Part I

Introduction: The Social Background

Chapter 1

An Example of Everyday Technical Risk Analysis

We recently discussed the danger of using cell phones on gas station forecourts on a mailing list of professionals interested in safety-critical systems involving computers, to which I belong. Participants in the discussions include some of the world's best-known computer-related safety researchers, as well as some of those who participated in writing the new standard IEC 61508 for the development of safety-critical systems with "programmable electronic" components, regulators from organisations such as Britain's Health and Safety Executive (HSE) and colleagues from industries which have a significant interest, such as transportation and power.

1.1 Engineering Risk Analysis

Much of engineering is concerned with the production and analysis of artifacts, one of the aspects of which is an assessment of the safety of the artifacts in use. Because artifacts can be complex and large (think of a industrial chemical plant, or commercial aircraft), they are often called "systems". The word is typically used for artifacts in which complex behavior plays a major role.

A safety analysis of a system or a situation including systems usually proceeds along roughly the following lines.

We start by fixing the interpretation of the concepts "accident", "(accident) severity", "hazard" and "(hazard) latency". First, you have to say what you consider to be an accident There are many definitions, for example [vdM00, Entry for "Accident"], but basically it is up to you. People usually think in terms of significant damage to themselves, their surroundings, or their bank accounts, but it is probably better to be specific. You also have to classify accidents according

to their "severity".[1] Is it a bruise or a break? Will it cost a dollar or a million dollars? Then, you have to identify "hazards". These are situations, whether states of the system alone, of the environment alone, or of a combination, which are precursors of an accident (cf. [vdM00, Entry for "Hazard"]). A hazard doesn't actually have to result in an accident; but the idea is that if you deal with the hazards effectively, the accidents don't happen. Hazards persist for a length of time (called "latency", [vdM00] cites [Lev95, Chapter 9]).

At this point, you have the opportunity to deal with the hazards, either by removal, or mitigation, or even by reducing the severity of a resulting accident. This is not strictly an "analysis" step, but it does seem prudent to fix things as you go along.

After you have dealt with all the hazards, you have to figure out what the residual "risk" is. To do this, you try to determine the following:

- for each hazard, the likelihood that it will occur
- for each accident that can result from a given hazard, the likelihood, assuming the hazard has occurred, that the accident will ensue (the "conditional probability of the accident on the hazard", for those who know some probability theory)

Then

- for each hazard/accident combination, you "combine" (probability theorists, read "multiply") the likelihood of the hazard with the conditional likelihood of the accident given the hazard, and weight this by the severity of the accident (which has usually been given some numerical value, so "weight" means "multiply")
- You then "combine" (read "add") all these individual results into a general result, called the "risk".

(Probability theorists can note that, if you take the "severity" to be a number, and consider this number to be a measure of loss, then I have just described a calculation of the *expected value of loss*. That is not a coincidence. Further, one should note that my description of the process is only exact if there is a degree of probabilistic independence of the situations whose likelihoods are combined. If not, you have to tweak the numbers appropriately to achieve an accurate assessment of the expected loss. I discuss this in Chapter 5.

Now let's go to the discussion of phones on forecourts.

[1] op. cit., entry for "Severity".

1.2 Phones on Forecourts: Causal Analysis

Granted by all discussants, a gas station forecourt can contain flammable vapors (a suitable fuel-air mix). The question is how and if mobile phones could be an ignition source. (Fire safety experts generally agree that three things are needed for a fire: a combustible substance, a supply of oxygen, and an ignition source.)

Nobody could say exactly why the *use* of phones per se might be dangerous. Yes, they are electromagnetic (EM) radiators, and the use of EM radiators in potentially explosive atmostpheres is usually regulated. But exactly what hazards do mobile phones engender when in use?

John Ganter of Sandia Labs in the US pointed out that tales of explosions in gas stations being caused by use of mobile phones appeared to be an urban myth:

> The explosion story was thoroughly debunked on National Public Radio in November 1999. All versions traced back to an apocryphal story from Southeast Asia. They reported that oil companies, however, are still pressing ahead with warning sticker campaigns.
>
> `http://www.darwinawards.com/legends/legends1999-04.html`
>
> Exxon spokesman Crawford Bunkley said the company does not know of any fires caused by cell phones. But, "although the likelihood of ignition remains remote, nonetheless, we believe this warning is appropriate in light of statements make by some phone manufacturers, who have cautioned their customers to switch off their phones when refueling."
>
> `http://www.austin360.com/entertainment/features/legend_cellphone.html`

Simon Brown of the UK Health and Safety Executive (HSE) suggested that the possibility of dropping the phone presents a hazard. Why? Because the battery could come loose and produce a spark. This concern was echoed by an HSE colleague to whom Simon referred me.

