when performed diligently and correctly, allow experts to interrogate the world accurately and definitively. (And, when not carried out diligently and correctly, produce errors that can, and will, be identified by a community of expert peers.)

⁸ The logic of construing the safety of reactors in this way is problematically reductive (see Downer 2015), but it has become orthodox in formal discourse about atomic energy and we will not take issue with it here.

This intuitive and common conviction goes by many names, but we will follow Collins (1985) in calling it the “canonical rational-philosophical model” of technoscientific knowledge. Its role in establishing the credibility of expert assertions not only stems not only from its promise of objectivity and certainty, but also from its ability to supersede the need to trust in the qualities of specific individuals and organizations. It can perform this latter function because variables like ‘insight’ and ‘judgment’ become less important when experts are seen as merely applying rigorous rules and formula, while questions of ‘honesty’ and ‘neutrality’ become less important when expert conclusions are understood as being objectively verifiable by third parties. So insofar as this canonical model of knowledge holds sway, it is possible to replace trust in people with trust in process (Porter 1995; Gross 1990; Fox-Keller 1994; Miller 2003).⁹

The more practical, and less remarked upon, rationale for disproportionately conferring authority on engineering claims (and, with caveats, technoscientific claims in general) lies in the fact that empirical, lived experiences frequently affirm the efficacy of those claims.¹⁰ Put simply, engineering *works*. Cell phones connect; bridges don’t collapse; and cars start. Our artifacts might not always be perfect, but longstanding experience suggests that nothing fundamental in the underlying epistemology of their creation makes them inherently incapable of perfection, or that expert claims about their properties are deserving of systematic skepticism. No modern polity would lightly ignore an official engineering determination that a proposed bridge would not support its own weight, or that a proposed jetliner was too underpowered to fly. And any that did would quickly learn some salutary lessons. Engineering claims of this nature are so demonstrably credible and efficacious that state bureaucracies, courts, and other rule-making bodies have long been disciplined to accept them as established facts, lest they ‘politicize’ them with skepticism.

Both the rationales above are intuitive. It is easy to imagine why expert engineering assertions are afforded considerable administrative authority on issues pertaining to technological capabilities. It is also easy to imagine why reactor safety assessments would naturally inherit

⁹ As Porter (1995) observes, public bodies often use a trope of impartial professionalism and quantitative assessment to justify their decisions “Arbitrariness and bias are the most usual grounds upon which such officials are criticized.” He writes. “A decision made by the numbers (or by explicit rules of some other sort) has at least the appearance of being fair and impersonal. Scientific objectivity thus provides an answer to a moral demand for impartiality and fairness. Quantification is a way of making decisions without seeming to decide” (Porter, 1995, p. 8).

¹⁰ This, in turn, often reifies the canonical understanding of their epistemology.

this authority. Superficially at least, the nature and provenance of reactor assessments look equivalent to those of any other engineering claim about technological performance or capability. If anything, they look more impressive than most: the levels of expertise, effort and expenditure involved in producing them all clearly being exceptional. They differ from most other such claims, in that empirical experience cannot speak directly to their efficacy. (Given the small number and diverse designs of reactors in operation, it would take thousands of years to gather the empirical data required to statistically demonstrate the ultra-low failure probabilities that experts ascribe to them [Raju, 2016]). But they are proffered by the same community of experts (i.e. engineers) who speak with demonstrable credibility in a vast number of other circumstances where the efficacy of their claims are visible. It should not be surprising, then, that audiences of all kinds treat reactor assessments no differently than the many other engineering assertions that inform their reasoning and decisionmaking.

It is important to understand, therefore, that the nature and provenance of nuclear safety assessments *are* significantly different from those other engineering calculations. And in ways that directly affect their credibility. To recognize this difference, and its necessity, it helps to consider the STS position on certainty and its relationship to expert authority.

2.2 STS and Authority

STS has long had a complicated and nuanced relationship to questions of technoscientific authority.

On one hand, the discipline is all but defined by its rejection of the ‘canonical model’ of technoscientific knowledge, on which, as outlined above, the authority of its claims is often thought to rest. STS was built on prior work — by logical philosophers such as Wittgenstein (2001 [1953]) and Feyerabend (1975), and philosophically-minded historians such as Kuhn (1996 [1962]) — who, beginning in the mid-Twentieth Century, began to systematically explore fundamental limitations of the scientific method. By demonstrating, in different ways, the formal impossibility of the ‘ideal experiment’ or ‘perfect proof,’ these scholars showed how even the most rigorous and objective knowledge-claims necessarily rest on subjective

judgments and unprovable assumptions. Their work gave way to a new orthodoxy in epistemology, often referred to as a 'constructivism' (as opposed to 'positivism').¹¹

Constructivism made STS possible. Scholars used its insights to get inside technoscientific knowledge and examine the interpretive labor involved in forging discrete and meaningful 'facts' from the unruly fabric of reality (e.g. Bloor 1976; Kusch 2012; Stanford 2009; Sismondo 2010; Collins 1985; Latour 1987). Over a period now spanning decades, STS scholars have explored practical manifestations and ramifications of dilemmas that philosophers had previously identified in principle. This literature has many facets, but especially pertinent to the argument that follows are a series of studies illustrating how engineering tests and models imperfectly represent the real situations they are being used to explore, and how such imperfections can become a source of uncertainty, contestation and value-laden judgement (e.g. Wynne, 1988; Pinch, 1993; MacKenzie, 1996; Downer, 2007). In more depth than we can afford here, these studies explain why no STS scholar maintains (or could maintain) that expert engineering claims, reactor reliability assessments included, are ever inviolably objective or rule-governed.

On the other hand, however, STS's commitment to constructivism has never led it to deny the *efficacy* of technoscientific knowledge. This is not a contradiction. The conviction that 'facts' always contain irreducible uncertainties and subjective judgements is not incompatible with the belief that modern medicine controls infections more effectively than the rites of faith-healers. STS scholars freely concede, in other words, that the epistemological limitations of proof, while fundamentally unavoidable, are often inconsequential in real-world circumstances where practical questions of authority and credibility are at stake. Contrary to the suggestions of some of their critics (e.g. Gross & Levitt, 1998), none would endorse a 'post-truth' society, where the claims of technoscientific experts are shown no institutional deference over those of laypeople.

Even though the STS literature speaks eloquently to the paucity of the 'canonical' understanding of knowledge, therefore, it rarely challenges the practical authority of technoscientific experts. This is especially true in respect to the authority afforded to engineers;

¹¹ Constructivists don't deny there are ontological truths about the world, as critics sometimes suggest. They simply hold that all access to those truths is mediated by observation and experiment, which necessarily contain fundamental ambiguities and require subjective judgements: an irreducibly social component to every 'fact' (e.g. Bloor 1976; Collins & Pinch 1993).

whose claims are sometimes seen as being especially resistant to constructivist misgivings about certainty. Constant (1999) gives a good account for why. He concedes — per the constructivists — that absolute truth is an unrealizable goal in engineering,¹² but points out that engineers are generally more concerned with ‘application’ than ‘discovery,’ and inherently less interested with what is ‘true’ than with what ‘works’ (1999: 352-5).¹³ In an engineering context, he argues, theories, tests and models need not be perfect in order to be practical; and their practicality is easily observed. Thus, he concludes, the limitations of proof rarely have practical implications for the credibility of engineering assertions. “[T]here is profound difference between reliable knowledge and unreliable stuff,” he writes, and, despite the concerns of modern epistemologists, “most of our stuff more or less works most of the time” (1999: pp.335; 331).

Committed to the inherent fallibility of technoscientific knowledge but unwilling to reject its practical efficacy, STS has long walked a fine line on questions regarding expert authority. An early tenet of the discipline was that scholars should remain agnostic on the ‘truth’ or ‘falsity’ of the knowledge they explore: seeing both conditions as ‘outcomes to be explained’ as opposed to ‘explanations for outcomes’ (Bloor 1976). In keeping with this stance — usually referred to as ‘symmetry’ — the discipline had a longstanding methodological tradition of not trying to intervene normatively in contemporary technoscientific debates; scholars would rarely opine on the credibility of specific claims or the authority of specific experts.

This position has evolved over the last decade, largely due to the rise of ‘post-truth’ phenomena such as the ‘antivax’ and ‘climate denial’ movements. Concerned that their deconstructions of the scientific method were being understood as sympathetic to antiscientific projects, STS scholars started exploring ways to reconcile constructivism with a more normative position on expert authority. This has not been an easy or uncontentious project, given the discipline’s longstanding commitment to symmetry, and has led to complex and occasionally acrimonious internal debates (e.g. Latour 2004; Lynch 2017; Collins & Evans 2002; Rip 2003; Wynne 2003).

It is not our intention in this paper to engage deeply with these debates. To do so meaningfully would require considerably more space than we can allocate. We do, however, wish to offer one

¹² All our “engineering facts” are “immutably corrigible, hypothetical and fallible,” he writes (Constant 1999: 352-5).

¹³ Newtonian mechanics, for example, remain perfectly functional for most engineering applications despite Einstein having contested their veracity.

modest — but we would argue significant — observation with bearing on the issue. This is to point out that, unlike most technoscientific knowledge-claims, constructivism unproblematically supports a strong normative position on the credibility of reactor safety assessments.

The following section outlines why.

2.3 STS and Reactor Safety

To understand why STS can afford to take a normative position on claims about the likelihood of catastrophic reactor accidents, it is necessary first to recognize that they are essentially claims about *certainty*: a matter on which constructivism has always maintained a strong position. To say this is not simply to say that societies act as if they are certain those claims are correct. This is part of it, but there are other instances where societies bet lives and livelihoods on the correctness of technoscientific knowledge-claims, and we would contend (in the manner of Constant [1999] above) that most are often relatively unproblematic.¹⁴ It is to say, rather, that predictive assessments of a reactor's reliability are, *constitutively*, expressions of confidence (i.e. in its future failure behavior). We will speak more of this distinction below. For now, however, consider the dilemma that would arise if experts were to assert they were 99.9999999% certain a reactor will not fail over a given period, but only 80% certain that this number is correct. Such a statement would be almost nonsensical, as the second variable should already be implicit in the first.¹⁵

Admittedly again, other engineering assertions have this property — claims about the structural integrity of buildings, for instance — and we would again argue that most are relatively unproblematic on a practical level. The distinctiveness of reactor safety assessments, we contend, stems from the combination of the property above with a range of further

¹⁴ When states add fluoride to public drinking water, for instance, they are evincing similarly high levels of confidence in expert claims about its toxicology.

