



Anatomy of a Disaster: Why Some Accidents Are Unavoidable

John Downer

Anatomy of a Disaster: Why Some Accidents Are Unavoidable

John Downer

Abstract.....	1
Introduction.....	2
1. <i>Aloha 243</i>	3
2. Disaster Theory.....	5
3. Normal Accidents	6
i. Normal Accidents are unpredictable and unavoidable.	7
ii. Normal Accidents more likely in ‘tightly-coupled’, ‘complex’ systems.	8
iii. Normal Accidents are unlikely to reoccur.	9
iv. Normal Accidents rarely challenge established knowledge.	9
v. Normal Accidents are not heuristic.....	9
4. The Epistemology of Failure	10
5. <i>Aloha 243</i> Revisited.....	12
Tests & Theories	15
6. Epistemic Accidents.....	17
i. Epistemic Accidents are unpredictable and unavoidable.	18
ii. Epistemic Accidents are more likely in highly innovative systems.	19
iii. Epistemic Accidents are likely to reoccur.	19
iv. Epistemic Accidents challenge design paradigms.	19
v. Epistemic Accidents are heuristic.....	19
Epiphenomena.....	20
Discussion.....	20
Caveats and qualifications	21
The Politics of Blame.....	21
The Limits of Control	22
References.....	24

The work was part of the programme of the ESRC Centre for Analysis of Risk and Regulation.

Published by the Centre for Analysis of Risk and Regulation at the
London School of Economics and Political Science
Houghton Street
London WC2A 2AE
UK

© London School of Economics and Political Science, 2010

ISBN 978-0-85328-403-1

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, without the prior permission in writing of the publisher, nor be otherwise circulated in any form of binding or cover other than that in which it is published and without a similar condition including this condition being imposed on the subsequent purchaser.

Printed and bound by Kube, March 2010

Anatomy of a Disaster: Why Some Accidents Are Unavoidable

John Downer

Abstract

This paper looks at the fateful 1988 fuselage failure of Aloha Airlines Flight 243 to suggest and illustrate a new perspective on the sociology of technological failure and the question of whether such failures are potentially avoidable. Drawing on core insights from the sociology of scientific knowledge, it highlights, and then challenges, a fundamental principle underlying our understanding of technological risk: idea that ‘failures’ always connote ‘errors’ and are, in principle, foreseeable. From here, it suggests a new conceptual tool for Disaster Theory, by proposing a novel category of man-made calamity: what it calls the ‘Epistemic Accident’. It concludes by exploring the implications of Epistemic Accidents and sketching their relationship to broader issues concerning technology and society, and social theory’s approach to failure.

If you can meet with Triumph and Disaster / And treat those two impostors just the same...
~Rudyard Kipling

Introduction

‘The more human beings proceed by plan, the more effectively they may be hit by accident.’ Or so wrote Friedrich Dürrenmatt, the late Swiss playwright. The line is memorable for its counter-intuitiveness: modern societies invest heavily in the belief that good plans will protect them from accidental disasters. Nowhere is this more true than with complex and potentially dangerous technologies, such as nuclear reactors and civil aircraft. Such technologies cannot be allowed to fail, and so we plan them meticulously, and invest enormous effort in testing and analyzing those plans. Then, when accidents come, as they invariably do, engineers return to their drawing-boards to remedy the flaws in their blueprints. As Hutter and Power (2005: 1) put it: there is ‘...widespread recognition that disasters and accidents are in a very important sense organized’ and a corresponding belief that, in principle, it should be possible to ‘organize’ them out of existence. Together, these ideas are central to what Jasanoff (1994; 2005: 11) calls our ‘civic epistemology’ of technological risk, and are deeply implicated in societal choices about technology.

This paper challenges these ideas. It draws on the sociology of scientific knowledge, first to highlight and challenge a fundamental principle underlying our understanding of technological risk, and then to contribute to the Disaster Theory by proposing a new category of man-made calamity: what it will call the ‘Epistemic Accident’.

To outline the argument in brief:

1. The paper begins by outlining a well-known 1988 aviation misadventure: the dramatic final flight of *Aloha Airways Flight 243*, and its subsequent investigation by the US National Transportation Safety Board (NTSB).
2. It then draws on this accident to outline and illustrate the social study of the causes of man-made accidents (or ‘Disaster Theory’). In particular, it explores Normal Accident Theory (NAT), which it argues to be a significant exception within the broader field in certain meaningful respects. (Most significantly in that it suggests that some accidents are, in principle, unavoidable).
3. Having sketched the broad contours and ontological premises of Disaster Theory, the paper then introduces the sociology of scientific and technical knowledge, and invokes it to contest a core premise of Disaster Theory. More specifically, it challenges the implicit ‘rational-philosophical model’ of engineering knowledge, which assumes that engineering facts are, in principle, objectively ‘knowable’, and that ‘failures’ are theoretically foreseeable.
4. The paper then illustrates the above argument by revisiting *Aloha 243* in more depth: examining the NTSB’s conclusions in light of the sociology of knowledge, and exploring the accident’s ‘engineering-level’ causes in more detail. It concludes that, on close inspection, the accident fits awkwardly into existing sociological models of disaster.

5. Considering this disjuncture, the paper then suggests a new category of disaster: the 'Epistemic Accident'. This reconciles the sociology of disaster with the sociology of knowledge by arguing that some accidents are caused by engineering beliefs that prove to be erroneous, even though those beliefs are logical and well-founded. Having outlined Epistemic Accidents, it then explores their implications by systematically comparing and contrasting them with the properties of Normal Accidents, outlined in an earlier section.

6. Finally, the paper suggests how the idea of Epistemic Accidents offers an important critical perspective on far-reaching civil narratives about technological risk and disaster.

And so to 1988...

1. *Aloha 243*

On April 28, 1988, *Aloha Airlines Flight 243*, a passenger-laden Boeing 737, left Hilo airport on a short hop between Hawaiian islands, climbed to an altitude of 24,000 feet, and fell apart. The pilots would later report a violent lurch followed by a tremendous 'whooshing' of air that ripped the cabin door from its hinges, affording the copilot a brief but surely memorable glance of blue sky where she expected to see first class. A closer look would have revealed first-class passengers silhouetted absurdly against the emptiness, still strapped to their seats but no longer surrounded by an actual airplane. All of them perched on a nightmarish roller-coaster, hurtling through the sky at hundreds of miles per hour with the ocean far below.

The pilots, although confused by the chaos, found they were still in control of the aircraft. And so -- unable to communicate with air traffic control over the howling winds -- they set the emergency transponder and landed as abruptly as they dared at nearby Kahului Airport (NTSB 1989: 2).

The airplane's condition when it arrived at Kahului astonished the aviation community and immediately entered the annals of engineering folklore. An eighteen-foot strip of the fuselage, thirty-five square meters, had completely torn away from the airframe, like the top from a sardine can, severing major structural beams and control cables and exposing passengers to the sky on all sides (see fig. 1).

Fierce winds and flying debris had injured many of those on board but, by grace and safety belts, only one person had died: senior flight attendant Clarabelle Lansing, who



Fig. 1. *Aloha Airlines Flight 243* (Source: Hawaii State Archives)

disappeared into the void as the fuselage tore.¹ The airframe itself was terminal, however: never before or since has a large civil aircraft survived such a colossal insult to its structural integrity.

The US National Transportation Safety Board (NTSB) promptly launched an investigation to determine the cause. When their report appeared the following year it inculpated a fateful combination of stress fractures, corrosion and metal fatigue. It concluded that imperfections in the layered aluminum fuselage had allowed salt water to creep between the sheets, corroding the metal panels and forcing them apart.² Over time, this ‘disbonding’ stressed the rivets binding the fuselage together and fostered fatigue cracks in the sheets themselves, eventually causing the entire structure to fail.

The ‘causes’ of accidents are invariably nested, however -- as Jasanoff (2005) illustrates -- and the above explanation raised as many questions as it answered. Structural failures of this magnitude were not supposed to happen so suddenly, whatever the underlying cause. The aviation world had long understood that some early 737 fuselages had imperfections that induced skin cracking, but they also understood that fatigue cracks progressed gradually enough for routine inspections to raise alarms long before a fuselage could even come close to rupturing.

Faith in the airplane’s structural integrity was further buttressed by the understanding that, even if inspections fell spectacularly short and the fuselage ruptured, then the rupture would be




Fig 2. Boeing 737 Fuselage (Source: NTSB 1989)

modest and relatively safe. This belief was premised on the ‘fail-safe’ design of the 737’s metal skin. The airplane’s architects divided its fuselage into small (10-inch by 10-inch) panels, each bounded by ‘tear straps’ designed to constrain ruptures by channeling and redirecting cracks (see fig. 2). In theory, a tear or rupture would cause the fuselage to ‘flap’ open around a single panel, releasing internal pressure in a way that was limited, controlled and -- importantly -- not threatening to the fuselage’s basic integrity (NTSB: 34). The fail-safe design was classic ‘design-in-depth’ aviation engineering: the failure of one panel would have been a clear alarm but would not have endangered the airplane’s structural integrity. For extra security, the design allowed for cracks of up to 40 inches encompassing two panels simultaneously. (This reflected an

¹ Sixty-five of ninety passengers and crew were injured, eight seriously. The event would later be dramatized in a 1990 TV movie – ‘Miracle Landing’ – starring Wayne Rogers and Connie Sellecca as the pilots, and Nancy Kwan as star-crossed Clarabelle.