A member of a major UK safety consultancy asked if there are "published data of the electromagnetic field strengths required for ignition of various types of fuels or fuel/air mixtures". Simon Brown referred him to British Standard (BS) 6656:1991 'Guide to the prevention of inadvertent ignition of flammable atmospheres by radio-frequency radiation'.

Of course, it is not the spark alone that is crucial to ignition, but the energy it contains. Nobody seems to be worried about people taking off sweaters in garage forecourts, despite that, if the sweater is knitted from artificial fiber, tens or hundreds of sparks could easily be produced. The field strength required

to produce a spark in air (to generate the "avalanche effect" as it is known to electrical engineers and physicists) depends for all practical purposes only on air pressure, and also the distance between conductors if that distance is under about 1cm.

American Petroleum Institute guidelines for fuel ignition (at least for aircraft fuel) say 0.25 milliJoules of energy is required for ignition [Swa00]; the "generally accepted" figure according to an aviation fuel-ignition expert employed by NASA to report on sparks for TWA 800 is 0.2 mJ [Fis00].

So then I asked a question: "Spark or whatever, how do you suddenly get 0.2 mJ out of a battery classified as 1000mAh at 3V?"

The answer, from John Dalton, a consultant to Reflex Technology UK Ltd, is that the contacts could be shorted. His calculation was "[...] 2V x 10A x .00001s = 0.2mJ. [A typical mobile phone] battery stores over 10KJ [...]". The answer, that an "incendive spark" could theoretically be generated from a mobile phone battery, was confirmed by 1992 HSE investigations (according to Simon Brown's colleague) as well as by my colleague Dr. Willi Schepper, an electrical engineer and physicist who is expert in EM fields in enclosed spaces. Willi also pointed out that the same was true of any of the everyday batteries used in, say, flashlights, cameras and Walkmans. My colleague Professor Harold Thimbleby (then) of University College, London, wondered nevertheless whether one could contrive to obtain such a spark, and volunteered to spend the weekend experimenting with batteries and propane (he is, fortunately, still with us :-).

Well, maybe an incendive spark is not excluded by simple energy calculation, but the question remains, if and how one could "short" the battery contacts inadvertently, to create an incendive spark, by dropping a mobile phone.

Physically, I find it hard to see how the terminals can be shorted inadvertently in any of the designs of phone that I looked at. In two of them, my Nokia 9110 Communicator and my Ericsson SH888, the battery is external, part of the body. In two others, a Siemens S35 and a Motorola L7039 Timeport, the battery is internal, covered by a detachable part of the body shell. In all four phones, the battery contacts are four flush metal strips, indented from the battery body, of size approximately 3mm x 2mm. In the 9110, SH888 and S35, these are indented some 1mm from the battery body (in the S35, each individual contact is separately indented) and protected further by an extruded rand on the body, for a total of some 3mm or more. In the 9110, the "stepped" battery body shape in this area ensures that the contacts are some half-centimeter inside the "convex hull" of the body.

The contacts on the battery of the L7039 are only protected by the thickness of a layer of thick film that encloses the battery. However, the shell covering the battery is protected by a double-action mechanism: first a knob, flush with

the body, must be depressed some 5mm; simultaneously the shell must be slid some 5mm along the body, at right angles to the direction in which the knob is depressed. The entire body shell seems to be robust. It is hard to believe that it could open, even when thrown against the floor. (Simon Brown informs me, though, that he has dropped his L7039 and the battery has fallen out.)

Further, in all these phones, pins internal to the phone make contact with the battery. These pins are spring-loaded; I presume that any movement of the pins, conformant with release of the battery, acts as a switch to turn the phone off before contact is lost, amongst other things to protect the phone electronics from surges such as may happen with spark events. There is obvious motivation for phone designers to enhance the reliability of their products in this manner. First, they wish their products to be known to be reliable, even under abuse; further, they are statutorily responsible in, say, Europe, for replacement of defective devices up to a year after sale, and it can be very hard to prove that a phone has been physically abused even if they think it may have been.