¹⁵ If experts could test the failure rate of a reactor empirically from service data, then it would become possible to say it has a level of reliability that is independent of expert confidence in its accuracy. Insofar as such tests are impossible, however, then a measure of confidence will always be a constitutive element of any reactor reliability assessment.

considerations. None of these are wholly distinctive by themselves, but, taken together, they make those assessments uniquely sensitive to constructivist arguments about the limits of proof.

The first such consideration is the fact that engineers must infer the reliability of reactors from first principles, without recourse to statistically-significant service data. (Both because the performance of reactors must be established as a precondition of their operation, and because reactors never accrue enough service data to statistically establish the performance demanded of them.) This is important because in most engineering contexts, reliability is an *actuarial* property of systems: experts examine how often a system fails in operation and then make extrapolations from that data about its future performance. Assessing a system's reliability in this way always involves some fundamental uncertainty, as data are always imperfect, and extrapolations always involve subjective judgements about the equivalence between past and future conditions. But again, *per* Constant (1999) above, we would argue that such uncertainties are rarely determinative for most practical purposes. Assessing a system's reliability *without* recourse to significant service data, by contrast, introduces much greater levels of uncertainty. In these circumstances, experts must extrapolate a system's failure-behavior entirely from their understanding of its functioning: inferring it from (often test-based) failure data on the system's component elements, using elaborate theoretical models of how those elements will interact, as apart of a system,¹⁶ under predicted conditions (Miller 2003: 177). And insofar as their understanding of that system's functioning is necessarily imperfect, the reliability they extrapolate from it might be wrong (without this error becoming visible until it results in a catastrophe). This need to infer failure rates from function rather than from service data makes constructivist misgivings about proof a much more practical concern. It greatly expands the number of representations — tests, theories, models — on which reliability-predictions rest, while simultaneously untethering those claims from any empirical feedback.

A second consideration, which compounds (and, to some extent, creates)¹⁷ the dilemma above, is simply the extraordinarily high levels of reliability required of reactors. We saw above that reactor reliability assessments can be understood as expressions of confidence (in the system's future failure behavior). We also saw that experts must derive that failure behavior from their

¹⁶ Modeling system behaviors includes calculating the efficacy of measures like redundancy and defense-in-depth.

¹⁷ The absence of useful actuarial data flows from this, as does the untestability of conclusions.

understanding of each reactor's functioning. It follows from these two conditions that claims about extreme reliability imply commensurately extreme confidence in the 'correctness' of expert understandings of the system (and thus in the representativeness and accuracy of all the tests, models and theories that underpin that understanding). Not only are reactor reliability assessments unusually dependent on representations, in other words, they are unusually dependent on those representations being *perfect*. For when predicting potential failure-behavior over billions of hours of operation, even the smallest uncertainties about the accuracy of a theory, the completeness of a model or the representativeness of a test become significant.¹⁸ Contra to Constant's (1999) portrayal of engineering as a pragmatic discipline, this makes the task of establishing reactor reliability more akin to a search for 'truth' than a search for 'utility.'

A third and final consideration is that reactors would be inherently 'difficult' technologies even if they weren't required to be extraordinarily reliable (and knowably so). Unlike, for instance, bridges, they are extremely complex, tightly-coupled systems, consisting of many interdependent social and technical elements that interact in non-linear ways (properties with implications that Perrow [1999] explores at length). They operate with uncommonly extreme temperatures and pressures. They harness rare materials like enriched uranium, and imperfectly understood phenomena like radioactivity. They operate over extremely long timeframes, making their designs (and assessments of those designs) dependent on challenging and often untestable theories about future conditions (regarding a location's climate, for instance, or seismology), and about long-term system behaviors (regarding corrosion, for instance, or operator performance). And they must actively negate the effects of atomic fission: a powerful and dangerous process that is inherently unforgiving of failure and requires the system do constant work simply to remain stable.

These considerations compound each other. In combination, they imply that constructivist critiques of the 'perfect proof' should *matter* to the practical efficacy of assertions about reactor safety, even if those critiques have little bearing on most other engineering assertions.¹⁹ Decades

¹⁸ It is also worth noting in this context that reliability is unusual in being a *negative* property of artifacts. This is to say that — unlike mass, density, or almost any other engineering variable — it denotes an 'absence' (i.e. of failure). This is a subtle but significant distinction, because absences are difficult to demonstrate empirically with tests and models (see Popper 1959).

¹⁹ For more on how matters of 'scale' — e.g. of timeframes — can have disproportionate implications for logic of risk analyses, see Jasanoff (1993: 126-7).

of epistemology — not to mention common sense — speaks to the impossibility of formally calculating, via myriad representations and abstractions, the future failure behavior of a highly ‘difficult’ socio-technical system, over a decades-long timeframe, to an extreme degree of confidence. The task, simply requires too many judgments to be made with too much perfection to be plausible. The truth of this should be self-evident.

Given that the challenges of ‘knowing’ the reliability of a reactor are closely aligned to the challenges of ‘designing’ it to be reliable, moreover — the key to both lying in understanding the system itself — it follows that reactors should be unreliable for the same reason that their reliability is unknowable. This relationship is itself unusual and worth noting. In most contexts, difficulties of expert assessment have little bearing on the phenomena being assessed. (Meteorologists’ inability to accurately predict the weather in London ten weeks from now, does not affect the likelihood of it raining.) Engineers’ inability to predict a reactor’s failure-behavior, by contrast, does affect the likelihood of that reactor failing. The knowledge that would have informed the prediction would have also have informed the design. So unlike the meteorologists, they must grapple with a problem of recursion: any cause they have to believe their reliability assessments are uncertain is also cause to believe that their reactors are unreliable.

Let us take stock. Expert understandings of reactors and their failure behaviors represent unfathomably vast tapestries of knowledge-claims: about the diligence of operators; about the long-term fatigue characteristics of specific materials in varying configurations and environments; about the seismology of their locations; about climate change and its implications for future flood requirements; and much, much else besides. Any errors in these knowledge-claims — arising, perhaps, from incomplete model variables, unrepresentative test conditions or imperfect theoretical assumptions — represents a potential source of unexpected failure, not only of reactor assessments, but of reactors themselves. In light of this, it would be remarkable if reactors were truly as reliable as their manufacturers claim and their regulators attest. And even if (in a hypothetical world where we build thousands of near-identical reactors) millions of years of cumulative service-data should one day prove them both correct — by demonstrating, actuarially, that experts had built ultra-reliable reactors and accurately predicted the true

frequency of their failures — those experts would still have been misleading to have asserted they knew the correctness of their predictions with confidence at the time they were made.

In this context if in few others, therefore, it is reasonable and necessary for STS scholars to take a strong position on the validity of a contemporary knowledge-claim. The assertions that nuclear regulators make about the improbability of catastrophic reactor accidents are fundamentally at odds with the way the discipline understands knowledge as being dependent on subjective judgements.

Since arguments from epistemology can often seem conjectural, the following section will ground this conclusion in an example. It illustrates the dilemma above by looking at the NRC's certification of a new reactor 'type': the Westinghouse 'AP1000.'

3. Assessing The AP1000

3.1 *Uniquely Safe*

In fulfilling its mandate to police (and testify to) the safety of US nuclear plants, the NRC evaluates and, if satisfied, 'licenses' generic reactor designs. This process, known as 'design certification,' involves an exhaustive analysis of the proposed reactor, wherein the applicant (i.e. the manufacturer) provides volumes of data for the NRC to evaluate, much of it derived from tests and analyses pertaining to potential vulnerabilities (NRC, 2009, p.18). If the regulator is persuaded by the reasoning and conclusions of this data it issues a 'generic design approval,' thereby endorsing the applicant's safety claims and approving its design for sale (contingent on further site-specific approvals). In 2005, the Westinghouse Electric Company's AP1000 reactor became the first to be granted such an approval from the NRC since the Three Mile Island accident in 1979.²⁰

Conscious of Three Mile Island's long shadow, Westinghouse designed the AP1000 with an emphasis on safety. According to its regulatory filings, the AP1000 is an extraordinarily reliable design, with a 'core melt frequency' of 4.2×10^{-7} per reactor-year, and a 'large release frequency'

²⁰ Its generic certification was originally issued in December 2005, and then reissued in December 2011 to reflect a modified design.

of 3.7×10^{-8} . This puts their expected probabilities (per reactor) at around once-in-every-2.4-million-years, and once-in-every-27-million-years, respectively (Sutharshan et al., 2011, pp. 297–298). In other words, Westinghouse claims that the US could expect to operate an AP1000 for millions of years before experiencing a Fukushima scale catastrophe.

Unsurprisingly perhaps, some experts remain skeptical of these numbers, the processes by which they were derived, and the approvals they legitimate (Feldman, 2010).²¹ To understand their concerns, however, it is necessary to first understand something about the distinctiveness of the AP1000's design.

3.2 PCCS

An essential basis for the unprecedented safety assertions made about the AP1000 is its claim to being what some nuclear officials have described as the first “passively safe” reactor (Lyons 2016).²² In the parlance of nuclear engineering, safety systems are broadly classified as ‘active’ or ‘passive’ by virtue of a range of characteristics. These include, but are not limited to: the number of mechanical parts they utilize; the amount of input they require from an ‘intelligent’ controller (human or otherwise), and the extent to which they require external power (IAEA, 1991, p. 10). The ‘passiveness’ of any given design is always open to interpretation,²³ therefore, but the AP1000's claim to uniqueness in this regard lies in the design of its emergency cooling system.