² The bonding process Boeing used on its early 737s relied on an epoxy glue that came on a ‘scrim’ tape. Assembly workers had to keep the tape refrigerated to suspend the glue’s adhesive reaction until it was in place, then allow the glue to ‘cure’ by letting it warm. If the tape was too cold when it was applied it gathered condensation that impaired proper adhesion. If it got too warm then it began to cure too early, again impairing adhesion. Even under Goldilocks conditions, the epoxy tended to bind to oxide on the surface of the aluminum rather than to the metal itself, another difficult issue (Aubury 1992).

understanding that the worst ruptures would come from a broken engine blade ripping through the fuselage, and an estimate of how much damage this might cause.) (NTSB 1989: §1.17.2).

Arguably the most significant level of causation in this accident, therefore, lies in the failure of the fail-safe design.³ Yet social scientists would almost certainly propose a further level. They routinely look beyond engineering explanations altogether, by locating accidents and the actions that contribute to them, in wider social contexts, to which they attribute an ultimate degree of causation. This broad and heterogeneous endeavor is usually referred to as 'Disaster Theory'.

2. Disaster Theory

Until the late 1970s, the question of what caused technological disasters belonged almost exclusively to engineers. By the 1980s, however, social scientists had begun to recognize that such accidents had social and organizational dimensions. The British sociologist Barry Turner (1976, 1978) investigated a series of 'man-made disasters' and noted that they were invariably preceded by a legacy of unheeded warnings that, if acted on, would have averted a misadventure. This realization -- that experts were overlooking clear and intelligible danger signals -- begged the question: 'Why?'. It made accidents a legitimate object of social enquiry by recasting them as 'social' rather than 'engineering' problems. Turner thus began thinking of accidents as 'failures of foresight' and began a project to explain why their warnings go unnoticed. This was important because a disaster with warning signs is, in principle, a disaster that might be avoided. (Even if, as Turner himself argued, there are often intractable problems associated with doing this *in practice*).

Turner's insight was compelling and, in his wake, social scientists from a range of disciplines -- organizational sociologists, psychologists, political scientists and others -- began scrutinizing accidents in the collective endeavor that became known as Disaster Studies. A group centered on the University of California at Berkeley, for instance, pioneered an approach known as High Reliability Theory (HRT), which looked for the distinctive characteristics of organizations -- such as aircraft carriers and air traffic control systems -- that consistently perform reliably in demanding circumstances (and so, by deduction, for the hallmarks of systems likely to fail) (e.g. Rochlin et al. 1987; La Porte & Consolini 1991; Roberts 1993; Wieck 1987). Psychologists, such as Reason (1990), began looking for auguries of disaster by isolating the conditions that foster human error. Diane Vaughan (1996), notably, drew on the sociology of relationships to frame 'failures of foresight' in terms of the incremental acceptance of warning signals, or what she called the 'normalization of deviance.' Not to mention a constellation of others (many of them as notable), who started pursuing the same end by varying means.⁴

I will henceforth refer to this body of scholarship as 'Disaster Studies'. Collectively, these works endeavor to hone what Vaughan (2005: 63) calls our 'technologies of control' -- the models and

³ The convention in US accident reports is to apportion blame (Galison 2000), and, accordingly, the NTSB pointed to Aloha Airlines's 'deficient' maintenance program (NTSB 1988: §2.3), which, in keeping with every close examination of technological practice, showed itself to be 'messier' than its idealized portrayal (Wynne 1988; Langewische 1998). Aloha protested that the airplane was compliant with all the relevant FAA mandates, and that the report gave a false impression that their maintenance practice was untypical (Cushman 1989). It is probably telling that the report led to few, if any, sanctions on the airline.

⁴ With some luminary stars that this narrative omits reluctantly, because of its constraints.

practices by which organizations govern complex systems -- by looking for errors that can then be corrected. All of them implicitly concede Turner's basic tenet: that accidents are '*foresight failures*', in that they are preceded by auguries that organizations could, *in principle* (even if not in practice), recognize in time to prevent catastrophic accidents.⁵

One major strand of Disaster Theory breaks with this ontological premise, however, and approaches disasters from a different perspective: Normal Accident Theory (NAT).

3. Normal Accidents

NAT, first proposed by Yale sociologist Charles Perrow, and expounded in his (1984 [1999]) book *Normal Accidents*, is a major strand of Disaster Studies but a signal exception to the 'foresight-failure' model.⁶ Perrow's argument is that seemingly trivial events and non-critical failures sometimes interact in unexpected ways that thwart the very best engineering designs and cause catastrophic system-wide failures. He calls these failures 'Normal Accidents', (or, sometimes 'System Accidents').⁷

No complex system can function without trivial errors: that is, with absolutely no deviation from platonically perfect operation. They must tolerate minor irregularities and small component failures: spilt milk, small sparks, blown fuses, misplaced luggage, stuck valves, obscured warning lights and suchlike. Engineers might not like these events, but eliminating them entirely would be entirely impracticable, and they are considered irrelevant to system safety.⁸ Hollywood might enjoy inventing scenarios where a succession of such events in perfect synchrony lead inexorably towards a disaster, but in practice the odds against such scenarios are astronomical: literally billions-to-one against, far too insignificant to affect an engineering calculation.

Perrow's core insight, however, is that even these 'trivial' errors -- the background noise of normal technological practice -- have inherent catastrophic potential. This is because where potentially dangerous systems have millions of interacting parts that allow for billions of unexpected events and interactions, then some fateful billion-to-one coincidences are almost inevitable (or 'normal'). Nuclear reactors fit this description perfectly and Perrow suggests that the 1979 near-catastrophe at Three Mile Island (TMI) is an exemplary Normal Accident. By his account, the incident began when leaking moisture from a blocked water filter inadvertently tripped valves controlling the flow of cold water into the plant's cooling system. Redundant backup valves, that should have intervened, were also inexplicably closed, which should have

⁵ Although Vaughan (2005: 64) suggests that disaster studies might benefit from more micro-level research into how 'signals of danger' are identified, categorized and defined.

⁶ Normal Accidents, the concept, are the heart of Normal Accidents the book. If this sounds tautologous, realize that only relatively small fraction of the book focuses on this specific issue. Perrow draws on evidence from a wide range of domains (nuclear power; shipping; aviation; mining and more) to make many technological, sociological, psychological and political observations about technological accidents. Despite these variegated insights, however, the Normal Accident is the book's leitmotif and signal contribution to Disaster Studies.

⁷ He contrasts these accidents with what he calls 'component failure accidents,' which he defines as accidents where one failure causes other failures through predictable and linear relationships between the elements of a system (Such as when an electrical surge affects the wiring in a civil aircraft.) (Perrow 1991: 70).

⁸ Engineers and policy-makers do envisage 'failure-free systems', but only as systems that can to err without failing. These systems have redundancies and fail-safe elements that accommodate the vagaries of normal technological practice (Downer 2009).

been clear from a (rarely-needed) indicator-light in the control room, if it had not been obscured by a tag hanging from a switch above. The final line of technological defense was a tertiary system: a relief valve which should have opened but did not, while a malfunctioning indicator-light erroneously indicated that it did. None of these failures were particularly noteworthy in themselves; complexity simply colluded with cruel coincidence, and the reactor's controllers understandably struggled to comprehend its condition in time to prevent a catastrophic meltdown.

NAT is important because Normal Accidents are impossible to predict. This is in principle as well as in practice. The type of errors that combine to forge these accidents are not significant enough, in themselves, to make them distinguishable on an ontological level. Investigators can note the significance of specific events in retrospect, as Turner observed, but their significance is an artifact of hindsight. Without this vantage there is nothing to distinguish them from a million other unavoidable technical idiosyncrasies or 'deviances' that constitute the 'messiness' (Wynne 1988) of a typically functional technological system.⁹

One highly significant but under-recognized consequence of suggesting that some failures occur without meaningful errors is that NAT breaks, ontologically, with the 'foresight-failure' premise. This puts NAT at odds with wider Disaster Theory, which implicitly assumes that accidents, whatever their cause, are preceded by warning signals, and then looks further downstream to the social question of why those warning signs go unheeded.

This has far-reaching consequences. Perrow describes NAT as an effort to 'chart the world of organized systems, predicting which will be prone to [Normal] accidents and which will not' (1991: 62).¹⁰ But what is the world that Perrow charts? Which systems are vulnerable to Normal Accidents? What are NAT's implications for Disaster Studies? These questions can be addressed by highlighting five distinct but interrelated properties of Normal Accidents:

i. Normal Accidents are unpredictable and unavoidable.

