Applying numbers to the cheese

Disclaimer

This paper reflects the author's opinion, and does not necessarily reflect any opinions held by the Australian Transport Safety Bureau.

The chance of a bad thing

Statisticians have an easier life if the distribution of individual measurements from a population forms a 'normal' distribution. A normal distribution is popular because it is relatively easy to describe in statistical terms, and allows a statistician to predict the characteristics of a whole population based on observations from just a sample of that population.

Many references describe a normal distribution as 'a distribution that fits the normal distribution curve'. This is not a practically helpful definition.

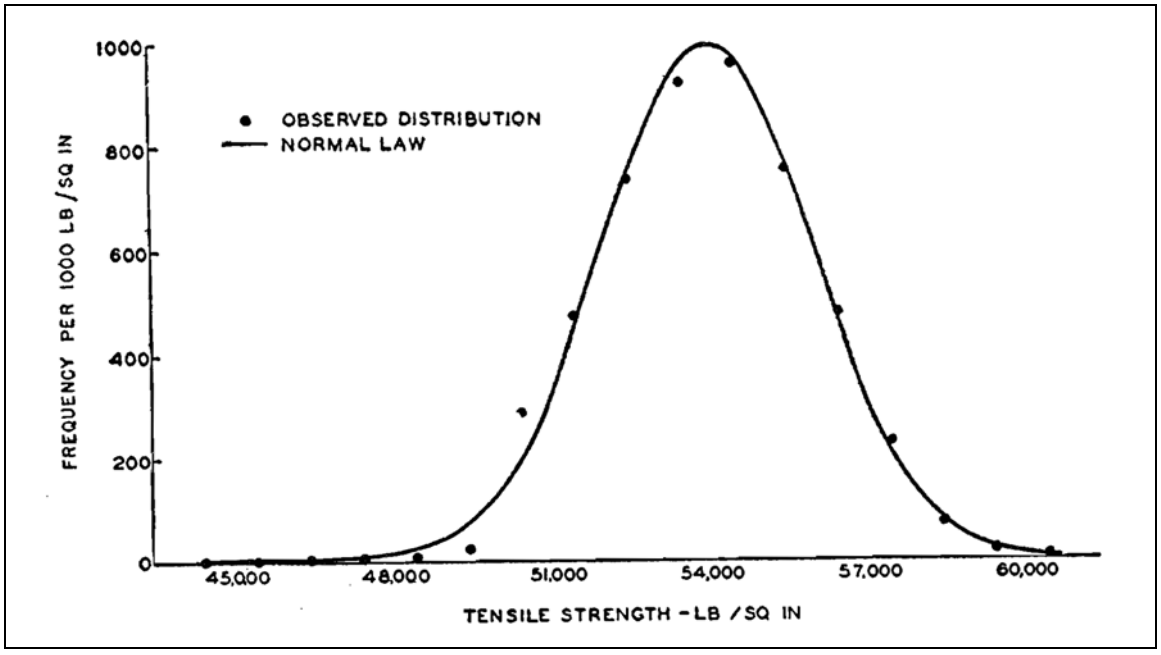

Figure 1: Distributions compared with the normal law, from *Statistical method from the viewpoint of quality control*, p55.

Normal distributions generally reflect measurements of a population in which independent influences apply to the population to direct a measured characteristic toward a desired result. This wide-ranging definition can be applied to manufactured products, whether it relates to measured tolerances such as the actual quantity of beer in a bottle of manufactured beer, or the performance of a manufactured product such the life of a light bulb or the structural strength of a mass-produced component. Indeed, the use of normal distributions formed the basis of quality control in mass production, to make the fortunes

of American manufacturers in the early 20th century. The same principles are still in use today, as described in Walter A Shewart's book *Statistical method from the viewpoint of quality control,* first published in 1939 and still in print today.

Normal distributions can also be seen in the natural sciences. However I have found some academic papers that refer to IQ, dexterity test results, or specific anatomical measurements within a population, and then say'...**assume** a Normal distribution...' [my emphasis]. They then continue to make further assumptions which are reliant on the original assumption that a normal distribution existed in the first place.

Presumed normal distributions have been used in statistics, economics, manufacturing quality control, even forecasting the distribution of IQ in populations where a researcher or manager may be trying to use the distributed data to create useful predictions or conclusions. The normal distribution is known as a 'thin-tailed' curve, as the number of measurements that are distant from the preferred measurement diminish quickly, compared with other types of distribution.