Emergency cooling systems, somewhat self-evidently, are designed to keep reactors from overheating in the event of an emergency situation involving a loss of primary coolant. In traditional reactor designs these systems involve extensive networks of pipes, pumps and valves, and require electrical power to function. Westinghouse designed the AP1000's “Passive

²¹ Even the NRC's own Chairman, Gregory Jaczko, voted against his agency's license for the new reactors on safety grounds. He was concerned that they carried no binding commitment to implement changes in federal requirements arising from the NRC's post-Fukushima work (Abernethy, 2012).

²² The ‘AP’ in its name stands for ‘Advanced Passive.’

²³ All reactors rely on passive safety to some degree. The most notable structure in most nuclear plants, the containment building — an inert concrete and metal edifice enveloping the reactor-vessel — is straightforwardly a passive safety system. The IAEA distinguishes four different degrees of “passiveness” to which a system might aspire, each corresponding to the number of pre-determined criteria it embodies (IAEA, 1991).

Containment Cooling System" (PCCS) without most of these components; designing it to operate without external electrical power (albeit for a limited time), by utilizing a combination of convection, conduction, evaporation and gravity.

The PCCS is tightly integrated with the AP1000's containment structure. As in most reactors, this structure consists of two parts: an inner 'containment vessel' made of a welded steel dome, and an outer 'shield building' made of reinforced concrete. Unlike other reactors, however, the AP1000's shield building supports a large water storage tank in the roof, and has openings at the side and top. These features are key elements of its cooling mechanism. Should normal cooling be lost, valves in the storage tank are designed to open, allowing the water to pour onto the hot steel of the containment vessel below. The water would cool the vessel, and, as it evaporated, would cause convection in the space (or 'annulus') between the vessel and the shield building. This convection would draw cool air into the building through the side openings, which would absorb more heat, rise, and exit via the opening at the top, thereby creating a 'chimney effect' that would further accelerate the convection. In theory, this process — illustrated in a diagram below, from the International Atomic Energy Agency [Figure 1] — would keep the vessel safe for as long as there was water, with no external power required.




FIGURE 1. Source: IAEA, Status report 81 - Advanced Passive PWR (AP 1000) <https://aris.iaea.org/PDF/AP1000.pdf>

Due to its limited water, the PCCS cannot function indefinitely without electricity or human intervention,²⁴ but Westinghouse contends that harnessing gravity and convection in this manner confers significant safety benefits. This claim is intuitive. A major weakness of the Fukushima reactors was that flooded generators could not provide electricity to drive emergency cooling. It is not inevitable, however. As the IAEA (1991, p. 10) has put it, “[p]assive safety is not synonymous with inherent safety or absolute reliability.” There are various reasons for caution,²⁵ but probably the most general of them pertains to the newness of the design. The nuclear industry has valuable operational experience with active safety systems, which have been the backbone of its safety engineering for decades. This experience, while insufficient to demonstrate statistically the performance the systems require, has nevertheless offered many insights regarding their functioning and performance.

By eschewing a mature safety philosophy, in other words, the AP1000 is inescapably more *conceptual* than other reactors.²⁶ The ramifications of this are reflected in doubts and criticisms that arose about the design during the reactor’s certification. In the sections that follow we will sketch the details of two prominent lines of such criticism, both of which, in different ways, relate to the assessment of the shield building and the radical cooling system it accommodates.

The first, which is relatively brief, involves the shield building’s ‘open’ design.

3.3 Convection and Corrosion

As outlined above, the AP1000’s novel cooling system requires that its shield building have openings at the sides and top to allow for airflow. These openings are a significant departure from traditional shield building design, however, and some critics have argued that the NRC has misconstrued their ramifications. Probably the most prominent voice in this regard is that of a former nuclear reactor operator and industry executive who turned whistleblower in 1990,

²⁴ Under ideal conditions, the tank contains enough water to cool the containment for 72 hours: long enough to create a useful window for emergency workers but not to render it safe (Lyman, 2011). (Fukushima was without power for longer than 72 hours; although there were people on the site providing water within this time.)

²⁵ Forces like gravity and convection are generally less powerful and more variable than ‘active’ forces like electricity, for instance, and some experts have argued that this should be a cause for concern (e.g. EPRI, 2007, pp. 1–2).

²⁶ For this reason, France’s Institute for Radiological Protection and Nuclear Safety (IRSN), maintains it is more difficult to assess the performance of passively safe systems (IRSN, 2016; WNN, 2016).

Arnold Gunderson. In two reports on the AP1000 (Gundersen, 2010b, 2010a), Gunderson has argued that the openings undermine the integrity of the reactor's containment by making radioactive leaks more consequential and, simultaneously, more probable.

The building's openings make radioactive releases more consequential, Gunderson contends, because they make it incapable of mitigating any leaks that might occur due to corrosion or cracking of the containment vessel. Traditional shield buildings serve a dual function; they insulate the reactor from the world by protecting it from external trauma (tornado debris, for instance), and they insulate the world from the reactor by acting as a final barrier against any accidental radioactive release from the primary containment (Ramsey & Modarres, 1998, p. 145; IAEA, 2004, p. 1). The openings in the AP1000's shield building stop it from performing this latter function. (Gunderson [2010b, p. 3] argues that the 'chimney' effect they are designed to foster would actually accelerate the spread of unfiltered radioactive materials into the environment.) The AP1000's open design makes its safety especially dependent on the integrity of its steel primary containment vessel, therefore, but Gunderson argues that the same feature also exposes that vessel to an unusually moist, corrosive environment, thereby making it more prone to failure than equivalent structures in other plants. The shield building's design thus makes the primary containment vessel simultaneously more vital and more vulnerable, he concludes, and this is not adequately reflected in the safety analysis.

Gunderson is an outside observer to the certification process, and neither the NRC nor Westinghouse were formally required to address his misgivings. When confronted by reporters, however, Westinghouse did not directly contest his main points: that the 'open' shield building might exacerbate corrosion and facilitate any radioactive releases. Instead it asserted that such concerns were misplaced because maintenance programs would ensure that any corrosion "... would be readily identified and corrected during regular inspections well before it could in any way become an issue" (Feldman, 2010).²⁷ The NRC agreed, finding that the reactor's planned maintenance program "...is acceptable and is expected to ensure against undetected corrosion of the CV [containment vessel] pressure boundary" (Harki, 2011, p. 5).

²⁷ It also maintained that the vessel had unusually thick steel, which would help mitigate the issue.

Gunderson disputes this maintenance assurance on two grounds. Firstly, he contends that the AP1000 has uncommonly difficult maintenance requirements because it has a number of features — such as airflow-channeling metal ‘baffles’ — that are unusually likely to collect corrosive moisture, and, at the same time, are unusually difficult to inspect for corrosion (Gundersen, 2010b, pp. 1–2). Secondly, he observes that inspections in the past have not always readily identified corrosion even in traditional reactors. “Recent experience with the current generation of nuclear reactors” he writes, “shows that containment corrosion, cracking, and leakage are far more prevalent and serious than anticipated by the [NRC]” (Gundersen, 2010a, p. 1). He substantiates this claim by highlighting known occasions when maintenance protocols have failed to detect dangerous corrosion in the past. These include numerous examples of vessels that were found to be corroded to a degree that would have prevented them from performing their primary emergency function,²⁸ and even several ‘near-misses,’ where corrosion was found to have almost caused a direct failure (Gundersen, 2010a; NRC, 2010a).²⁹ A significant inner-containment breach due to corrosion and inadequate maintenance inspections “is not a low probability event,” he concludes (2010a, p. 3); and, given the open shield building, it is an event that regulators should have taken more seriously.

This is essentially where the disagreement rests and is likely to remain. Neither Westinghouse nor the NRC are formally obliged to recognize or rebut Gunderson’s criticisms, whatever their merit. Westinghouse stands by its safety calculations, with their ahistorical expectations of the AP1000’s maintenance program. The NRC concurs. And the matter stands.

3.4 Modularity and Ductility

The second, slightly more substantial, line of concern comes from within the certification process itself. It is more extensive because it based on objections raised by an NRC official — John Ma — who was directly involved in the AP1000’s assessment. Ma, a structural engineer, had been with the NRC since the agency was formed in 1974, and was well respected, having

²⁸ Salem Nuclear Plant in New Jersey, for instance, where a 2009 inspection found late-stage corrosion that had accumulated because the area was considered ‘inaccessible.’

²⁹ Beaver Valley Nuclear Plant in Pennsylvania, for example, where, in 2009, corrosion was found to have penetrated the entire metal surface of the containment vessel before being discovered. Also, similar incidents at D.C. Cook nuclear plant in Michigan (in 2001), and Brunswick nuclear plant in North Carolina (in 1999) (Gundersen, 2010a, p. 5; 2010b, p.5).

won many awards, including the NRC Administrative Service Excellence Award. He outlined his misgivings to the AP1000's design approval in a formal "non-concurrence" report (NRC, 2010b) to which the regulator was obliged to respond.

The principle concerns raised in Ma's report relate to the manner of the shield building's construction. For reasons of cost efficiency, the AP1000's is constructed from 'modules': steel-reinforced concrete sections, which are prefabricated in factories and then shipped to the plant by rail (Schulz, 2006, p. 1554). The modules come in two types — referred to in the certification literature as "Module 1" and "Module 2" — and differ primarily in the spacing between their internal steel "tie bars." Module 1s have closely packed bars, making them stronger but also heavier and more expensive. They are intended for use in areas of high loads: about 40% of the building's walls, mostly at the bottom of the structure where they must bear the weight of the modules above. The remaining 60% of the building is constructed from Module 2s, which have looser spacing between bars, making them cheaper and lighter but also weaker. Ma worried that this modularity had uncertain and potentially dangerous implications for the shield building's structural integrity (and thus the resilience of the emergency cooling system that depended on it).