NAT, as outlined above, argues that some accidents occupy a blind-spot in our 'technologies of control' because they are failures without conspicuous errors. Since no single failure in a Normal Accident -- water filter, coolant valve, warning light, etc. -- falls outside the accepted definition of 'typical' technological practice (in all its unavoidable messiness), then none can constitute a distinguishable 'signal' or 'deviance', and the 'foresight-failure' model that presupposes that accidents are somehow 'allowed' to happen (because warning signals are missed, ignored or normalized) does not apply.

⁹ Hindsight might explain the injury that foresight would have prevented, but hindsight, as the proverb goes, is 'a comb the world gives to men only after they are bald'.

¹⁰ Idiosyncratically, perhaps, for the work of a sociologist, NAT is equally applicable with or without a social component. 'Perhaps the most original aspect of the analysis,' as Perrow puts it, 'is that it focuses on the properties of systems themselves, rather than on the errors that owners, designers, and operators make in running them' (1991 63). It is, as Hopkins (2001: 65) points out, 'an unashamedly technological determinist argument'. Perrow describes it as a 'first attempt at a 'structural' analysis of risky systems' (1991: 63).

The ontological unpredictability of Normal Accidents makes them unavoidable.¹¹ *Normal Accidents*, the book, contains many suggestions as to how designers might reduce the risk of technological disasters, but these are secondary to its central thesis. Perrow's core argument, NAT, is unequivocal: it is that disasters are an inevitable property of certain technological systems, no matter how well managed. 'In addition to being unforeseeable, incomprehensible and not amenable to knowledgeable intervention,' he writes, 'the Normal Accident cannot be prevented' (1982: 176).

Hopkins (1994) criticizes NAT for offering little insight into accident prevention, but herein lies a core virtue of Perrow's insight: it forces analysts to recognize, *a priori*, that not all accidents are preventable. The only absolute and inviolable defense against Normal Accidents is to avoid building certain systems altogether.

ii. Normal Accidents more likely in 'tightly-coupled', 'complex' systems.

NAT offers two yardsticks for assessing the vulnerability of systems to Normal Accidents: their 'complexity', and the degree to which they are 'coupled'. Both are derived from the organization and interaction of elements in a system.

'Complexity' here is a diffuse measure encompassing factors such as the number of elements in a system and their relationship to each other. Complex systems, by this definition, have many interacting elements. (So a sheet of advanced composite laminate material is 'simple' even though it might be very sophisticated.) The number of interacting elements is significant, but so too is their manner of interaction. Perrow compares 'complex' with 'linear' systems, explaining that interactions in linear systems are 'expected', 'familiar' and 'quite visible even if unplanned', whilst those in complex systems are 'unfamiliar,' 'unexpected' and '[...] either not visible or not immediately comprehensible' (1984: 78). A flight-control system (or 'autopilot') is much more 'complex' than an assembly line, for instance, because, even though both consist of many elements, the latter is very linear (because its interactions happen in a simple sequence), whereas the former is characterized by 'pooled interdependence'.

'Coupling', meanwhile, is a measure of the 'slack' in a system. A system is 'tightly-coupled' if interactions between its elements happen rapidly, automatically and inflexibly, with little room for intervention or adjustment by human actors. 'Loosely coupled' systems, by contrast, interact more slowly and flexibly, offering both time and opportunity for intervention. Universities, for instance, are 'complex', with their many interacting elements, but 'loosely coupled' because the failure of a sociology department, although unfortunate, is unlikely to immediately or uncontrollably affect a physics department or a football team. Assembly lines, by contrast, are relatively 'simple', because of their linear, sequential and predictable interactions, but 'tightly coupled' because a single point of failure will halt the entire system.

¹¹ Engineers could design a near-perfect, ten-million-dollar safety valve and thereby prevent TMI's specific cascade of errors, but they could not give every component the same treatment.

Both high complexity and tight coupling contribute to Normal Accidents, so a grid with each on different axes offers a simple framework for determining their likelihood. Perrow argues that technologies such as nuclear power plants and large aircraft occupy the star-crossed corner-point on this grid, where high complexity meets tight coupling and disasters are especially likely.¹²

The other three characteristics of Normal Accidents follow from the same logic of understanding them as ‘perfect storms’ of otherwise insignificant events, but are less explicit in Perrow’s writing and less explored in the Disaster Studies literature more generally. They can be outlined relatively simply:

iii. Normal Accidents are unlikely to reoccur.

That is to say: ‘specific’ Normal Accidents are highly unlikely to reoccur. Where there are a billion possible ‘billion-to-one events’ that can instigate an accident, then it is logical to anticipate an accident, but not the same accident twice.¹³ The exact confluence of faulty valves, miswired warning lights, and errant tags that were implicated in TMI might never reoccur if the same plant ran, unaltered, for ten thousand (‘groundhog’) years.

iv. Normal Accidents rarely challenge established knowledge.

Normal Accidents do not challenge common engineering understandings and theories about the world: their ‘design paradigms’ (Petroski 1994). This is because the errors that combine to produce Normal Accidents are rarely surprising in themselves. A stuck valve usually reveals nothing about valves, for example, and would do almost nothing to challenge the engineering underlying their design. Valves sometimes stick; this is why they come with warning lights and redundancies.

This leads naturally to the final point:

v. Normal Accidents are not heuristic.

A close corollary of point 4 above -- that Normal Accidents do not challenge engineers’ understandings of the world -- is that they have minimal heuristic value. Our ‘technologies of control’ have little to learn from Normal Accidents. Insofar as they are lessons at all, then it is always the same lesson: that unpredictable, improbable and disastrous confluences of errors are likely to occur in complex and tightly-coupled systems.

These elementary properties of Normal Accidents set NAT apart from other strands of Disaster Theory, which, as discussed above, are deeply wedded to an ontological understanding of

¹² This helps explain the oft-overlooked virtues of straightforward and unadorned systems: offering a useful counter-narrative to a worldview that often equates intricacy with progress.

¹³ Perrow offers a few examples of reoccurring Normal Accidents in his book, such those related to a flawed cargo door latch on the DC-10 airplane (1990: 137-140). On close consideration, however, these accidents are not Normal Accidents in the strict sense that he describes. The DC-10 latch, for instance, was a single point of failure and the accidents related to it are best understood as ‘component failure accidents’ in Perrow’s taxonomy.

technological accidents as -- in principle, if not in always practice -- predictable, preventable and heuristic.

This disparity leads some Disaster Theorists (e.g. Sagan 1993) to argue that NAT is antithetical to other approaches, but this is misleading. NAT is ontologically different from other perspectives, as argued above, and drawing contrasts can be instructive, as Sagan shows; yet there is room for different perspectives to co-exist. Few Disaster Theorists, when pushed, contest the logic of NAT or the existence of Normal Accidents. Instead, criticisms tend revolve around the scope and prevalence of Normal Accidents, the utility of NAT, or its socio-political ramifications (e.g. La Porte, 1994; Perrow, 1994; La Porte & Rochlin, 1994; Hopkins 1999).¹⁴ Perrow, meanwhile, (like other advocates of NAT), does not suggest that all accidents are Normal. Indeed, he argues that Normal Accidents are not the norm. (The term ‘normal’, here, connoting ‘inevitability’ rather than ‘commonness’ (Perrow 1984: 174).) He credits Three Mile Island as a Normal Accident, but suggests that many other signal technological disasters of the last century -- Challenger, Bhopal, Chernobyl -- were not.

NAT has limited scope, therefore, but rather than viewing this as a ‘major limitation,’ as Hopkins (1999: 95) does, we reconcile NAT with other approaches by saying it delineates an important but restricted category of ‘unavoidable’ accidents, to which other approaches do not apply. The fact that it does not fit all accidents, or even most, means it need not detract from efforts (such as those of the High Reliability Theorists) to interrogate, theorize and prevent ‘non-Normal’ accidents.

All Disaster Theorists, including Perrow, therefore, implicitly hold that all non-Normal accidents are, in principle, foreseeable. Which is to say that the only barriers to avoiding them are social, psychological and organizational. These non-Normal accidents include those that are caused by technological faults but which do not qualify as ‘Normal’ because they involve linear and predictable interactions: what Perrow calls ‘Component Failure Accidents’ (Perrow 1999 [1984]: 70-71). These accidents, he argues, are caused by institutional shortcomings and so fit squarely into the sphere of theoretically avoidable disasters: if engineers are perfectly rigorous with their tests, thorough with their inspections and assiduous with their measurements, then such accidents should not happen.

This claim is wrong. Some accidents are not ‘Normal’ in the Perrovian sense, but not avoidable either. To understand why, it helps to look at the epistemology of failure.