Most flights are completed safely. They fall within the area of the highest frequency density under the distribution curve. The more a flights moves away from the highest point of the curve and toward the tail, the higher the probability that one or more safety-critical parameters may be exceeded, and safety can no longer be assured. As aircraft accident investigators, we should threfore be working toward designing aircraft systems to be operating as far away from the thinnest end of the frequency distribution curve as we possibly can, because that will show our success in reducing the frequency of parameter exceedences and aircraft accidents. The rest of this paper will be examining the very end of the frequency distribution curve, where the subjects that we study will normally reside.

Over the past century or so, occasional statistical observers have described **power law**, or exponential distributions that are *scale-invariant*, unlike a normal distribution. A scale invariant distribution is one where the change in frequency will be the same across a defined change in the measurement, wherever you look at the frequency graph. Here are some examples:

- In 1897, Vilfredo Pareto discovered the 'wealth' power law, showing that the ratio between the number of people who earned a certain income and the number people who earned double that income was the same, across a very broad spectrum of income earners.
- In 1913, Auerbach discovered a similar type of frequency distribution existed across the frequency of town sizes across the US (Figure 2). Krugman repeated the study (1996), and said, *We are unused to seeing regularities this exact in economics*.

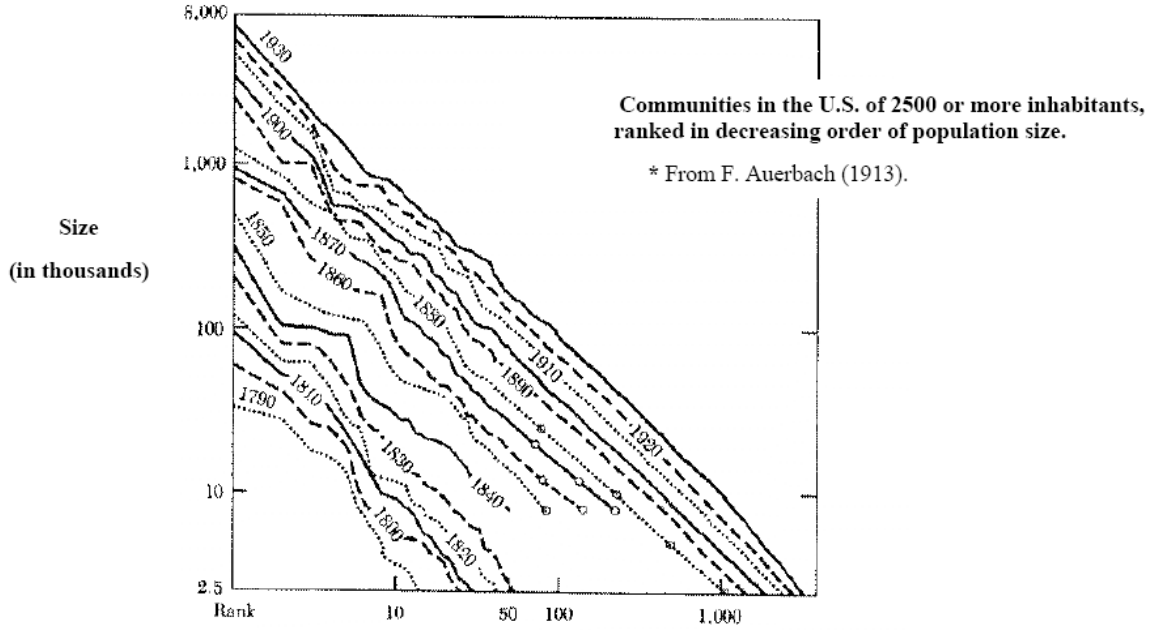

Figure 2: Log log depiction of town rank by size as a (-1) slope. *From Auerbach, (1913)*

- In 1949, Zipf found the same type of frequency distribution existed across the frequency of particular words used in a language, whatever the language. A constant ratio existed between the number of words used at a certain frequency and the number of words used at half that frequency. The ration remained constant, whatever the starting frequency of word use.
- In 1975, Benoit Mandelbrot developed fractal geometry, in which 'a pattern or shape whose parts resemble the whole', demonstrated in the shapes of coastlines or snowflakes, irrespective of the size of the observed pattern.
- In 1996, Per Bak described the frequency of different avalanche sizes as each grain of sand was dropped on the top of a sand pile, again forming the same type of frequency distribution known as a power law distribution. A similar result has been demonstrated for the frequency of earthquake magnitudes.