His specific concerns related to the modules' 'ductility': their ability to flex without breaking when subject to earthquake vibrations or an external blow (such as from tornados or airplanes). Ductility has important ramifications for the structural integrity of a reinforced concrete building, and is closely assessed in the reactor certification process. To this end the NRC traditionally invokes a set of standards and tests outlined by the American Concrete Institute (ACI), known as the "ACI Code."³⁰ In the planning stages of the certification process, Westinghouse and the NRC had agreed to use the ACI code for assessing the AP1000's modules.³¹ When they came to test Module 2, however, it failed something called the "out-of-plane shear test," indicating it was more brittle than the ACI would accept. A further problem arose when Westinghouse was unable to complete another key trial prescribed by the code — the "in-plane shear test" — on either module, due to the limitations of their facilities (NRC,

³⁰ Specifically Part 349; otherwise known as "ACI-349". The code does not prescribe tests for structures that meet specific prescriptive requirements. The modules were too innovative to meet these requirements, however, so the NRC agreed that Westinghouse would use tests to demonstrate compliance (NRC, 2010b).

³¹ Due to their varying compositions, the module types differ in their ductility.

2010b). These failures and gaps formed the basis of Ma's non-concurrence report. He pointed out that the shield building's modules did not fulfill key elements of the ACI code, and concluded from this that the structure's ability to withstand potential earthquakes and impacts was insufficiently proven (NRC, 2010b).³²

The NRC, as required by the non-concurrence process, offered detailed rebuttals to Ma's concerns.

In relation to Module 2 failing its 'out-of-plane' test, it argued that the result was misleading and unrepresentative because the reactor was a "first of a kind design" for which the test was not intended. The ACI did not develop the tests for composite steel-and-concrete modules, it noted, so the failure did not prove they were unsafe or even noncompliant (NRC, 2010b). It further argued that the tests were premised on an erroneous assumption: that earthquakes create out-of-plane loads. Contra to the tests' assumptions, it claimed, earthquakes primarily produce in-plane loads, so the structure's out-of-plane ductility only needed to resist tornado-generated missiles, which exert less force (NRC, 2010b, p. 5). Module 2 might have technically failed its out-of-plane ductility tests, in other words, but, by this view, it had nevertheless demonstrated the resilience the tests were intended to ensure.

In relation to Westinghouse's inability to perform the 'in-plane' tests, the NRC argued that the tests had been adequately replaced by simulations, which demonstrated adequate compliance in a different way. It explained that Westinghouse had gathered test data on different structures of similar composition to its modules, and then used computer models — called 'Finite element method structural analyses' — to extrapolate from that data to the AP1000 modules. It further pointed out that the modelers had used conservative margins in case their simulations were less than perfect (NRC, 2010b, p. 5-6).

Finally, and most generally, the NRC argued that the ACI code was not, in fact, a mandatory regulatory requirement. It conceded that its regulatory guidelines "endorsed" the code, and that "...changes to the design that would satisfy [Ma's objections] would result in a more robust

³² The tests are especially important, he argues, because the shield building has "highly irregular configurations" that "...good engineering practices would try to avoid" (NRC, 2010b, p. 19).

structure” (NRC, 2010b). Having established that such changes were unnecessary, however, it claimed to be powerless to insist on their implementation.

The non-concurrence process also allowed Ma to respond to the NRC’s rejoinders.

To the argument that the out-of-plane shear tests were unrepresentative, Ma — citing a report by outside consultants hired by the regulator itself — raised three objections: two of them technical, and the third a point of principle. Firstly, he argued that the NRC was misconstruing the underlying logic of the test (by relating it directly to shear forces rather than to the ‘flexural strength’ of the building). Secondly, he argued that the NRC’s assertion that earthquakes would not exert out-of-plane shear was wrong, due to the cylindrical shape of the building.³³ Finally, he critiqued the claim that the modules deserved exceptional treatment because they were a “first of a kind” design. This novelty should be a reason to be cautious with the requirements rather than a reason to relax them, he argued, especially the requirements concerning ductility, since it counteracts a phenomena (torsion) that is poorly understood analytically (NRC, 2010b).

To the argument that experts substituted computer models for the ‘in-plane’ shear tests it was unable to perform, Ma questioned the models’ representativeness. He argued that the commercial, general-purpose computer models Westinghouse used for the analysis were an inadequate (or not knowably adequate) substitute for the tests they replaced. The software was designed for analyzing regular, uniform structures, he claimed, and was inappropriate for analyzing complex, irregular structures made from interlinked modules like the shield building. (Not least because the connections between modules introduced important variables that could not be accounted for in the models.) He also dismissed assurances about using conservative values to compensate for uncertainty; arguing it was impossible to adequately define ‘conservative’ without a good understanding of the system (which the test itself was supposed to provide.) So it was, he concluded, that “[t]he analysis results should be viewed as approximations at the best, and not as accurate representations of the actual behavior of the shield building” (NRC, 2010b).

³³ “Due to the three-dimensional cylindrical shape of the shield building,” he writes, “any point in the shield building is subjected to both in-plane and out-of-plane shears concurrently and simultaneously during earthquakes.”

To the NRC's final argument, that the ACI code was not actually mandatory, Ma responded that this was a matter of interpretation, and that by his interpretation the regulator's decision did not reflect best practice: being inconsistent with guidance requiring that structures meet "... quality standards commensurate with the importance of the safety functions to be performed" (NRC, 2010b).³⁴ "NRC staff acceptance of a design which does not fulfill the [ACI] code would set a regulatory standard below that applied to other designs," he concluded, "... in spite of the greater importance of the shield building to the AP1000 design."

The outcome of Ma's non-concurrence report is perhaps best characterized as a (very one-sided) agreement to disagree. The NRC was given the last word. In its closing arguments it again appealed to formal constraints on its decision-making; pointing to rules and procedures that placed the burden of proof on those who would deny the application.³⁵ The reactor had met the required standards, it concluded, and, absent objective proof of hazard, its hands were tied. Officially at least, Ma's objections, having been formally recognized and recorded, were thus considered to be resolved.

4. Reckoning With Uncertainty

4.1 False Promises

For the purposes of this argument, the most important takeaway from these brief accounts of Ma's and Gunderson's disputes with the NRC is that neither was 'solved' logically and empirically, so much as they were 'settled' organizationally and bureaucratically. When following their arguments from claim to counterclaim, moreover, it becomes clear that neither dispute *could* have been resolved on empirical and logical grounds alone. Some of the issues they raise rest on technical questions that might have been amenable to further study and clarification — whether earthquakes create out-of-plane loads, for example — but key differences hinge on questions for which there can be no knowably 'correct' answers.

³⁴ GDC 1 10 CFR 50.55a (a)(1) and 10 CFR Part 50, Appendix A.

³⁵ Regarding the computer models, for example, it asserted that "absent specific information that the results are materially inaccurate, we have no basis for rejecting the applicant's approach" (NRC, 2010b).

Even the most technical issues Ma and Gunderson raise sometimes offer glimpses of irreducible uncertainties. These are most explicit, perhaps, in the appendices of Ma's non-concurrence report, the precise language of which the NRC had little control over. A document labeled 'Attachment 1,' for instance, contains feedback on problems of testing, construction and welding. It includes the noteworthy line: "The staff has no confidence that a potential success of carefully mockup [sic] [ductility] tests would be replicated during construction." Another — an external report commissioned by the regulator — straightforwardly states that: "A generally accepted criterion for assessing adequate ductile behavior does not exist amongst jurisdictions or design code [sic] worldwide" (NRC, 2010b). While a third — an email exchange around the viability of the modeling software — includes the line: "So, there are no definite reports or studies addressing this issue. But the anecdotal evidence [about the capabilities of the commercial software used by Westinghouse] is compelling" (NRC, 2010b).

Buried beneath both Ma's and Gunderson's settlements, therefore, there remain unanswered (arguably unanswerable) questions and unresolved disagreements: about the representativeness of computer models; the meaning of 'best practice'; the long-term diligence of maintenance workers; the applicability of the 'ACI code,' and much else besides. These questions and disagreements, along with their unsatisfying resolutions, are masked by the canonical ideal of rationality, which reduces messy debates to pristine numbers and declarative conclusions. But for anyone willing to look behind the curtain, they demonstrate that even the most rule-governed reactor safety-claims hinge on subjective judgements and unprovable assumptions.

Our purpose in this paper is not to single-out the AP1000 as being distinctively unsafe or opaque relative to other reactors. The kinds of disputes and misgivings outlined above are not unique to the AP1000.³⁶ The reactor might indeed be safer than its peers, and the NRC's assessment of that safety might be similarly exceptional. Our point, rather, is that all contemporary assertions about reactor safety are problematic — even those pertaining to designs touted as being unprecedentedly safe and straightforward like the AP1000. Moments of expert contention like Ma's and Gunderson's disputes with the NRC are worth exploring, not because they are exceptional, but because they shed light on opaque properties of reactor assessment work in general.

³⁶ For the case of ambiguities in the safety case of another reactor design, see (Ramana & Seshadri 2015).

All reactor safety assessments contain similar subjective judgements, and, as a result, all contain uncertainties that materially affect their credibility. The truth of this can be glimpsed without examining those assessments closely. It can be seen in the fact that the NRC's analyses, when performed at different times, have arrived at very different conclusions about the same reactors (e.g. NRC 1990, 1–4. See also: Lochbaum 2000, 13–17; Ramana & Seshadri 2015; Miller 2003, 175). It can be seen in the fact that different regulators from different jurisdictions draw divergent conclusions from the same evidence (Simola 2002; Amendola 1986).³⁷ And it can be seen in the fact that past assessments have assigned extremely low probabilities to events that have already occurred on multiple occasions (Miller 2003, 171; Cooke 1982, 334-5).

While recognizing the generalizability of this insight across reactor assessments, however, it is also important to remember the uniqueness of reactor assessments relative to other engineering ventures. The nuclear sphere is not alone in grappling with uncertainty. Every engineering test, model and theory (and every system or analysis built on them) rests on subjective judgements and contains a degree of epistemological uncertainty. On a practical level, however, this uncertainty does not affect all analyses or systems equally. As we saw, the greater the complexity of a system, the larger the network of knowledge-claims on which it rests, and the more likely it (or its assessment) is to have been built on an erroneous assumption or misjudgment. And the more reliability a system demands, the more important it becomes that every judgement and assumption implicit in its design (or assessment) be perfect, (and the harder it becomes to examine those assumptions and judgements against empirical experience). And when these properties come together — as they do, to an almost unique extent, in the context of reactor safety — then they have powerful epistemological ramifications. Reactors and their safety assessments are not exceptional in resting on unprovable assumptions and subjective judgments, in other words, they are unique because their properties make those assumptions and judgments determinative (and, indeed, prohibitive).