4. The Epistemology of Failure

If a bridge collapses or an airplane fuselage breaks apart because it experiences forces beyond those its designers anticipated, then it is easy to interpret the accident as an engineering error. Such accidents invariably reveal flaws in engineering assumptions, models, or data, but

¹⁴ Another criticism observes that the ‘Normal-ness’ of any given accident, is open to interpretation and contestation (Hopkins 1994). The Normal Accident is best understood as an ‘Ideal Kind’, however, and its ambiguity in specific cases need not detract from its usefulness as a conceptual lens through which to view technologies and the causes of disaster. Critics might debate the applicability of NAT in specific instances, therefore, without contesting the argument in principle. Beauty might be in the eye of the beholder but this hardly negates its existence, its value as an explanatory category or its tangible consequences.

engineers work in an empirical realm of measurable facts. Facts are knowable. Facts are binary. True or false. Ontologically distinct. So when the facts are wrong, this wrongness, and any disaster it contributes to, can be viewed as a methodological, organizational, or even moral failing: one that proper engineering discipline should have avoided and one that social theorists might one day prevent.

This picture of facts is rooted in what the sociologist Harry Collins (1992 [1985]: 185) has called the ‘canonical rational-philosophical model’ of expert knowledge. This model construes engineering as a process governed by formal rules and objective algorithms that promise incontrovertible, reproducible and value-free facts grounded in measurements rather than expert opinion.

Porter (1995) calls this ‘the ideal of mechanical objectivity’. It is intuitive, convincing, and entirely rejected by most epistemologists.

The idea that erroneous knowledge claims are fundamentally and recognizably distinct from those that are true (such that perfect tests and calculations should be able to eliminate errors entirely) was a cornerstone of Western philosophy for centuries. Logical philosophers -- from Francis Bacon, through William Whewell to Karl Popper and beyond -- continuously honed the ‘scientific method’ to ensure it led ineluctably towards truth: fiercely debating the nature of proof, the foundations of evidence and the essence of facts.

The several hundred years that separate Bacon and Popper, however, speak to the elusiveness of the ‘ideal experiment,’ ‘perfect proof’ and ‘indubitable fact’. Indeed, beginning with Wittgenstein’s later work, logical philosophers started rejecting the entire enterprise. ‘Finitist’ philosophers, such as David Bloor (1976), and historians, such as Thomas Kuhn (1962), began to argue that no facts -- even scientific or technological -- are, or could be, completely and unambiguously determined by logic or experiment. Today, few doubt that the canonical rational-philosophical model of scientific knowledge is inadequate for explaining science in action, and the ‘facts’ it produces.

As Pinch (1991: 151) observes, therefore, it is ‘very tame’ of social scientists to merely accept the rational-philosophical vision of technological knowledge. In fact, there is no need for it to do so because in the mid 1970s Bloor (1976) made an argument for a sociology of scientific knowledge (as opposed to a sociology of scientific practice) built on the idea that -- from the perspective of the actors who define them -- ‘true’ beliefs can be ontologically indistinguishable from those that are ‘false’.¹⁵

Bloor’s finitist argument was influential and, over roughly same period as some sociologists were encroaching disaster investigations, others began to explore the properties and implications of scientific and technological knowledge.¹⁶ Through a series of epistemologically-conscious ethnographies (or ‘epistemographies’), they convincingly demonstrated that the seemingly

¹⁵ The argument here is not, as critics often claim, that there are no deep ontological truths about the world, just that actors can never access them directly because their access is filtered through perceptions, theories and preconceptions that undermine appeals to ‘objectivity’ and add an indelible ‘social’ dimension to all knowledge-claims.

¹⁶ This endeavor eventually grew to encompass other social sciences, became the focus of specialist research centers and departments, and is now usually referred to as ‘Science and Technology Studies’ (STS) (e.g. Hackett et al. 2008).

abstract concerns of philosophers have very tangible consequences in practice. Successive studies, for instance, illustrated the surprising degree to which credible and informed experts often disagree over seemingly objective ‘facts’ and the frequency with which expert communities reverse their opinion on well-established and apparently inviolable ‘truths’ (e.g. Collins 1985; Collins & Pinch 1993, 1998; Latour 1987).

A subset of these studies look directly at engineering knowledge (eg: Pinch & Bijker 1984; Collins & Pinch 1998; MacKenzie 1996; Bijker, et. al. 1987; Latour 1996). They subvert the idea of inviolable engineering ‘facts’ with the observation that engineering’s orderly public image belies a ‘messy reality’ of real technological practice. There is more to understanding technology, they argue, than can be captured in standardized tests. Or, as Wynne (1988: 153) puts it:

In all good ethnographic research [on] normally operating technological systems, one finds the same situation. [...] Beneath a public image of rule-following behavior [...] experts are operating with far greater levels of ambiguity, needing to make uncertain judgments in less than clearly structured situations.

Sociologists of engineering knowledge explore the epistemology of bench tests, much as sociologists and philosophers of science examine laboratory experiments. They look to how engineers extrapolate from tests to the real world, and from past to future performance, and they leverage epistemic dilemmas, such as the ‘problem of relevance’ or the ‘experimenter’s regress’ (Pinch 1993; Collins 1985), to argue that tests cannot ‘objectively and definitively’ interrogate a technology or reveal the ‘truth’ of its functioning (Pinch 1993).¹⁷ By this view, technological ‘facts’ are best understood as ‘...hypothes[e]s to be tested first on paper and possibly in the laboratory, but ultimately to be justified by [their] performance’ (Petroski 1992: 104).

Although it is rarely remarked upon, these studies imply a finitist account of technological failure and a new understanding of disaster. To appreciate this argument and its implications, it first helps to look again at *Aloha 243*.

5. *Aloha 243* Revisited

Aloha 243, by any account, was *not* a Normal Accident. It was the product of the failure of a single major structural element, the fuselage, from a known cause, fatigue, rather than a confluence of ‘trivial’ errors compounding each other and cascading confusingly towards disaster. In Perrowian terms it was a Component Failure Accident, and so, by the logic of NAT, it falls outside Perrow’s bounded category of unavoidable accidents. Simply put, all of disaster theory suggests that, in principle, it should have been both preventable and foreseeable. In fact, *Aloha* was neither. But it was illustrative of the epistemic limitations that underly all engineering knowledge.

¹⁷ Many of these findings are echoed in introspective studies by engineers. Henry Petroski (1994; 1992), for instance, has written extensively about the inevitable limits of engineering knowledge.

When *Aloha*'s fuselage failed in a way that completely circumvented the 'fail-safe' tear-strips designed to prevent major ruptures, it challenged basic aeronautical engineering beliefs about the propagation of cracks, and revealed shortcomings in fundamental assumptions about metal fatigue and its relationship to aircraft manufacture.

A range of intricate engineering theories -- what Petroski (1994) calls 'design paradigms' -- framed common understandings of the 737's metal skin. Intricate formulae defining the tensile strength of aviation-grade aluminum, for instance, were a lens through which engineers could interpret its design and predict its failure behavior. One such paradigm concerned what is known as Multiple Site Damage (MSD): microscopic cracking that appears around fastener holes as a fuselage fatigues. 'The *Aloha* accident stunned the industry,' as report put it, 'by demonstrating the effects of undetected multiple site damage.'¹⁸

At the time of the accident, neither Boeing nor the airlines considered MSD a particularly significant safety issue. This was because, even though it was virtually undetectable at an early stage, engineers widely believed that no MSD crack could grow from an imperceptible level to 40 inches (the fuselage 'fail-safe' design level) in the period between maintenance inspections.

Fatefully, however, this assumption failed to recognize that, in certain areas of the fuselage, and under certain conditions,¹⁹ MSD had a tendency to develop along a horizontal plane between a line of rivets, such that a string of almost imperceptible but adjacent cracks could abruptly coalesce into one big crack, longer than 40 inches, that nullified the fail-safe tear straps (see fig. 3).²⁰ As the NTSB report puts it:

'...the Safety Board [...] conclude that, at the time of the accident, numerous fatigue cracks in the fuselage skin lap joint [...] linked up quickly to cause catastrophic failure of a large section of the fuselage. The Safety Board believes that sufficient fatigue cracking [...] in the lap joint [served] to negate the design-intended controlled decompression of the structure (NTSB 1989: §2.3).

But why did this come as a surprise? It is easy to assume that aeronautical engineers should have understood the fatigue behavior of airplane fuselages, given the accumulated metallurgical research available to them. After all, the prevailing orthodoxy was hardly idle: it was grounded in extensive laboratory experiments and decades of experience




Fig 3. MSD cracks (Source: NTSB 1989)

¹⁸ 'Applying Lessons Learned From Accidents Accident/Incident Summary' Available online:[AlohaAirlines_B737_Flight243_Accident_Summary_RV1-A.doc]

¹⁹ Specifically, when there was significant disbonding between fuselage sheets and subsequent corrosion.

²⁰ This is, by necessity, a simplification of a complex and multifaceted engineering explanation; albeit hopefully not a distorted one in respect to the central argument being illustrated. A fuller account (including a more recently proposed alternate, 'fluid hammer', hypothesis) would complicate the narrative without undermining the main premise.

with aluminum airplanes.

Engineers who work closely with metal fatigue, however, are generally unsurprised by its 'surprises'.²¹ When fatigue felled a UK Royal Air Force transport aircraft outside Baghdad in January 2005, for instance, a Ministry of Defense spokesperson echoed a long-standing engineering lament: 'It's hard to believe that there might be a structural failure after all the close monitoring we do,' he said, 'but there is a history of surprises about metal fatigue' (Air Safety Week 2005).