There are good reasons for preventing arc faults in PEDs such as mobile phones, besides the motivation of increasing reliability. Arcs are potentially powerful sources of electromagnetic noise. There are many EM-sensitive environments, some of them, such as commercial aircraft, with safety-critical electronic components which may be influenced by fields inside the cabin. The UK Civil Aviation Authority has recently confirmed by experiment that a radiating mobile phone inside a cabin can generate field intensities inside an avionics bay or flight deck of between about 2 Volts/meter and 5 Volts/meter, which exceeds the demonstrated interference immunity levels of avionics qualified to pre-1984 standards, some of which is also installed in newly-built aircraft [Aut00]. There are also many anecdotes collected by aviation authorities (such as NASA's ASRS or the UK CHIRPS systems) from flight crew about in-flight interference apparently related to passenger PEDs, as often tested in flight using Mill's Method of Differences (cabin crew asks the passenger to turn the device off, you observe the effects go away; and then back on, and the effects return).

Back to phones and forecourts. Yes, there may be flammable fuel-air mixture at unpredictable places on a station forecourt. Yes, mobile phone batteries, similarly to other household electrical devices, contain enough "juice" in their batteries to produce a 0.2mJ or greater discharge. Yes, such a discharge can ignite a flammable fuel-air mixture. But no one has identified an actual mechanism to produce a discharge. A simple switch design ensures that, when a battery moves in relation to the phone body, the electronics are isolated. And the battery terminal structure cannot easily be distorted. "Short circuits" between battery terminals through accidental contact with conductors appear to be effectively hindered by the physical design of the batteries, as are thereby sparks engendered through such activities.

I think that is the status of the risk analysis. Let us consider the accident to be avoided as simply the ignition of flammable vapors, whether or not this event causes injuries. One may identify two nominal hazards here: use of a phone, and dropping a phone. But a causal mechanism that could lead from either of these hazards to an accident has not been identified. A putative "hazard" that doesn't appear to be able to lead to an accident is not a hazard. So it seems likely that use of a phone or dropping a phone simply aren't hazards.

End of story? Well, no, there is the matter of public policy.

1.3 Phones on Forecourts: Safety Policy

Some gas stations explicitly prohibit use of phones on forecourts. Simon Brown's HSE colleague suggested that this might be a policy of station owners. An HSE Guide [HE97] says that employees must "make sure that ..... no one uses portable electric/electronic equipment such as a CB radio or portable telephone". The force of such "guidance" comes from the *1974 UK Health and Safety at Work etc Act*, which establishes inter alia the responsibility of an employer towards employees to take all reasonable steps to ensure employees' safety and to equip them to do their job without danger to themselves or others (paraphrasing op. cit.).

The document does not say where an employee should attempt to ensure that no one uses RF transmitting devices. Without further qualification, it would be reasonable to interpret it as meaning anywhere on the premises. It turns out it could mean either that, or in an area around filling pumps and nozzles. I base this on the following observations.

Simon Brown's colleague suggested that one is only permitted to use RF transmitters in environments defined to be at risk of explosive conflagration if the RF device has been individually certified for such use. For example, UK Home Office Guidance HGN(F)15 issued from the Communications Advisory Panel advises that no transmitter shall be used within 10 metres of an area "zoned" as "potentially explosive", which, according to Brown, includes the areas around filling pumps and nozzles. That gives us one interpretation. The latest UK Institute of Petroleum Guidance for the construction and operation of petrol filling stations requires an assessment of any transmitting equipment to be made by the site operator before permission for use on site is granted. That gives us the second interpretation. I am informed further that, in practice, this means assurances from the equipment supplier that the equipment is certified for operation in a zoned area.

There is a third statutory interpretation, namely the forecourts of gas stations. This occurs in the Highway Code, the UK "guidance" document to road users, whose violation is used as a basis for police warnings and prosecutions. Mark Coates of BAe Systems quoted the Highway Code:

> Petrol stations
>
> Never smoke or use a mobile phone on the forecourt of petrol stations as this is a major fire risk and could cause an explosion.

Notably, it doesn't say "No CB radios". Maybe there just aren't enough of them to justify the extra words.

Compare all this advice about use of phones with John Ganter's quote from Exxon spokesman Bunkley that

> "some phone manufacturers [...] have cautioned their customers to switch off their phones when refueling."