This paper is far from being the first to conclude that experts cannot realistically predict the frequency of reactor accidents, and that reactor assessments are less authoritative (and reactors probably less safe) than is usually believed. Scholars who have looked directly at the logic of the

³⁷ When the UK's Office for Nuclear Regulation (ONR) certified the AP1000 in March 2017, for example, it identified a number of different issues to those raised by the NRC. It also required that some of the same issues be addressed in different ways. For a hand wringing account of these national differences by an advocate for nuclear power, see (Conca 2015).

industry's assessments have almost all come to similar conclusions. This includes social scientists (e.g. Miller 2003; Shrader-Frechette 1991; Clarke 2006; Perrow 1999; Perkins 2014), but also prominent physicists and engineers from outside the industry (see Wellock 2017: 705-708). Most notable in the latter regard, perhaps, was a blue-ribbon panel of experts — the so-called “Lewis Panel” — convened by Congress to investigate the NRC’s probabilistic approach to assessing reactor safety when it was first introduced. The panel’s report (Lewis et al. 1978) came to very similar conclusions to those outlined above; being highly critical of idea that regulators could definitively establish the failure-behavior of reactors, and for many of the same reasons.³⁸ Some of the regulator’s key premises were “conceptually impossible,” the report argued, and it had greatly understated its own uncertainties (Lewis et al. 1978).

Despite this prior work, we nevertheless believe the argument of this paper to be valuable. Partly because its conclusion, like that of many critical arguments, is routinely challenged, and, as Jasanoff (1993: 123) puts it, is “in danger wearing thin without continual monitoring and periodic repair.” The Lewis report was primarily written for an expert audience, and in the decades since its publication the nuclear engineering community has closed ranks around PRA. More fundamentally, however, because even if other critics have challenged the efficacy of reactor assessment, the *necessity* of its shortcomings, and (especially) the *exceptionality* of those shortcomings, have both struggled to find much purchase in the mainstream discourse, scholarship and governmentality around reactors. Constructivism helps illuminate both properties, and, in so doing, helps accentuate the shortcomings themselves. It sheds light on logical infirmity of reactor assessments, as other accounts do, but also on why this infirmity is easily overlooked and why it is more consequential than in other, seemingly equivalent, contexts.

The latter point, regarding exceptionality, is easily lost amid the debates that surround all expert knowledge, but it is important. A more nuanced understanding of why and how nuclear safety differs from most engineering problems — its acute epistemological challenges and uncertainties — poses far-reaching questions. The US, like all nuclear polities, has built

³⁸ “Since [core melt] has never occurred in a commercial reactor, there are no direct experimental data on which to base an estimate.” The report states. “[...] Therefore it is necessary to resort to a theoretical calculation of the probability. But since the system is so complex, a complete and precise theoretical calculation is impossibly difficult. It is consequently necessary to invoke simplified models, estimates, engineering opinion, and in the last resort, subjective judgments” (Lewis et al. 1978 :6).

considerable intellectual, administrative and organizational empires on the certainty that catastrophic reactor accidents are too unlikely to be worth considering. These, we contend, are foundations of sand.

To fully explore the complex implications of reactors without certainty would be an endeavor that far exceeds the scope of this paper, which primarily aims to provide a basis for further scholarship. If only as a roadmap for future consideration, however, we will close by sketching some preliminary thoughts regarding that endeavor's scope and direction.

4.2 Reactors Without Certainty

4.2.1 STS

Before exploring the wider implications of a constructivist understanding of nuclear safety, it is worth reiterating its implications for STS itself. For even though the idea that reactor assessments are imperfect and judgement-laden is unremarkable in an STS context, the practical significance of this imperfection and subjectivity is more noteworthy. Those assessments (or, more broadly, any predictive assessments of extreme reliability in a complex system) represent a category of contemporary knowledge-claims on which the discipline can take a strong normative position. Here, if nowhere else, STS scholars can make pronouncements on the credibility of contemporary experts and their assertions without problematizing their commitment to constructivism. This is, we recognize, a relatively narrow category of knowledge-claims. One that, although not limited to reactor assessments exclusively (see e.g. Downer [2017] on the safety of jetliners; or Lynch et al. [2008] on the accuracy of DNA testing), encompasses few of the contentious issues that drive many of the discipline's internal debates about normativity. It is, however, an important category of knowledge. Societies only require that experts attest to extreme levels of failure performance when the implications of failure are commensurately meaningful.

4.2.2 Risk Studies

Looking outwards to wider academic 'risk' literatures, it quickly becomes clear that problematizing the certainty of reactor assessments has complicated consequences. Pull on this

thread and a lot of established ideas begin to unravel. Take, for instance, the literatures concerning ‘risk perception’ (most notably those that apply the ‘psychometric paradigm’) (e.g. Slovic et al. 1982), or those that apply ‘cultural theory’ (e.g. Douglas & Wildavsky 1983), to interpret public attitudes to technological risk. These literatures are very different in most respects, but foundational elements of both are derived from studies that examine seemingly ‘irrational’ (or ‘boundedly-rational’) public concerns about atomic energy, and seek to make that irrationality intelligible.³⁹ Such studies define rationality with reference to formal risk analyses, crucial elements of which, as we have seen, are premised on the improbability of catastrophic disasters. Their conclusions would undoubtedly be different if scholars reinterpreted skepticism of those analyses as evidence, not of irrationality, but of unusual perspicacity, technical literacy or critical engagement. (This, in turn, would have ramifications for the prominent vein of pro-nuclear journalism that draws on these literatures to portray the industry’s critics as fear-mongers, beholden to irrational prejudices [see Miller 2003: 187]).

A similar argument could be made in relation to the more principles-based and policy-oriented scholarship that explores the ‘social-’ or ‘ethical acceptance’ of technological risk (e.g. Sjöberg 2004; Taebi 2017; Grunwald 2000; Hanson 2003). This literature looks at more explicitly normative questions regarding the levels and distributions of risk that should be acceptable to publics. Much like the literatures above, however, it routinely makes a distinction between the ‘actual risk’ of reactors (as it is assessed and defined by experts), and the public perception and definition of that risk. While it routinely problematizes the way experts ‘define’ risk,⁴⁰ therefore, it rarely questions whether expert risk assessments are credible on their own terms. To do so might lead scholars to more actively accommodate the possibility that lay publics could share expert definitions of risk while (rationally) rejecting expert conclusions about it.

4.2.3 Costs and Benefits

Scholars in other disciplines might build on the infirmity of reactor assessments by exploring questions that are usually rendered moot by the certainty that catastrophic accidents will not occur. Economists who price energy options, for example, might endeavor to more

³⁹ For more on this tension in wider discussions of risk see Jasanoff (1993: 123).

⁴⁰ Exclusively in terms of physical hazards, for instance, without considering wider concerns such as job security.

systematically consider the costs of such accidents. For even if the likelihood of incurring those costs is fundamentally uncertain, it could be useful to explore the point at which they would become untenable. Environmental scholars might ask similar questions: exploring how improbable meltdowns must be before the carbon benefits of nuclear energy cease to outweigh their other environmental hazards. Both endeavors, in turn, would require substantial scientific research into the myriad potential consequences of reactor accidents; work that is systematically underfunded at present because it is deemed unnecessary.

4.2.4 Agnotology

As well as shining new light on the consequences of catastrophic accidents, understanding the exceptional epistemology of reactor safety assessments might also offer new perspectives on the discourse and rhetoric that surrounds them. The NRC's public portrayal of its safety assessments, for example, could itself be an object of study.

The regulator is not unaware that its assessments contain judgements and uncertainties. In private discourse, its leading assessment experts will acknowledge the subjectivities of PRA and its limitations as a tool for establishing the absolute risks of a system, preferring to characterize it as a means of comparing the relative risks of different design options (Miller 2003: 183-4; Wellock 2017: 694; Apostolakis 1990: 1363; Wu & Apostolakis 1992: 335).⁴¹

In public contexts, however, there is little question that it vigorously promulgates an understanding of its assessments as rule-governed, objective and definitive determinations of accident risk: routinely eliding their subjectivities and uncertainties (Apostolakis 1990: 1363; Wellock 2017; Miller 2003: 183-4). And it is clear from any account of PRA's history that the regulator intended its assessments to be understood in this way: their explicit rhetorical function being to bound accidents out of public discussion (Wellock 2017; Miller 2003; Apostolakis 1988: 248). This attitude is made most evident, perhaps, in the NRC's 'Guidelines for External Risk Communication' (NRC 2004). Intended for internal consumption, the document advises regulators to "[a]void making statements such as 'I cannot guarantee...' or '[t]here are no guarantees in life...'" when speaking about reactor safety; because "[...]

⁴¹ George Apostolakis — MIT professor, and risk assessment expert and former NRC commissioner — is uncommonly explicit on these limitations (e.g. Apostolakis 1988; 1990).

statements like these contribute to public outrage [...]” (NRC 2004: 38). (Notably, the report does not suggest that policymakers are less in need of reassurance. If anything, 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]).

The distinctive epistemology of reactor assessments becomes important when interpreting this portrayal, because the NRC does not intuitively seem unusual in presenting its conclusions this way. It has long been understood that scientists and engineers, across a range of domains, routinely hide the ‘messiness’ of their knowledge-work and project a misleading sense of objectivity and completeness about their conclusions (e.g. Latour 1987; Wynne 1988; Barnes & Edge 1982). The NRC does not even seem unusual in leveraging a performance of objectivity to bolster its institutional authority over matters pertaining to public policy. This too is common, as an extensive and multidisciplinary literature attests (e.g. Hilgartner 2000; Jasanoff 2003; Porter 1995; Power 1997).