Sociologists of technology have long noted that 'outsiders' tend to be less skeptical of facts and artifacts than the engineers who produce them, simply because outsiders are not privy to the epistemological uncertainties of knowledge production (MacKenzie 1998). 'Distance lends enchantment,' as Collins (1985) puts it, and metal fatigue is no exception in this respect.²² From afar it looks straightforward, but from the perspective of engineers it is striking in its formidable complexity and ambiguity.

The formal study of metal fatigue is known as 'fracture mechanics'. Its acolytes examine the properties of metal (principally tensile strength and elasticity)²³ and their relationship to different kinds of stress. In the laboratory -- where minimally defective materials in straightforward forms are subjected to known stresses -- fracture mechanics is a complex but broadly manageable science. It boasts carefully-honed models that work tolerably well in predicting where, when and how a metal form will fail (even if the most accurate of them must venture into the rarified algorithms of quantum mechanics). In operational aircraft, however -- where elaborate forms, replete with rivet holes and imperfections, experience variable and uncertain stresses, the vicissitudes of time and the insults of human carelessness -- the calculations involved are vastly more forbidding.²⁴ So much so that engineers confess that even the most elaborate calculations essentially amount to informed guesses (Feeler 1991: 1-2). Even today, some 20 years since *Aloha*, aircraft fatigue management is an inexact science: one that Air Safety Week recently described as 'a dark black art [...] akin to necromancy' (2005).²⁵

This complexity makes fatigue very difficult even to monitor. Most ferrous metals have a 'fatigue limit'; which is to say they do not fatigue at all until they reach a specific stress threshold. Aluminum alloys like those used in the 737, however, have no clear fatigue limit and

²¹ These surprises began in the mid-19th century, when it became apparent that seemingly good machine parts were failing as they aged. By 1849, engineers were researching 'metal fatigue', and by the 1870's had carefully documented the failure behavior of various alloys, even though the causes remained opaque (Garrison 2005). Fatigue retained its capacity to surprise, however, shaking the nascent civil aviation industry almost a century later by felling two airliners (both de Havilland Comets) in the first four months of 1954 (Faith 1996: 158-165).

²² MacKenzie (1998) speaks of a 'uncertainty trough,' based on this principle, which maps confidence in a technology against an actor's epistemic distance from it.

²³ 'Tensile strength' is the measure of a metal's 'resistance to being pulled apart', usually recorded in pounds per square inch. 'Elasticity' is the measure of a metal's ability to resume its original form after being distorted. The point at which this no longer happens being known as the metal's 'elastic limit'.

²⁴ Certain design choices can create unanticipated stress-points in the fuselage that instigate fatigue, as can slight imperfections such as scratch marks or defects in the metal itself. Holes in the fuselage, such as those required by bolts and rivets, are particularly vulnerable because they tend to concentrate stress around their periphery (Feeler 1991: 1-2). All these things are a challenge to accurately model in advance.

²⁵ It is worth noting, moreover, that even a Platonically perfect understanding of fatigue in aluminum alloys would not have been sufficient for understanding fatigue in the 737 airframe. Amongst other things, engineers would also have had to wrestle with the nuances of the adhesives used to bind the aluminum sheets together.

are thought to fatigue, often imperceptibly, at any stress level. Nor is this a linear process. As metals fatigue they retain their original strength until they reach another threshold: a tipping-point where actual cracks begin. Pre-crack fatigue was originally thought to develop at the microscopic level of the metal's crystalline structure, but is now understood to accumulate at the atomic level. At this level, it is imperceptible to maintenance engineers. Modern 'nondestructive evaluation' techniques enable inspectors to detect cracks as small as 0.04 inch, beyond which they are effectively blind (Maksel 2008). As one engineer puts it: 'Until cracking begins, it is for all practical purposes impossible to tell whether it will begin in 20 years, or tomorrow' (Garrison 2005).

Tests & Theories

With such uncertainty at the heart of metallurgy it is understandable that the 737's designers might fail to predict how their airframe would fatigue, but it remains difficult to intuit why their errors did not reveal themselves sooner. Engineers have long understood that they cannot perfectly deduce technological performance from algorithms and blueprints. This is why they perform tests. Yet Boeing extensively tested the 737 fuselage for its fatigue resilience and decompression behavior, under the FAA's keen regulatory gaze. They even acquired and retested an old airframe as some aircraft in service approached their original design life. And in each case, the aircraft performed just as calculations predicted.²⁶

This disparity, between the aircraft's performance in the laboratory and its performance in the practice, reflects a tension between tests and the phenomena they attempt to reproduce. Tests simplify by necessity. They get engineers 'closer to the real world,' as Pinch (1993: 26) puts it, 'but not all the way'. This is because the world is complex and unwieldy, and 'knowing' it requires what Scott (1998: 11) calls 'a narrowing of vision'. Engineers cannot exactly reproduce the real world in their laboratories so they must reduce its variables and make subjective judgements about which are relevant.²⁷

Aloha exemplifies the problems imposed by these necessary reductions. Boeing's understanding of fatigue, for example, led to a belief that the very largest fuselage ruptures would be the result of escaped engine-blades. This, in turn, led them to interrogate the fuselage by 'guillotining' a fully pressurized section with two 15-inch blades. Such tests created punctures smaller than 40 inches (the specified design limit) that traversed one tear-strap, and, as anticipated, allowed the skin to 'flap' safely (NTSB 1989: §1.17.2). Of course, it also left engineers blind to the danger of MSD coalescing into cracks longer than 40 inches.

The airplane's manufacturers and regulators did explicitly test their understanding of fatigue, including MSD, but these tests were similarly framed and constrained by engineering doxa. The aviation community had long understood that the key determinant of fatigue in an aircraft fuselage was its number of 'pressurization cycles' (usually one full cycle for every takeoff and landing). To test the 737's fatigue life, therefore, the designers pressurized and depressurized

²⁶ In general, the structural components of an airplane (such as the airframe and wings) are designed such that 'an evaluation of the strength, detail design, and fabrication must show that catastrophic failure due to fatigue, corrosion, manufacturing defects, or accidental damage, will be avoided throughout the operational life of the airplane.' (NTSB 2006: 37).

²⁷ MacKenzie (1990), for instance, explores high-level doubts about whether successful test launches proved that America's nuclear missiles would hit their targets under wartime conditions.

(cycled) a half-section of a 737 fuselage 150,000 times (representing twice the design-life goal) in a test facility. This produced no major cracks (NTSB 1989: §1.17.2).²⁸

Again, however, the NTSB suggest that this test was unrepresentative because its theoretical foundations were flawed. By isolating ‘cycles’ as the limiting factor in fatigue, the test excluded a range of variables that were significant to the *Aloha* incident. Flight 243 was, to be sure, a highly cycled aircraft (because of its short routes), but there were other factors that contributed to its decay. One was its age: manufactured in 1969, it was one of the longest serving 737s in operation. Another was its operating setting: the warm, saltwater environment of Hawaii, a misery for metal. A third was its flawed construction: as outlined above, imperfect bonding in the airplane’s fuselage allowed saltwater to creep between its alloy sheets. These three factors -- age, environment, and manufacture -- set *Aloha* 243 apart from Boeing’s test fuselage. The disbonding created gaps that allowed saltwater to creep in, and, over time, this corroded the metal in a way that stressed the aircraft’s rivets and nurtured cracks in its fuselage. Unsurprisingly, therefore, it fatigued differently from the new, properly-bonded fuselage that engineers repeatedly pressurized in a dry laboratory.²⁹

In an ideal world, the airplane’s designers would have been able to perform perfectly representative fuselage tests -- recreating every possible combination of environmental condition and manufacturing imperfection. In this ideal world, the tests would have run for decades to accurately reproduce the vicissitudes of time (as some clinical trials do), and would have maximized their statistical representativeness by involving hundreds, if not thousands, of aircraft.

This ideal world was clearly impractical, however. The 737’s designers had to strike a balance between ‘control’ and ‘authenticity’: between making their tests accurate and making them useful (Downer 2007). As we saw above, this balance necessarily involved judgements and assumptions -- for instance, that the airframes would be properly ‘bonded’; that fatigue cracks grow slowly; and that fuselage punctures would stretch no further than 40 inches. These assumptions were based on the best information available, which came from fracture mechanics: a laboratory science with unknowable applicability to operational aircraft, and where experts freely conceded that even the most sophisticated models were little more than ‘informed guesses’. This left ample room for epistemic uncertainties; some of which had consequences. All technologies are ‘real life experiments’, as Weingart (1991) puts it, and it is the nature of experiments that their outcomes are uncertain. ‘Even years of successful application’ writes Petroski, ‘[do] not prove an engineering hypothesis to be valid’ (1992: 104-5).³⁰

From the vantage of hindsight, *Aloha* 243, like most accidents, looked replete with early warnings. Engineers might have spotted and acted on the MSD cracks, for instance, had they known to look and how to interpret what they found. Lacking this knowledge, however, such ‘auguries’ could hide in plain sight: invisible even to experts who were literally staring straight at them. Just as if *Aloha* 243 were a Normal Accident, therefore, its warning signals were only

²⁸ And fulfilled all FAA certification requirements.