If there is a risk from RF transmissions from phones, this risk also occurs when the phones are switched on, whether or not they are being used for a conversation. Mobile phones which are switched on, in so-called "standby" mode, transmit protocol signals to maintain contact with a base station. If transmission were to be a problem, then phones should be switched completely off in the area of concern.

I tried to see if anyone could tell me of a difference relevant to a risk assessment between dropping a phone when in use and dropping a phone which is switched on but not in use. That is, whether the likelihood of an "incendive spark" differed between the two cases, and thus whether these should be considered two distinct hazards.

I guess that a phone in use is drawing more current; so the resistance of the internal electronics must be lower; so the current contained in any spark between battery terminals and electronic components of the phone when it is being used may contain more energy. But we have already ruled out the possibility of such a spark through considering the design of the phone.

We need to consider possible arcs occurring through current passing between different battery contacts directly, not by current passing between contacts and pins. Let us assume the mechanics and situation of accidentally dropping a phone are relevantly similar whether a call is being made or the phone is on standby. Then any relevant difference in the risk calculation could only be put down to differing likelihoods of dropping a phone in use, when you are holding on to it relatively firmly, and dropping a switched-on phone when not in use, say when rummaging around in a handbag or in your pockets for money or car keys. Well, no prizes for guessing whether anyone has estimated *those* likelihoods.

There appears to be no reason anyone knows why a situation in which a phone is in use on a forecourt is more risky than a situation in which a phone is present but on standby. This seems to be acknowledged by unidentified phone manufacturers in the quote from Bunkley; they want you to turn the phone off.

But the prohibition against use, but not against standby, through the HSE or the UK Highway Code guidance does not match this concern. Furthermore, let us recall that no one appears to know whether these situations are indeed hazardous according to the technical definition.

Simon Brown's HSE colleague wrote to me that, in 1992, the construction of certain mobile telephones was examined, and compared against "recognised standards for EX type equipment" (which I take to mean standards for RF equipment qualified for use in certain explosive atmospheric zones) and drew the conclusions that mobile telephones did not meet any recognised standard for intrinsically safe electrical equipment for use in a "zone 1" area. The reasons given were: excessive battery output, inadequate segregation between critical components, and unacceptable voltage limits on the supply. When assessed against British Standard (BS) 6941 for use in "zone 2" areas, the levels of current in normal operation were shown to be safe but the protection afforded by the plastic enclosure and battery connections did not meet the requirements.

I take it there have been considerable changes in phone design since then. However, the expert informed me that the consensus view of his colleagues working in the area was that they would probably arrive at the same conclusions for equipment available today. (To emphasise: that's a view, not a finding).

That explains why the ban on use; it is radiative RF equipment in a high-fire-risk area, and such use by employees is controlled statutorily through HSE guidance. Employees of gas stations are also supposed to stop you from using your phone on the premises. This does raise the question why this ban is not clearly advertised, and its semi-statutory nature made clear, on every station forecourt in the UK.

Companies must conform to this guidance, or risk being prosecuted by HSE if an accident ensues. However, it is not clear whether or how ordinary citizens are required to follow it; it does not have the force of law.

So the following appears to be the situation in the UK. You may use a phone on gas station premises if you are not an employee of the gas station, but you must expect an employee to attempt to get you to stop, presumably by informing you that your phone is not known to meet acceptable emissions standards for areas around filling pumps and nozzles. You may ignore himher with impunity. However, if you use your phone on the forecourt, you may expect a policeman to warn you about a offence associated with not following the Highway Code. If you should care to annoy all these people, you may terminate your phone conversation, and start practicing your juggling with your, your partner's and your children's phones, all on standby, content in the knowledge that the policeman can do nothing about it, the employee can ignore you, and that you are undertaking no risk.

If that all seems a little daft to you, consider that common risk assessments as performed by HSE are based on intuitively plausible principles, but not necessarily on technical risk analyses such as explained at the beginning of this paper. Let us then consider some common ways of formulating policy in the absence of knowledge.

1.4 Some Principles

Here is the "common sense" analysis. Phones have batteries: batteries contain enough "juice" so that, if discharged all at once in an arc, that arc could potentially contain sufficient energy to ignite a fuel-air mixture. Such an arc could be allowed by detachment of the battery and exposure of the contacts to a bridging conductor during abuse of a phone. Principle: better to be safe than sorry (BBSTS). (This is often called the "precautionary principle", which would lead to a different acronym used by my small son for something else entirely.) Based on this reasoning, BBSTS would seem to indicate that phones be switched off (SO) on station forecourts.