Again, however, the NRC’s elisions and performances are made exceptional by the nature of the claims involved. When the expert climatology community makes public predictions about anthropogenic climate change, it — like all expert communities — undoubtedly downplays uncertainties and internal debates (Grundmann 2011). Strictly speaking, this practice is misleading. It may occlude significant value-laden judgements, and may even be counterproductive.⁴² For all this, however, the community’s predictions still arguably represent good-faith, epistemologically-plausible conclusions: ‘normal’ science as Kuhn (1962) might put it.⁴³ When the NRC downplays the ambiguities of its safety assessments, by contrast, it is miscommunicating the *essence* of those assessments, and making consequential claims that are knowably implausible. If the climatologists are analogous to a doctor promising her patient that his fractured arm will definitely heal, we might say, then nuclear regulators are more akin to a doctor promising her patient he will definitely never fracture his arm again. Both promises

⁴² Consider, for example, the ‘Climate-gate scandal’ that arose when internal emails from the Climate Research Unit at the University of East Anglia were leaked in 2009 (Grundmann 2011).

⁴³ And in an age where bad-faith actors actively misrepresent expressions of normal scientific dissent, the practice of hiding that dissent might reasonably be interpreted as an effort to more fully inform publics.

contain some measure of concealed uncertainty (doctors know that a fraction of fractures never fully heal for unanticipated reasons), yet they are far from equivalent.

In light of this distinction, it becomes difficult to fold the NRC's portrayal of its assessments unproblematically into wider narratives about the "technicizing" of risk problems: just another instance of a much more general phenomenon. It would not be unreasonable, in fact, for scholars to look at documents like the NRC's 'risk communication guidelines' through the lens of what is sometimes known as the 'agnotology' literature (e.g. Proctor & Schiebinger 2008; McGooney 2012; Oreskes & Conway 2010). This literature explores efforts by powerful organizational actors, such as the tobacco and oil industries, to further their interests by shaping the public perception of technoscientific claims. To date, the discourse around atomic energy has largely escaped its gaze. This is understandable, given that it rarely questions officially-sanctioned knowledge-claims, and usually examines efforts to foster misleading impressions of uncertainty. For these same reasons, however, it might be interesting for scholars to look more directly at officially-sanctioned efforts to foster misleading impressions of certainty.

4.2.5 Capture

In a related vein, the exceptionalism of reactor assessments also calls for careful reconsideration of the role and institutional standing of nuclear regulators themselves. In the US, as in most polities, the regulator currently serves as the exclusive arbiter of all authoritative public knowledge pertaining to reactor safety.⁴⁴ Since safety is integral to almost every question regarding atomic energy, and since there is no statistically-meaningful way of testing the NRC's assertions about it, this monopoly gives the regulator enormous influence over the democratic decision-making around reactors.⁴⁵ This influence needs recalibrating.

It would be wrong, of course, to dismiss offhandedly the expertise of regulators; who, despite the limits of their knowledge, might still reasonably claim privileged insight into reactor safety. But that expertise should not be accepted blindly. Reactor safety is a unusual class of engineering problem, and experts cannot speak to it with the same authority that they speak to

⁴⁴ As the State of Vermont discovered at the top of this paper.

⁴⁵ As scholars of audit processes have long attested, experts and their calculative practices become highly influential when an important property is knowable only through its assessment (Hopwood & Miller, 1994; Power, 1997)

most other engineering variables. It is important that publics, policymakers and bureaucracies recognize this difference in the way they engage with bodies like the NRC and the claims they wield.

The question of how the rules and structures of nuclear governance should incorporate regulatory expertise deserves far more nuanced consideration than we could offer here.⁴⁶ We do note, however, that a constructivist understanding of reactor assessment implies an uncommon need, when addressing this question, to understand nuclear regulators in relation to the organizational- and value-structures in which they work. For while all regulatory knowledge has ‘politics’ (in the sense that it contains subjective judgements) (Jasanoff 1990), the fact that reactor safety assessments cannot be tested against experience greatly exacerbates the significance and scope of those politics.

To recognize this difference, consider, for example, the close relationship between the NRC and the industry it regulates; wherein experts routinely move between the two spheres, and manufacturers perform the majority of the analyses that regulators use to make their rulings. Drawing upon (Jasanoff 1996), this relationship has been described as co-production (Slayton and Clark-Ginsberg 2018). Such arrangements are common across a range of regulatory domains, and academics often view them as problematic due to the hidden subjectivities of formal regulation (e.g. Power 1997).⁴⁷ In most contexts, however, the scope for hidden interests to influence regulatory work is constrained, at least to an extent, by the visible results of that work.⁴⁸ (Sudden outbreaks of food poisoning, for instance, can create ‘crises’ for food regulators [Hutter & Lloyd-Bostock 2017].) The nuclear realm is uncommonly insulated from this feedback. Reactor accident predictions could be off by orders-of-magnitude without this error becoming visible for decades (until it became tragically and irrevocably apparent).

⁴⁶ The specifics of such questions would be unique to reactors with their outsized hazards and exceptional epistemology, but we might imagine tentative answers building on the literature that grapples with regulation and policymaking under conditions of uncertainty (e.g. Jasanoff 1990; Collingridge and Reeve 1986; Rip 1986).

⁴⁷ Organizations usually justify such arrangements by invoking (misleadingly) the ‘objectivity’ of the work itself.

⁴⁸ In other domains, such as civil aviation, the effects of capture are further offset by the industry’s structural incentives being aligned with the public interest (e.g. Downer 2010). But this is a difficult case to make for the nuclear industry.

In light of this, the oft-made assertion that nuclear regulators are ‘captured’ by their regulatee⁴⁹ should arguably carry more weight than similar assertions in other spheres.⁵⁰ More so than in most contexts, it matters that these regulators are not faceless intermediaries following applying impersonal standards but industry insiders with direct interests in its future, who interpret contestable rules and make complex choices about ambiguous evidence. Scholars since Marx have argued that structural interests colonize values, and it is a foundational claim of the STS literature that values come to permeate even the most rigorous knowledge-claims (e.g. Bloor 1976). To imagine that NRC regulatory determinations are an exception to this rule would be to defy generations of social research.⁵¹

4.2.6 Governance

Last, but far from least, reassessing the certainty of reactor safety assessments would have far reaching implications for nuclear governance itself: the decisions societies make regarding energy investments, reactor siting, emergency planning, and anything else that even obliquely touches the integrity of nuclear infrastructures. At present, all such decisions are premised on an understanding that catastrophic accidents are too improbable to merit consideration, and that this improbability is *knowable* with a high degree of objective certainty. The latter condition is fundamental. States do not site reactors downwind of major cities on the belief that meltdowns are “probably” or “arguably” too unlikely to be a concern; the costs of catastrophic reactor accidents are too grave.

It seems probable, therefore that such decisions might change if assertions about the probability of such accidents were understood as contested judgements, made by experts with clear interests in the industry’s future. And it seems almost inevitable that such decisions would change if those assertions were understood as being *a priori* implausible: impossible promises,

⁴⁹ Assertions of capture are common across nuclear regulatory regimes (Ramana & Seshadri, 2015). And observers routinely describe the NRC in these terms; including one of the agency’s former chairs (Jaczko 2019; also, e.g. Katz, 1984; von Hippel, 2011; Madrigal, 2011).

⁵⁰ The idea of ‘regulatory capture’ originated in studies of efforts by governments to stop corporations from forming monopolies (Stigler 1971; Peltzman 1976). Broadly speaking, it refers to a situation where the regulators interests become aligned with the organizations they are expected to police rather than those of the publics they are expected to protect.

⁵¹ It is worth noting that constructivist understanding safety assessment offers a nuanced understanding of capture’s mechanisms. Discussions of capture usually construe it as a process wherein regulators knowingly subvert their mandate. When understood in relation to the interpretive-flexibility of technical decision-making, however, the process does not require such connotations. Seen in this light, we might think of it as a more subtle process wherein aligned sympathies and shared world-views come to colonize irrevocably ambiguous rules and interpretations.

epistemologically unsupportable. For while public decisions are routinely premised on expert assertions with hidden uncertainties and subjectivities, very few are premised on assertions that are demonstrably unrealistic; and none so consequential.

So it is that the limits of reactor safety assessments raise important questions about the legitimacy of contemporary nuclear policy. Whether or not reactors would, or should, remain politically viable under these circumstances is a complicated question that we will leave to others. We do contend, however, that maintaining the political viability of atomic energy is not a justifiable reason to miscommunicate the nature and certainty of reactor safety assessments. The costs and hazards of reactor accidents demand democratic engagement in decisions about atomic energy, and, as scholars like Rip (1986) and Beck (1992) argued years ago, this requires transparency about uncertainty. For when the stakes are exceptional, the question of 'how safe is safe enough' can only be addressed with informed discussions of 'how certain is certain enough,' and knowing what we don't know is more useful than believing we know what we don't. Meretricious certainty might be comforting, but it offers little real guidance for navigating the high technological frontier.