²⁹ And differently from the later tests with an old but properly-bonded fuselage, which had not been flying routes in Hawaii.

³⁰ Newton’s mechanics, for example, would have served engineers just as well as Einstein’s or Bohr’s until at least the mid-Twentieth Century, and, even today, very few technologies would reveal the error.

visible in retrospect.

The accident allowed engineers to revisit their understanding of fatigue and its specific relationship to the 737 fuselage, so that the conditions that caused the accident might be identified in the future. This is not a process that could have worked in reverse, as we saw above, but it is significant nevertheless. Since *Aloha*, the aviation engineers have modified their understanding of metal fatigue, fail-safe tear panels, fuselage life and much else, and because of this there has never been another incident like it.³¹ But if *Aloha* 243 had disappeared over an ocean, cause undiscovered, then other aircraft would have surely have failed in the same way. When metal fatigue felled a de-Havilland Comet in 1954, for instance, the aircraft was lost in the water and investigators were unable to deduce that the airplane's square windows were concentrating stress at the corners and inducing fractures. It took another disaster from the same cause, just four months later, before flights were halted, the wreckage recovered, and the causes determined (Faith 1996: 158-165). Aircraft have had oval windows ever since.

In essence, therefore, regulatory assessments of *Aloha*'s fuselage embodied a complex and detailed understanding of metal fatigue; this understanding was based on theories; those theories were built on (and reified by) tests; and those tests were inescapably theory-laden. These properties of *Aloha*'s accident and investigation have broad implications that need to be theorized in a new way. Indeed, *Aloha* 243, I suggest, both requires and exemplifies a new sociological category of disaster: the Epistemic Accident.

6. Epistemic Accidents

Disaster Studies, as Pinch (1991: 155) puts it, needs to 'bite the bullet of technical uncertainty'. Turner recognized that engineering knowledge is based in simplifications and interpretations, (Turner 1976: 379; Weick 1998: 73), but assumed that these simplifications 'masked' warning signals. The truth is more subtle, however. It is not that simplifications 'mask' warning signals, as Turner suggests, but that -- on a deep epistemological level -- there need be nothing that makes 'warning signals' distinguishable from the messy reality of normal technological practice. In other words: that there might be nothing to mask.

If it is impossible to completely and objectively 'know' complex machines (or their behaviors), then the idea of 'failing' technologies as ontologically deviant or distinct from 'functioning' technologies is necessarily an illusion of hindsight. There is no inherent pathology to failure, and no perfect method of separating 'flawed' from 'functional' technologies. To paraphrase Bloor (1976): airplanes that fall must be understood in same way as those that soar.³²

³¹ Among other things, *Aloha* led to a shift in maintenance paradigms, away from a 'fatigue safe-life' approach, to an approach with greater emphasis on 'damage tolerance' and the physics of crack growth (NTSB 2006: 37-8).

³² It should be noted that in his discussion of complexity Perrow recognises that there are systems that engineers do not understand, and that this contributes to failures. For instance, he singles-out what he calls 'transformation processes' such as chemical and fissile reactions in rockets and power plants, respectively (Perrow 1999 [1984]: 85-86). These are known unknowns, however. NAT does not recognise that engineers do not understand the systems they do understand: that there are unknown unknowns in every technical system.

There is, of course, an ontological ‘truth’ of technological functioning, as becomes obvious in retrospect. It is simply that, epistemologically speaking, actors have no objective and unfiltered access to this truth. There can be no perspective, process, or procedure that will infallibly distinguish errors from ‘non-errors’, and so there can be no way of definitively knowing that errors exist until they manifest in a failure.

This insight, although straightforward to some scholars, has unexplored but far-reaching ramifications for disaster-studies. For if there need not be anything ontologically distinct about failure -- nothing identifiable that actors ‘miss’, ‘ignore’, or ‘normalize’ in the lead-up to an accident -- then there need not be anything for social theorists to fix. This means rethinking fundamental assumptions about failure, culpability and foresight.³³ It further implies the existence of ‘epistemic accidents’.

Put simply: ‘Epistemic Accidents’ can be defined as those accidents that occur because a technological assumption proves to be erroneous, even though there were reasonable and logical reasons to hold that assumption before (if not after) the event. Epistemic accidents are analogous to Normal Accidents in several ways, not least in that they both offer an ‘extra-social’ explanation for disaster. The two are not the same, however, as each has very different properties and implications.³⁴ Nevertheless, an economical way to explore these properties and implications is by comparison to those of NAT, as outlined above.

To wit:

i. Epistemic Accidents are unpredictable and unavoidable.

Like Normal Accidents, Epistemic Accidents are unavoidable. They circumscribe another subset of ‘unavoidable’ and ‘unforeseeable’ accidents that will inevitably elude our ‘technologies of control’ and obdurately resist the best-intentioned prescriptions of sociologists. This mutual unavoidability stems from fundamentally different causes, however. Normal Accidents are unpredictable because engineers cannot wholly predict the multiplicity of possible (yet unremarkable) interactions in complex, tightly-coupled systems. Epistemic accidents, by contrast, are unavoidable because engineers necessarily build technologies around fallible theories, judgements and assumptions.

In respect to their ‘avoidability,’ therefore, it helps to group Epistemic Accidents with Normal Accidents as two sociological models of disaster that eschew the ‘foresight-failure’ premise. In other respects, however, their different mechanisms translate into very different consequences. For example:

³³ Vaughan (1996; 1999; 2003) draws on the STS literature, but her central insight -- the ‘normalization of deviance’ -- is conventional Disaster Studies. It appeals to the sociology of experts rather than the sociology of expertise and fits squarely in the tradition of identifying warnings (‘deviances’) and explaining why they go unnoticed, rather than the STS tradition of problematizing concepts such as deviance.

³⁴ The category of Epistemic Accidents encompass some, but not all, of the disasters that are attributed to engineering ‘error’ and even some that Perrow describes as Normal Accidents (e.g.: the DC10 cargo door failures described in Perrow [1994 (1984): 137-140]).

ii. Epistemic Accidents are more likely in highly innovative systems.

If Normal Accidents are more likely when technologies are complex and tightly coupled, then Epistemic Accidents are more likely in systems that stretch the boundaries of established theory and prior experience. Which is to say, they vary with ‘innovativeness’.³⁵ An airframe panel made from a new composite material, for instance, is neither complex nor tightly-coupled in itself, but it’s newness is inherently risky. When engineers use traditional aluminum panels in airplane designs, they draw on decades of metallurgical research and service experience in a way that they cannot when working with new composite materials (Downer 2009). *Aloha* testifies to how experience reveals unanticipated properties even in well-established material. (Perrow, by contrast, cites the aviation industry’s ‘use of exotic new materials’ as a factor that directly contributes to the safety of modern aircraft (1999 [1984]: 128).)

iii. Epistemic Accidents are likely to reoccur.

Unlike Normal Accidents, which are -- almost by definition -- ‘one of a kind’ events, Epistemic Accidents are highly likely to reoccur. If ten-thousand identical reactors were run for a thousand years, it is unlikely that one would fail in the same way as Three Mile Island. But if *Aloha* 243 had not failed when it did, then it is highly likely that another airplane of the same design and similar age would have failed soon after in a similar fashion. (As did the de-Havilland Comet in 1954 (see above).)

iv. Epistemic Accidents challenge design paradigms.

Epistemic Accidents invariably reveal shortcomings in existing design paradigms, unlike Normal Accidents, which rarely challenge engineers’ understandings about the world. The fact that a freak combination of pipes, valves and warning lights could lead to the near-meltdown of TMI, did not force engineers to rethink their understanding of pipes, valves or warning lights. When a fatigue crack coalesced from nowhere to blow the roof off *Aloha* 243, however, this certainly challenged a number of engineering theories, facts and understandings.

v. Epistemic Accidents are heuristic.

The fact that Normal Accidents rarely challenge engineering paradigms and seldom

Normal Accidents	Epistemic Accidents
Unforeseeable	Unforeseeable
More likely in tightly-coupled, complex systems	More likely in highly innovative systems
Unlikely to reoccur	Likely to reoccur
Do not challenge design paradigms	Challenge design paradigms
Not Heuristic	Heuristic

³⁵ Of course, ‘innovativeness’, like ‘complexity’ and ‘coupling’, is a difficult yardstick. It is qualitative variable with few, if any, straightforward or objective ways to measure or quantify it. This should not detract from its merit, however. ‘Risk’, for instance, is a highly ambiguous variable that enjoys wide currency in public policy as an analytical construct.

reoccur gives them minimal heuristic value, but the reverse is true of Epistemic Accidents. If our ‘technologies of control’ have little to learn from the former, then they have much to learn from the latter. Epistemic Accidents force experts to confront the limits of their knowledge and adjust their theories. Because they are likely to reoccur, they allow engineers to leverage hindsight in a way that Normal Accidents do not: turning it into foresight and preventing future accidents. Maintenance crews are now wise to the implications of *Aloha*-esque fatigue.