Further, there is employee "guidance" for use of radio equipment in environments potentially at risk for fire. So another, well-established policy principle of Prior Coverage (PC) applies. PC: phones are radio equipment, gas station forecourts are described and covered environments; this situation is covered by existing guidance for companies which says "don't use" (DU).

Based on similar considerations, PC suggests "don't use", but BBSTS suggests "switch off". DU and SO are logically related. A phone which is switched off cannot be used, so any situation in which SO is followed is also a situation in which DU is necessarily followed. We can phrase this by saying that SO is directly comparable to, and stronger than, DU.[2]

Another policy principle, which we may call Like-Is-Like (LIL) would say: advise the public the same way that we advise companies on employee ("servant" in UK law) or subcontractor ("agent") behavior. PC + LIL indicate the weaker DU, but not the stronger SO, for public behavior with mobile phones on forecourts.

The reasoning on which SO under BBSTS is based and the reasoning on which DU under PC+LIL is based are not indentical. BBSTS is based on the battery-spark causal mechanism; PC is based on potential RF transmission dangers.

What has actually happened is that the measure implied by both, namely DU, has been chosen to advise the public. Such a procedure could be justified by a principle such as Multiple Justification (MJ): if different, partially independent, considerations justify some requirements, then recommend a requirement justified

[2] the state predicate expressing SO logically implies the state predicate expressing DU.

by all these considerations. A use of MJ was indirectly suggested in communications with HSE representatives. It seems to be felt that a requirement following from multiple independent considerations is somehow better justified that one following from just a single consideration.

All these principles, BBSTS, PC, LIL, MJ, seem intuitively reasonable. But, as is so often the case with such social principles, if unrestricted, they lead to anomalies. For example, a common criticism of BBSTS is that its uniform application, if unrestricted, can lead to social paralysis or to "infringement of basic freedoms" or various other phenomena regarded as undesirable by a society. The principle LIL, if applied indiscriminately, yields the "corporate state" so disliked by many cititens of Western countries from John Kenneth Galbraith onwards. PC likewise is a principle of strong conservatism, also disliked by many Western citizens, though not by all. So any use of these principles has to be carefully circumscribed: indiscriminate use is not seen to be justified by the polity. Although we (and organs such as HSE) must take these considerations seriously, we could reasonably leave the circumscription of these principles to social philosophers of safety and risk, e.g., [Bec86, Luh91, SF91]. I think it is true to say, however, that a policy reliance on BBSTS, PC, LIL and MJ is self-consistent, providing one gets the circumscription right. This is the combination of principles which I identified at work in my correspondence with HSE.

I wish to introduce another intuitively justifiable principle, Physical Causal Justification (PCJ). PCJ basically says that, if you suggest a safety measure R for a situation S, then that's because there are circumstances in which violating R in S leads to an accident. A little more detailed formulation of PCJ says that, for each safety recommendation R for circumstances S, there must be a physical process in circumstances S which results in an accident in some situation consistent with S if all recommendations except R are followed but R itself is violated. PCJ says, roughly, don't attempt to restrict on the basis of concern about a process that cannot physically happen. It is a principle of realism about causal mechanisms.

PCJ is the policy principle associated with technical risk analyses. If there is no causal mechanism leading from situation S to an accident, then no matter what are counted as "hazards" in the technical analysis, the likelihood of the hazard leading to an accident is null, because no accident can follow from the situation S. A very intuitive principle, which I shall call Don't Fret (DF), states that if the risk is null, no measures need be taken. PCJ describes the situation which results from performing technical risk analyses and applying DF. I think DF is incontrovertible.

technical risk analyses of the form mentioned at the start of this essay are required to be performed for safety-critical artifacts such as cars, power stations, commercial aircraft and air traffic control systems, and the chemical process industry. Many of these risk analyses are based on epidemiological studies (which

may estimate, for example, the health effects of small quantities of atmospheric pollutants such as radon gas in houses), or reliability studies (the likelihood that specific system components physically fail), but many are likewise not.