References

- Abernethy, C. (2012, February 9). NRC Approves Vogtle Reactor Construction - First New Nuclear Plant Approval in 34 Years. *Nuclear Street*. Retrieved from http://nuclearstreet.com/nuclear_power_industry_news/b/nuclear_power_news/archive/2012/02/09/nrc-approves-vogtle-reactor-construction-2d00-first-new-nuclear-plant-approval-in-34-years-2800-with-new-plant-photos-2900-020902#.W1aU6thKjOQ (accessed 03/04/2018).
- Amendola, A. (1986). Uncertainties in systems reliability modelling: Insight gained through European Benchmark exercises. *Nuclear Engineering and Design*, 93(2), 215–225. [https://doi.org/10.1016/0029-5493\(86\)90221-9](https://doi.org/10.1016/0029-5493(86)90221-9)
- Ansolabehere, S., Deutch, J., Driscoll, M., Gray, P. E., Holdren, J. P., Joskow, P. L., ... Beckjord, E. S. (2003). *The Future of Nuclear Power*. Retrieved from Massachusetts Institute of Technology website: <http://web.mit.edu/nuclearpower/>
- Apostolakis, G. (1990). The concept of probability in safety assessments of technological systems. *Science*, 250(4986), 1359–1364. <https://doi.org/10.1126/science.2255906>
- Apostolakis, G. E. (1988). The interpretation of probability in probabilistic safety assessments. *Reliability Engineering & System Safety*, 23(4), 247–252. [https://doi.org/10.1016/0951-8320\(88\)90034-8](https://doi.org/10.1016/0951-8320(88)90034-8)
- Balogh, B. (1991) *Chain Reaction: Expert Debate and Public Participation in American Commercial Nuclear Power, 1945–1975*. Cambridge, UK: Cambridge University Press.
- Barnes, B. & Edge, D. (1982) “Science as Expertise,” in Barnes, B. & Edge, D. (eds) *Science in Context: Readings in the Sociology of Science*. Milton Keynes: Open University Press.
- Beck, U. (1992) *Risk society: Towards a New Modernity*. London: Sage.
- Bloor, D. (1976) *Knowledge and Social Imagery*. London/Henley/Boston: Routledge & Kegan Paul.
- Clarke, L., & Perrow, C. (1996). Prosaic Organizational Failure. *American Behavioral Scientist*, 39(8), 1040–1056. <https://doi.org/10.1177/0002764296039008008>
- Clarke, L. (2006) *Worst Cases: Terror and Catastrophe in the Popular Imagination*. Chicago: University of Chicago Press.
- Clayton, M. (2011, April 20). Court clash in Vermont might alter balance of power over nuclear plants. Retrieved January 11, 2016, from <http://www.publicintegrity.org/2011/04/20/4287/court-clash-vermont-might-alter-balance-power-over-nuclear-plants>
- Cochran, T. B. (2011, April 11). Testimony at the Joint Hearings of the Subcommittee on Clean Air and Nuclear Safety and the Committee on Environment and Public Works on April 12, 2011. Retrieved May 24, 2012, from http://www.nrdc.org/nuclear/tcochran_110412.asp
- Collingridge, D. & C. Reeve (1986) *Science Speaks to Power*. St. Martin's, New York.
- Collins, H. (1985) *Changing Order*. London: Sage.
- Collins, H. & Evans, R. (2002) The Third Wave of Science Studies: Studies of Expertise and Experience. *Social Studies of Science*, 32 (2): 235–96.
- Collins, H. & Pinch, T. (1993) *The Golem: What Everyone Should Know About Science*. Cambridge: Cambridge University Press.
- Conca, J. (2015, March 12). Can't All Nuclear Just Get On The Same Page? *Forbes*. Retrieved from <http://www.forbes.com/sites/jamesconca/2015/03/12/cant-all-nuclear-just-get-on-the-same-page/>

- Constant, E. W. (1999). Reliable Knowledge and Unreliable Stuff: On the Practical Role of Rational Beliefs. *Technology and Culture*, 40(2), 324–357.
- Cooke, R. M. (1982). Risk Assessment and Rational Decision Theory. *Dialectica*, 36(4), 329–351.
- Douglas, M., & Wildavsky, A. (1983). *Risk and Culture: An Essay on the Selection of Technological and Environmental Dangers*. Berkeley, CA: University of California Press.
- Downer, J. (2007). When the Chick Hits the Fan: Representativeness and Reproducibility in Technological Tests. *Social Studies of Science*, 37(1), 7–26.
- Downer, J. (2010). Trust and technology: The social foundations of aviation regulation. *The British Journal of Sociology*, 61(1), 83–106.
- Downer, J. (2014). Disowning Fukushima: Managing the credibility of nuclear reliability assessment in the wake of disaster. *Regulation & Governance*, 8(3), 287–309. <https://doi.org/10.1111/rego.12029>
- Downer, J. (2015). The unknowable ceilings of safety: Three ways that nuclear accidents escape the calculus of risk assessments. In B. Taebi & S. Roeser (Eds.) *The Ethics of Nuclear Energy: Risk, Justice and Democracy in the Post-Fukushima Era*. Cambridge: Cambridge University Press (pp. 35-52).
- Downer, J. (2016). Resilience in Retrospect: Interpreting Fukushima’s disappearing consequences. In A. Herwig & M. Simoncini (Eds.) *Law and the Management of Disasters: The Challenge of Resilience*. London: Routledge (pp. 42-60).
- Downer, J. (2017). The Aviation Paradox: Why We Can ‘Know’ Jetliners But Not Reactors. *Minerva*, 55(2), 229–248.
- EPRI. (2007). *Probabilistic Risk Assessment Requirements for Passive Safety Systems* (Program on Technology Innovation No. 1015101). Palo Alto, CA: Electric Power Research Institute. Retrieved from <http://www.epri.com/abstracts/Pages/ProductAbstract.aspx?ProductId=00000000001015101>
- Fallows, J. (1985) The American Army and the M-16 Rifle. In D. MacKenzie & J. Wajcman (Eds.) *The Social Shaping of Technology (Second Edition)* Buckingham: Open University (pp. 382-394).
- Feldman, S. (2010, April 23). Coalition: Design for New Nuclear Reactor Less Safe Than America’s Current Fleet. *InsideClimate News*. Retrieved from <http://insideclimatenews.org/news/20100423/coalition-design-new-nuclear-reactor-less-safe-americas-current-fleet>
- Fox-Keller, E. (1994) The Paradox of Scientific Subjectivity. In A. Megill (Ed.) *Rethinking Objectivity*. Durham, N.C.: Duke University Press. (pp. 313–331).
- Feyerabend P. (1975) *Against Method: Outline of an Anarchistic Theory of Knowledge*. Redwood Burn; London.
- Fuller, J. G. (1975). *We almost lost Detroit*. New York: Reader’s Digest Press.
- Gross, A. G. (1990) *The Rhetoric of Science*. Cambridge, Mass: Harvard University Press.
- Gross, P. and Levitt, N. (1994). *Higher Superstition: The Academic Left and Its Quarrels With Science*. Johns Hopkins University Press; Baltimore.
- Grundmann, R. (2013). “Climategate” and The Scientific Ethos. *Science, Technology, & Human Values*, 38(1), 67–93.
- Grunwald, A. (2000) Technology Policy between Long-Term Planning Requirements and Short-Ranged Acceptance Problems: New Challenges for Technology Assessment. In J. Grin (Ed) *Vision Assessment: Shaping Technology in 21st Century Society*. Berlin: Springer (pp. 99–147).

- Gundersen, A. (2010a). *Nuclear Containment Failures: Ramifications for the AP1000 Containment Design*. Burlington, VT: Fairewinds Associates, Inc.
- Gundersen, A. (2010b). *Post Accident AP1000 Containment Leakage: An Unreviewed Safety Issue*. Burlington, VT: Fairewinds Associates, Inc.
- Hamblin, J. D. (2012). Fukushima and the Motifs of Nuclear History. *Environmental History*, 17(2), 285–299.
- Hansen, J. (2009). *Storms of My Grandchildren*. London: Bloomsbury Publishing.
- Hansson, S. O. (2003). Ethical Criteria of Risk Acceptance. *Erkenntnis*, 59(3), 291–309.
- Harki, G. A. (2011, August 15). AP1000 and Vogtle Inch Closer to Approval. *Nuclear Intelligence Weekly*, V(33), 5.
- Hilgartner, S. (2000) *Science on Stage: Expert Advice as Public Drama*. Palo Alto: Stanford University Press.
- Hopwood, A. G., & Miller, P. (Eds.). (1994). *Accounting social and institutional practice | Finance*. Cambridge: Cambridge University Press.
- Hutter, B. & Lloyd-Bostock, S. (2017) *Regulatory Crisis: Negotiating the Consequences of Risk, Disasters and Crises*. Cambridge University Press. Cambridge.
- IAEA. (1991). *Safety related terms for advanced nuclear plants, IAEA-TECDOC-626*. Vienna, Austria: International Atomic Energy Agency.
- IAEA. (2004). *Design of Reactor Containment Systems for Nuclear Power Plants* (No. IAEA Safety Standards Series NS-G-1.10). Vienna: International Atomic Energy Agency.
- IRSN. (2016). *Considerations on the performance and reliability of passive safety systems for nuclear reactors*. Paris: Institut de Radioprotection et de Sûreté Nucléaire. Retrieved from http://www.irsn.fr/EN/newsroom/News/Pages/20160107_Considerations-on-the-performance-and-reliability-of-passive-safety-systems-for-nuclear-reactors.aspx
- Jaczko, G. (2019) I oversaw the U.S. nuclear power industry. Now I think it should be banned. *Washington Post*. 16 May 2019. Online: https://www.washingtonpost.com/outlook/i-oversaw-the-us-nuclear-power-industry-now-i-think-it-should-be-banned/2019/05/16/a3b8be52-71db-11e9-9eb4-0828f5389013_story.html?utm_term=.a726492964f0 (Accessed 05/20/2019).
- Jasanoff, S. (1990) *The Fifth Branch: Science Advisors as Policymakers*, Harvard University Press.
- Jasanoff, S. (1993). Bridging the Two Cultures of Risk Analysis. *Risk Analysis*, 13(2), 123–129.
- Jasanoff, S. (1996). Beyond Epistemology: Relativism and Engagement in the Politics of Science. *Social Studies of Science*, 26(2), 393–418.
- Jasanoff, S. (2003). Technologies of Humility: Citizen Participation in Governing Science. *Minerva*, 41(3), 223–244.
- Johnson, W. (2011, July 24). U.S. heat wave causes new look at nuclear energy. *CBS News*. Retrieved from <https://www.cbsnews.com/news/us-heat-wave-causes-new-look-at-nuclear-energy/>
- Katz, J. E. (1984). The Uses of Scientific Evidence in Congressional Policymaking: The Clinch River Breeder Reactor. *Science, Technology, & Human Values*, 9(1): 51–62.
- Kuhn, T. (1996 [1962]) *The Structure of Scientific Revolutions* (3rd Edition). Chicago: The University of Chicago Press.
- Kusch, M. (2012) “Sociology of science: Bloor, Collins, Latour” In Brown, J. (ed.), *Philosophy of Science: The Key Thinkers*. London: Continuum Books.
- Latour, B. (1987). *Science in Action: How to Follow Scientists and Engineers through Society*. Cambridge, Mass: Harvard University Press.