Epiphenomena

Once the properties of Epistemic Accidents have been highlighted, it is easy to see them reflected in the networks that build and sustain complex technological systems. The correlation between Epistemic Accidents and innovativeness, for instance, is reflected in the civil aviation industry’s long-standing emphasis on conservative design, and how this contributes to its high levels of reliability. Civil airplane manufacturers are extremely conservative in their design choices. They were very slow to embrace composite materials, for instance, despite the obvious advantages of doing so, because abandoning aluminum meant abandoning decades of painfully-honed knowledge about sheet metal -- how it fatigued; how it behaved under stress -- that informed airplane design; manufacture; regulation; operation; repair; modification, and much else besides (Downer 2007; 2009a).

The knowledge that manufacturers are keen to preserve, meanwhile, owes much to the heuristic properties of Epistemic Accidents. Modern designs draw on decades of trials and errors, involving hundreds of tragic but instructive accidents, such as *Aloha*, and thousands of lost lives. Hard work and clever engineering were vital to the remarkable reliability enjoyed by civil aircraft today, but, with all the epistemic ambiguities of technological practice, it would have been impossible to anticipate every design flaw in the laboratory or on the drawing board. Understanding aircraft, as the FAA director testified to Congress in 1980: ‘...is not a year’s process, but a fifty-, sixty- or seventy-year process.’ This is why successful high-technologies invariably have long and inglorious legacies of failure behind them, and modern aviation, like liberal democracy, owes much to its tombstones.

In turn, this reliance on past failures -- derided and misrepresented by many commentators (Weir 2000; Nader & Smith 1994) -- has broad ramifications for technology regulation and policy. It suggests, for instance, a new mechanism by which the designs of safety-critical technologies become ‘locked-in’ (David 1985). If complex systems are built on the lessons of past failures, and those lessons become less relevant as systems change, then this will make systems difficult to displace, once they are established, without introducing new epistemic uncertainties and corresponding risks.

Discussion

Not everyone will agree that *Aloha* 243 was an Epistemic Accident. The details of *Aloha*, like those of all large technological accidents, are intricate enough to support a variety of

interpretations, and knowledgeable readers might prefer a more conventional explanation than the one offered above. Rather than explore every facet of the incident and challenge every rival explanation, therefore, let us note that its purpose in this narrative is illustrative more than evidential.

Simply put: the central thesis of this paper should stand on its own logic with or without *Aloha*. This is because, if we accept the premise that some accidents result from erroneous beliefs, and we further accept the argument that even the most ‘perfect’ knowledge-claims necessarily contain uncertainties, then Epistemic Accidents must exist. To demonstrate that *Aloha* was not an Epistemic Accident, even if it were possible, would not be not to refute the idea of Epistemic Accidents.

By way of framing any future discussion, however, it might help to briefly anticipate some further caveats and qualifications.

Caveats and qualifications

- *Not all accidents are ‘Epistemic’.* To say that epistemic accidents exist is not to say that all, or even many, accidents fit this description. As with Normal Accidents, Epistemic Accidents are but a subset of unavoidable accidents in a much larger pool of potentially avoidable failures, (many, perhaps most, of which result from organizational errors).
- *Other approaches to failure remain valuable and viable.* This is a corollary of the point above. To accept that Epistemic Accidents exist and that they are unavoidable is not to deny that some organizational structures are more effective than others, or that many accidents are preventable. It is vital, therefore, that social scientists continue to explore the organizational foundations of error.
- *Not all knowledge-claims are equal.* The idea that even the most rigorous scientific and technical knowledge contains fundamental ambiguities (such that ‘errors’ are impossible to eliminate entirely) does not imply that there are better (or even equal) alternatives to that knowledge. All facts are, to some degree, ‘under-determined’, as sociologists of knowledge say, but this is not to suggest that all knowledge-claims are equal.³⁶

The Politics of Blame

Epistemology aside, there is one social implication of Epistemic Accidents liable to critique: their relationship to the politics of blame. Technological accidents (indeed man-made disasters of all kinds) invariably demand a ‘civic theatre,’ or ‘ritual politics’, of blame (Vaughan 2005: 63-4; Jasanoff 2005). Yet, the idea that failures can occur without errors, and that even ‘perfect’ systems will occasionally fail, clearly has implications in this context because it suggests that there are essentially ‘blameless’ accidents. It is easy, therefore, to imagine actors invoking Epistemic Accidents to deny culpability.

³⁶ For more a detailed discussion of this point see Bloor (1976) and Collins & Pinch (1993)

It is important, in this context, to reiterate that the finitist view of failure does not, in any way, preclude the idea that some accidents involve failings: i.e. culpable errors. Nevertheless, it follows logically that some accidents are, indeed, less blameworthy than they appear. It will be very difficult to say which, especially as all Epistemic Accidents inevitably look anticipatable in hindsight. This paper can offer no straightforward rubrics about this, but it is worth remembering two things. Firstly, that in some spheres we are already comfortable distancing error from blame: we do not investigate, penalize or reorganize scientists, for instance, when ideas about cancer are overturned. And, secondly, that Normal Accident Theory poses the very same dilemmas, and yet the politics of blame continues apace.

The Limits of Control

Assuming the ‘finitist’ view of failure holds, therefore, and that Epistemic Accidents exist, then how does the concept relate to other discussions of risk, regulation and governance in complex systems?

Although a few implications of Epistemic Accidents and their relationship to technological practice have been outlined above, the detailed socio-political implications of a finitist understanding of technological failure are far too intricate to explore here in detail. Speaking very broadly, however, we might say that it challenges conventional ideas about civic society’s relationship with technology and with failure more generally.

Technology policy discourse favors an idealized model of regulation as a mechanical process governed by formal rules and objective algorithms: what Gherardi and Nicolini (2000: 343) call the ‘bureaucratic vision of safety’. NAT has long been a crack in this vision because it suggests that even ‘perfectly’ organized technologies will sometimes fail. This paper broadens and deepens that argument by suggesting another mechanism of ‘inevitability,’ and another category of accidents that elude even the best organizational solutions. For different reasons, both suggest that catastrophic potential is woven into the very fabric of technological progress. Technologies sometimes fail, and this is a levy of modernity.

It is worth noting that the argument here potentially reaches far beyond technology.³⁷ The argument that error might not be ontologically distinct or epistemologically identifiable has implications for the social analysis of all kinds of system failures, whether they be technological, medical or financial. Sociological accounts of everything from medical mishaps to market collapses, invariably construe them as ‘deviances’ from some ideal, which need explaining. The idea that, epistemologically speaking, there need not always have been a ‘knowably correct’ choice, offers an important counter-perspective.

Simply put, the designs of all complex systems rely on elaborate knowledge claims. These contain unavoidable uncertainties and ambiguities, which inevitably give rise to fundamentally unresolvable doubts and disagreements. Accounts of accidents that directly or indirectly explain them by reference to unresolved disagreements or ignored doubts, therefore, fundamentally misconstrue the nature of engineering knowledge. To working with real-life systems is, by necessity, to make judgments with imperfect information. To say that we should not fly airplanes

³⁷ Much as NAT potentially does.

unless all disagreements have been resolved and all facts are known is, in essence, to say that we should never fly airplanes.

References

Air Safety Week (2005) 'Hercules Crash In Baghdad Points To Metal Fatigue In C130's Wing Center'. *Air Safety Week*. Feb 21.

Aubury, Martin (1992) 'Lessons from Aloha'. *BASI Journal*, June. Available at: http://www.iasa.com.au/folders/Safety_Issues/others/lessonsfromaloha.html

Bijker, W.; Hughes, T. & Pinch, T. (eds) (1987) *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*. MIT Press: Cambridge (MA).

Bloor, David (1976) *Knowledge and Social Imagery*. Routledge & Kegan Paul: London/Henley/Boston.

Bundesstelle für Fluguntersuchung (BFU) (2004) 'Investigation Report AX001-1-2/02 MAY 2004'.

Collingridge, D. & Reeve, C. (1986) *Science Speaks to Power*. St. Martin's: New York.

Collins, Harry (1985) *Changing Order*. SAGE: London.

Collins, Harry (1992) *Changing Order: Replication and Induction in Scientific Practice*. University of Chicago Press: Chicago.

Collins, Harry & Pinch, Trevor (1993) *The Golem: What Everyone Should Know About Science*. Cambridge University Press: Cambridge.

Collins, Harry. & Pinch, Trevor (1998) *The Golem at Large: What You Should Know About Technology*. Cambridge University Press: Cambridge.

Cushman, John (1989) 'U.S. Investigators Fault Aloha Line In Fatal Accident' in *The New York Times*, May 24.

Davies, R. & Birtles, Philip (1999) *Comet: The World's First Jet Airliner*. Virginia: Paladwr Press.