Policy based on technical risk analysis (TRA) is not invariably non-controversial. In its guise as Risk Cost-Benefit Analysis (RCBA), it has come in for considerable discussion and criticism.[3] This criticism has centered mostly on the use of a one-dimensional metric for assessing accidents and their severity. Use of RCBA requires that one can directly compare, say, an injury to a person with physical damage to an artifact. It leads in the extreme to such disquieting phenomena as "the monetary value of a life". Whatever one's view on such topics, this discussion does not concern us here, because we have made no assumption about what kind of creature the severity of an accident is. In our pursuit of technical risk analyses, we are not obliged to compare a break of an arm with a break of a leg, or even either injury with the monetary cost of successful treatment, let alone with the dent in an airplane. But for RCBA , all these are directly comparable, as "costs". A policy based on RCBA says: minimise expected cost (MEC). It is the use of MEC as a decision procedure which comes in for criticism. And this issue simply does not concern us here. RCBA is supported by TRA, as are other assessment efforts, but it goes further than TRA requires. MEC is only coherent as a policy if these further steps required by RCBA are taken. We aren't considering those steps here, so the controversy concerning MEC is moot.

A major example of TRA different from RCBA occurs in the investigations of commercial aircraft accidents, required of signatories to the 1948 Chicago Convention which set up the International Civil Aviation Organisation (ICAO). The purpose of such investigations is to identify actual physical risks on the basis of careful causal analysis of accidents that have happened. Such identified risks are not required to be causally linked to the actual accident – it has happened that some risk phenomena have been identified by an investigation into an accident in which they were not causal. But the vast majority of recommendations are based on factors which are causal to the accidents. Technical risk analyses such as I outlined are complementary to causal accident investigations. The former are intended to identify all ways in which accidents could occur, and, because of this counterfactual modality, the goals are correspondingly harder to ensure than the goals of an accident analysis, which is amongst other things to identify *the* way in which a single accident *did* occur.

Let us now formulate the assertion that a (correct) technical risk analysis is always appropriate: it always yields incontrovertibly relevant risk assessment information, if performed. Let us call this assertion TRA Validity (TRAV). I would argue that TRAV is incontroverted amongst professionals. I don't know of

[3]Such discussion may be found, for example, in [SF85, SF91] and in the pages of the journal Risk Analysis.

an example in the literature arguing coherently against TRAV. Using TRAV, we can formulate the argument above justifying PCJ from a TRA along with DF as: TRAV + DF entail PCJ. Incontrovertibility should be preserved by entailment, at least by any criterion of rationality. So it follows that PCJ is incontrovertible.

One should not be tempted to view this line of argument as saying that PCJ is only incontrovertible in cases in which a TRA has actually been performed. The principle TRAV is universal. It speaks, not to whether a TRA has actually been performed, but simply to the definition of what a risk is. PCJ is unrestrictedly incontrovertible.

PCJ is the most well-justified of the principles we have discussed. Because of its relationship to technical risk analysis, and its unrestricted incontrovertibility, PCJ is arguably also the most important of the safety principles enumerated above. However, PCJ cannot be assured by the conjunction of the principles BBSTS, PC, LIL, MJ. The counterexample is the HSE recommendation concerning phones on forecourts. As we have seen, that recommendation is justified, both overtly and indirectly, by BBSTS, PC, LIL, and MJ, but I have argued that, as far as we can tell, it violates PCJ.

This leaves policymakers such as HSE in a difficult position. If BBSTS, PC, LIL and MJ cannot assure PCJ, but nevertheless PCJ must be assured because it is incontrovertible, then there is no shortcut to assuring PCJ. A TRA at some level must be performed. Furthermore, one must wonder about the justification for any of the principles BBSTS, PC, LIL and MJ. None of them alone, and not even all of them together, can reliably substitute for a TRA, else they would be able to assure PCJ. But if they cannot reliably substitute for a TRA, what, then, could be a justification for employing them?

I believe BBSTS, PC, LIL and MJ can substitute for TRA in cases in which it is not felt that the identified risk could justify the resources one would need to spend on performing a full TRA. But, as the phones and forecourts case shows, these principles do not suffice as a substitute for TRA. PCJ, at least in the weakest form of exhibiting a causal path to an accident, needs to be assured independently of these other four principles.

This book is concerned with the determination of and analysis of causal paths, both potential causal paths to failure or unsafe operation in systems being designed, and actual causal paths that have occurred in failures and accidents. The analysis of causality allows the use of a form of exact science which the other principles do not, at least not overtly.

First, I examine some of the social and pyschological background to technical risk.