- Latour, B. (2004). "Why Has Critique Run out of Steam? From Matters of Fact to Matters of Concern." *Critical Inquiry* 30 (2): 225-48.
- Lewis, H. W, R. J Budnitz, H. J.C Kouts, W. B Loewenstein, W. D Rowe, F. von Hippel, and F. Zachariassen. (1978). Risk Assessment Review Group Report to the U. S. Nuclear Regulatory Commission. NUREG/CR-0400. Washington, D. C.: Nuclear Regulatory Commission.
- Lochbaum, D.; Lyman, E. & Stranahan, S. (2014) *Fukushima: The Story of a Nuclear Disaster*. The New York Press. New York.
- Lochbaum, D. (2000). *Nuclear Plant Risk Studies: Failing the Grade*. Cambridge, MA: Union of Concerned Scientists.
- Lyman, E. S. (2011). Surviving the one-two nuclear punch: Assessing risk and policy in a post-Fukushima world. *Bulletin of the Atomic Scientists*, 67(5), 47–54. <https://doi.org/10.1177/0096340211421470>
- Lynch, M. (2017). "STS, Symmetry and Post-Truth." *Social Studies of Science*. 47(4): 593-599.
- Lynch, M.; Cole, S.; McNally, R. & Jordan, K. (2008) *Truth Machine: The Contentious History of DNA Fingerprinting*. Chicago: University of Chicago Press.
- Lyons, P. (2016). Challenges: Nuclear Power Today and Megawatt Size Reactors. Presented at the Workshop on Safe and Secure Megawatt-Size Nuclear Power Workshop. Retrieved from <https://arpa-e.energy.gov/?q=workshop/safe-and-secure-megawatt-size-nuclear-power-workshop>
- MacKenzie, D. (1996). How Do We Know the Properties of Artifacts? Applying the Sociology of Knowledge to Technology. In R. Fox (Ed.), *Technological Change: Methods and Themes in the History of Technology*. New York: Routledge (pp. 247–264).
- Madrigal, A. (2011). *Powering the Dream: The History and Promise of Green Technology*. Cambridge MA: Da Capo Press.
- McGoey, L. (2012). Strategic unknowns: Towards a sociology of ignorance. *Economy and Society*, 41(1), 1–16.
- Miller, C. R. (2003). The Presumptions of Expertise: The Role of Ethos in Risk Analysis. *Configurations*, 11(2), 163–202.
- Monbiot, G. (2006) *Heat: How to Stop the Planet Burning*. Penguin. London.
- NEI. (2013, August 15). Court Overrules State Laws Affecting Vermont Yankee. Retrieved January 11, 2016, from <http://www.nei.org/News-Media/News/News-Archives/Court-Overrules-State-Laws-Affecting-Vermont-Yanke>
- NRC. (2010a). *Containment Liner Corrosion* (No. Information Notice 2010-12). Rockville, MD: Nuclear Regulatory Commission. Retrieved from <http://pbadupws.nrc.gov/docs/ML1006/ML100640449.pdf>
- NRC. (2010b). *Safety Evaluation Report for AP1000 Shield Building: Non-Concurrence Process* (Non-Concurrence Process No. ML102630229). Rockville, MD: Nuclear Regulatory Commission. Retrieved from <http://pbadupws.nrc.gov/docs/ML1033/ML103370648.pdf>
- NRC. (2009). *Frequently Asked Questions About License Applications for New Nuclear Power Reactors* (No. NUREG/BR-0468). Rockville, MD: Office of New Reactors, Nuclear Regulatory Commission.
- NRC (2004) U.S. 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. (1990). *Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants*. NUREG-1150. Washington, D. C.: Nuclear Regulatory Commission.
- Oreskes, N. & Conway, E. (2010) *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*. Bloomsbury. London.
- Peltzman, S. (1976). Toward a More General Theory of Regulation. *Journal of Law and Economics*, 19(2), 211–240.
- Perrow, C. (1999). *Normal accidents : living with high-risk technologies* (Rev. ed.). Princeton, NJ: Princeton University Press.
- Pinch, T. (1993). “Testing - One, Two, Three ... Testing!”: Toward a Sociology of Testing. *Science, Technology & Human Values*, 18(1), 25–41. <https://doi.org/10.1177/016224399301800103>
- Popper, K. (1959) *The Logic of Scientific Discovery*. Basic Books, New York.
- Porter, T. M. (1995). *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton: Princeton University Press.
- Power, M. (1997). *The Audit Society: Rituals of Verification*. Oxford: Oxford University Press.
- Proctor, R. & Schiebinger, L. (2008) *Agnotology: The Making and Unmaking of Ignorance*. Stanford University Press; Palo Alto.
- Raju, S. (2016). Estimating the frequency of nuclear accidents. *Science & Global Security*, 24(1), 37–62.
- Ramana, M. V. (2011, April 20). Beyond our imagination: Fukushima and the problem of assessing risk. *Bulletin of the Atomic Scientists*. Retrieved from <http://thebulletin.org/beyond-our-imagination-fukushima-and-problem-assessing-risk-0>
- Ramana, M. V., & Kumar, A. (2014). ‘One in infinity’: failing to learn from accidents and implications for nuclear safety in India. *Journal of Risk Research*, 17(1), 23–42. <https://doi.org/10.1080/13669877.2013.822920>
- Ramana, M. V., & Seshadri, A. K. (2015). Negligence, capture, and dependence: safety regulation of the design of India’s Prototype Fast Breeder Reactor. *Journal of Risk Research*, 18(8), 1030–1050. <https://doi.org/10.1080/13669877.2014.1003958>
- Ramsey, C. B., & Modarres, M. (1998). *Commercial nuclear power : assuring safety for the future*. New York: Wiley.
- Rip, A. (1986). The Mutual Dependence of Risk Research and Political Context. *Science & Technology Studies*, 4(3/4), 3–15.
- Rip, A. (2003) Constructing Expertise in a Third Wave of Science Studies? *Social Studies of Science*. 33 (3): 419–34.
- Schulz, T. L. (2006). Westinghouse AP1000 advanced passive plant. *Nuclear Engineering and Design*, 236(14–16), 1547–1557. <https://doi.org/10.1016/j.nucengdes.2006.03.049>
- Shapin, S. (1995). Cordelia’s Love: Credibility and the Social Studies of Science. *Perspectives on Science*, 3(3), 255–275.
- Shrader-Frechette, K. (1991) *Risk and Rationality: Philosophical Foundations for Populist Reforms*. Berkeley: University of California Press.
- Simola, K. (2002). *Advances in Operational Safety and Severe Accident Research* [NKS-61]. Roskilde, Denmark: Nordic Nuclear Safety Research.
- Sismondo, S. (2010) *An Introduction to Science and Technology Studies*. Blackwell.
- Slayton, R., & Clark-Ginsberg, A. (2018). Beyond regulatory capture: Coproducing expertise for critical infrastructure protection. *Regulation & Governance*, 12(1), 115–130.
- Slovic, P; Fischhoff, B. & Lichtenstein, S. (1982) Why Study Risk Perception? *Risk Analysis* 2(2): 83–93.

- Sjöberg, L. (2004). Local Acceptance of a High-Level Nuclear Waste Repository. *Risk Analysis*, 24(3), 737–749.
- Stanford, K. (2009). *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*, Oxford: Oxford University Press.
- Stigler, G. J. (1971). The Theory of Economic Regulation. *The Bell Journal of Economics and Management Science*, 2(1), 3–21.
- Sutharshan, B., Mutyala, M., Vijuk, R. P., & Mishra, A. (2011). The AP1000™ Reactor: Passive Safety and Modular Design. *Energy Procedia*, 7: 293–302. <https://doi.org/10.1016/j.egypro.2011.06.038>.
- Taebi, B. (2017). Bridging the Gap between Social Acceptance and Ethical Acceptability. *Risk Analysis*, 37(10), 1817–1827.
- Von Hippel, F. (2011, March 23). It Could Happen Here. *The New York Times*.
- Wald, M. L. (2013, August 27). Vermont Yankee Plant to Close Next Year as the Nuclear Industry Retrenches. *The New York Times*. Retrieved from <http://www.nytimes.com/2013/08/28/science/entergy-announces-closing-of-vermont-nuclear-plant.html>
- Watson, S. R. (1994). The meaning of probability in probabilistic safety analysis. *Reliability Engineering & System Safety*, 45(3), 261–269.
- Weart, S. R. (1988). *Nuclear fear: a history of images*. Cambridge, Mass.: Harvard University Press.
- Wellock, T. R. (2017). A Figure of Merit: Quantifying the Probability of a Nuclear Reactor Accident. *Technology and Culture; Baltimore*, 58(3), 678–721.
- Wheatley, S., Sovacool, B., & Sornette, D. (2017). Of Disasters and Dragon Kings: A Statistical Analysis of Nuclear Power Incidents and Accidents. *Risk Analysis*, 37(1), 99–115.
- WNN. (2016, January 8). IRSN sees problems in assessing passive safety systems. *World Nuclear News*. Retrieved from <http://www.world-nuclear-news.org/RS-IRSN-sees-problems-in-assessing-passive-safety-systems-0801165.html#%2EVpJWxCiVNyI%2Elinkedin>
- Wittgenstein, L. (2001) [1953]. *Philosophical Investigations*. Blackwell Publishing. London.
- Wu, J. S., & Apostolakis, G. E. (1992). Experience with probabilistic risk assessment in the nuclear power industry. *Journal of Hazardous Materials*, 29(3), 313–345.
- Wynne, B. (1988). Unruly Technology: Practical Rules, Impractical Discourses and Public Understanding. *Social Studies of Science*, 18, 147–67.
- Wynne, B. (2003). Seasick on the Third Wave? Subverting the Hegemony of Propositionalism. *Social Studies of Science in Social Studies of Science* 33 (3): 401–17