David, Paul A. (1985) 'Clio and the Economics of QWERTY'. *American Economic Review Papers and Proceedings* 75: 332-337.

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

Downer, John (2009a) 'Watching the Watchmaker: On Regulating the Social in Lieu of the Technical.' LSE CARR Discussion Paper 54; June 2009; [ISBN 978-0-85328-396-6] Available online: <http://www.lse.ac.uk/collections/CARR/pdf/DPs/Disspaper54.pdf>

Downer, John (2009b) 'When Failure is an Option: Redundancy, Reliability, and Risk.' LSE CARR Discussion Paper 53; May 2009; [ISBN 978-0-85328-395-9] Available online: <http://www.lse.ac.uk/collections/CARR/pdf/DPs/Disspaper53.pdf>

Faith, Nicholas (1996) *Black Box*. London: Boxtree.

Galison, P. (2000) 'An Accident of History' in Galison, P. & Roland, A. (eds.) (2000) *Atmospheric Flight in the Twentieth Century*. Kluwer: Boston, pp. 3-43.

Garrison, Peter (2005) 'When Airplanes Feel Fatigued' in *Flying*, September. [available at http://www.flyingmag.com/article.asp?section_id=13&article_id=578&print_page=y]

Gephart, R.P., Jr., (1984) 'Making Sense of Organizationally Based Environmental Disasters'. *Journal of Management* 10 (2): 205-225.

Hackett, Edward; Amsterdamska, Olga; Lynch, Michael and Wajcman, Judy (eds) (2008) *The Handbook of Science and Technology Studies* (3rd edition). MIT Press: Cambridge (MA).

Hall, R. F. & Powell, Brian E. (2000) 'The Effects of LCF Loadings on HCF Crack Growth' M.O.D. Report number: A405493, October.

Hutter, Bridget & Power, Michael (eds) (2005) *Organizational Encounters With Risk*. Cambridge: Cambridge University Press.

Hopkins, Andrew (1999) 'The Limits of Normal Accident Theory'. *Safety Science* 32 (2): 93-102.

Hopkins, Andrew (2001) 'Was Three Mile Island a 'Normal Accident'?' *Journal of Contingencies and Crisis Management* 9 (2): 65-72.

Jasanoff, Shelia (1986) *Risk Management and Political Culture*. Russell Sage Foundation: New York.

Jasanoff, Sheila (2005) 'Restoring reason: causal narratives and political culture' in Hutter, Bridget & Power, Michael (eds.) *Organizational Encounters With Risk*. Cambridge: Cambridge University Press, pp. 207-232.

Kripke, Saul (1980) *Naming and Necessity*. Harvard University Press: Cambridge (MA).

Kuhn, Thomas (1962) *The Structure of Scientific Revolutions*. The University of Chicago Press: Chicago.

Langewiesche, William (1998) *Inside the Sky: A Meditation on Flight*. Pantheon: New York.

La Porte, T. (1982) 'On the design and management of nearly error-free organizational control systems', in Sills, D. Shelanski, V. & Wolf, C. (eds.) *Accident at Three Mile Island*. Boulder: Westview.

La Porte, T., & Rochlin, G. (1994). 'A rejoinder to Perrow', in *Journal of Contingencies and Crisis Management* 2 (4): 221-227.

La Porte, T. (ed) (1991) 'Social Responses to Large Technical Systems: Control or Anticipation' in *Proceedings of the NATO Advanced Research Workshop on Social Responses to Large Technical Systems: Regulation, Management, or Anticipation*, Berkeley, California, October 17-21, 1989: Nato Science Series.

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

Latour, Bruno (1996) *Aramis, or The Love of Technology*. Harvard University Press: Cambridge (MA).

MacKenzie, Donald (1990) *Inventing Accuracy: A Historical Sociology of Nuclear Weapon Guidance*. MIT Press: Cambridge (Mass).

MacKenzie, Donald (1996) *Knowing Machines: essays on technical change*. MIT press: Cambridge (Mass).

MacKenzie, Donald (1998) 'The Certainty Trough.' in Williams, Robin; Faulkner, Wendy & Fleck, James (eds) *Exploring Expertise: Issues and Perspectives*. Basingstoke, UK: Macmillan, pp. 325-29.

Nader, Ralph, & Smith, Wesley (1994) *Collision Course: The Truth About Airline Safety*. TAB Books: New York.

National Transportation Safety Board (NTSB) (1989) Bureau of Accident Investigation: 'Aircraft Accident Report: Aloha Airlines, Flight 243, Boeing 737-200, N73711, near Maui, Hawaii, April 28, 1988.' Report no: NTSB/AAR-89/03; Acc. No: PB89-91C404; June 14. Washington DC.

National Transportation Safety Board (NTSB) (2006) 'Safety Report on the Treatment of Safety-Critical Systems in Transport Airplanes' Safety Report NTSB/SR-06/02. PB2006-917003. Notation 7752A. Washington DC.

Perrow, Charles (1983) 'The Organizational Context of Human Factors Engineering'. *Administrative Science Quarterly* 28 (4): 521-541.

Perrow, Charles (1984) *Normal Accidents: Living with High-Risk Technologies*. Basic Books inc: New York.

- Perrow, Charles (1999) *Normal Accidents: Living with High-Risk Technologies* (2nd ed.). Princeton University Press: New Haven.
- Perrow, Charles (1994) 'The limits of safety: the enhancement of a theory of accidents' in *Journal of Contingencies and Crisis Management* 4 (2): 212-220.
- Petroski, Henry (1992) *To Engineer is Human: The Role of Failure in Successful Design*. Vintage books: New York.
- Petroski, Henry (1994) *Design Paradigms: Case Histories of Error and Judgment in Engineering*. Cambridge University Press: Cambridge.
- Pinch, Trevor & Bijker, W. (1984) 'The Social Construction of Facts and Artifacts: or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other'. *Social Studies of Science*. 14: 339-441.
- Pinch, Trevor (1991) 'How do we treat technical uncertainty in systems failure? The case of the space shuttle Challenger' in La Porte, T. (ed) *Social Responses to Large Technical Systems*. Nato Science Series: D: 58; March: 143-158.
- Pinch, Trevor (1993) '"Testing - One, Two, Three...Testing!": Toward a Sociology of Testing'. *Science, Technology, & Human Values* 18 (1): 25-41.
- Rasmussen, J. (1990) 'Human Error and the Problem of Causality in the Analysis of Accidents' in *Philosophical Transactions of the Royal Society of London* B327: 449-462.
- Rasmussen, J., Duncan, K. & Leplat, J. (1987) *New Technology and Human Error*. John Wiley and Sons: New York.
- Reason, James (1990) *Human Error*. Cambridge University Press: Cambridge.
- Sagan, Scott (1993) *The Limits of Safety*. Princeton: Princeton University Press.
- Scott, James (1998) *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. New Haven: Yale University Press.
- Shrivastava, P. (1987) *Bhopal*. Cambridge: Ballinger.
- Stoller, Gary (2001) 'Engineer has alternate theory on plane disaster' in *USA Today*. April. Available at http://www.iasa.com.au/folders/Safety_Issues/RiskManagement/alohaagain.html
- Turner, B.A. (1976) 'The Organizational and Interorganizational Development of Disasters' in *Administrative Science Quarterly* 21 (3): 378-397.
- Turner, B.A. (1978) *Man-Made Disasters*. Wykeham: London.

Turner, B.A. & Pidgeon, N.F. (1997) *Man-Made Disasters* (second edition). Butterworth-Heinemann: Oxford.

Turner, B. (1994) 'Causes of Disaster: Sloppy Management.' *British Journal of Management* 5 (3): 215-219.

Vaughan, Diane (1996) *The Challenger Launch Decision*. Chicago: University of Chicago Press.

Vaughan, Diane (1999) 'The Dark Side Of Organizations: Mistake, Misconduct, And Disaster' in *Annual Review of Sociology* 25: 271-305.

Vaughan, Diane (2003) 'History As Cause: Challenger and Columbia' Ch. 8, Report, *Columbia Accident Investigation Board*. Vol. 1. August.

Vaughan, Diane (2005) 'Organizational rituals of risk and error' in Hutter, Bridget and Power, Michael (eds) *Organizational Encounters With Risk*. Cambridge: Cambridge University Press, pp 33-66.

Weingart, Peter (1991) 'Large technical systems, real-life experiments, and the legitimation trap of technology assessment: the contribution of science and technology studies to constituting risk perception' in La Porte, T. (ed.) *Social Responses to Large Technical Systems*. Nato Science Series.

Weick, K.E. (1987). 'Organizational culture as a source of high reliability.' in *California Management Review* 29: 112-127.

Weick, Karl (1998) 'Foresights of Failure: An Appreciation of Barry Turner' in *Journal of Contingencies and Crisis Management* 6 (2): 72-75.

Weir, Andrew (2000) *The Tombstone Imperative: The Truth About Air Safety*. Simon & Schuster London.

Wynne, Brian (1988) 'Unruly technology: Practical Rules, Impractical Discourses and Public Understanding'. *Social Studies of Science* 18: 147-67.