Power law distributions are notably different from normal distributions in that the population frequency in a power law distribution does not diminish as quickly as the graph moves away from the highest frequency density, compared with normal distributions. A graph of a power law distribution is known as a fat-tailed curve (Figure 3). If you look at the examples described above, they are not influenced so much by one variable, as being capable of being influenced by a number of variables.

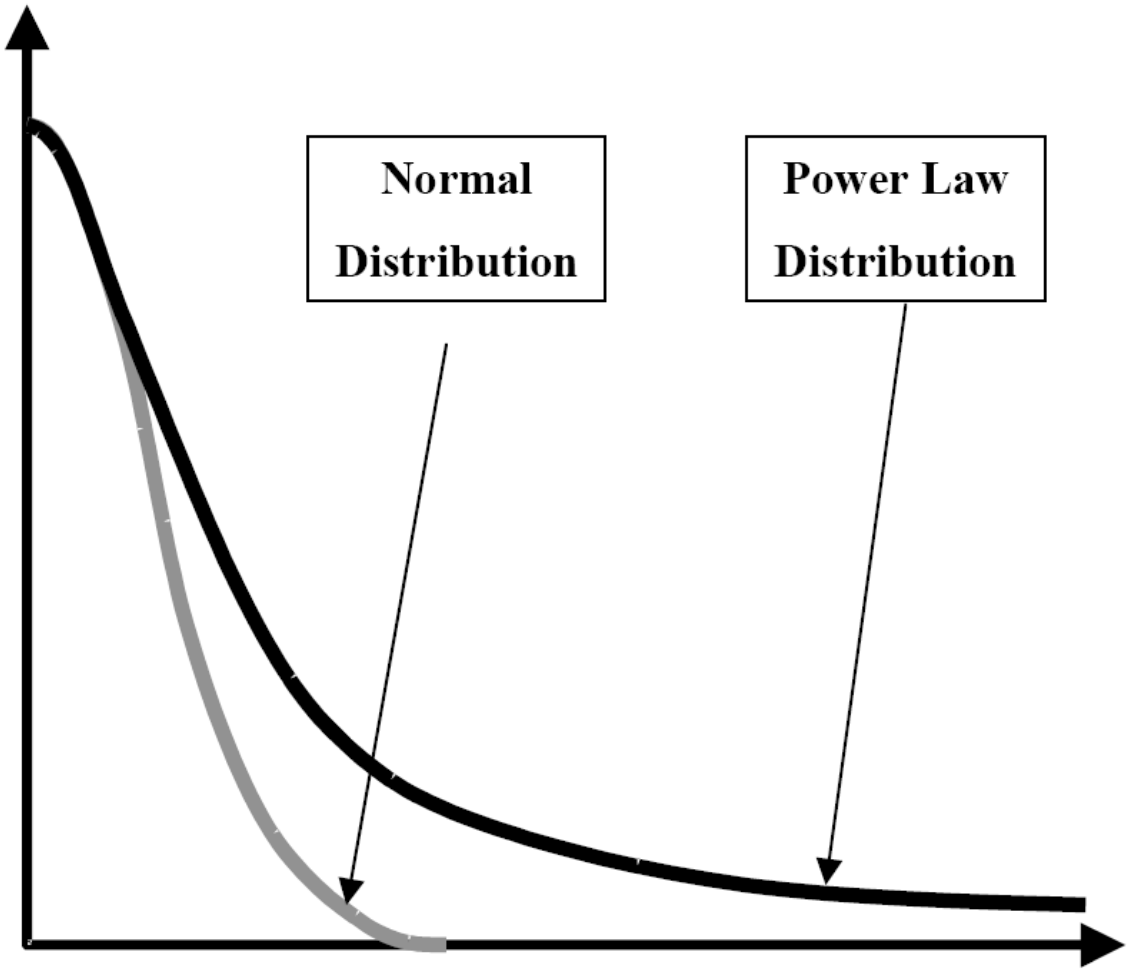

Figure 3: Comparison between normal law and power law distributions *Andriano and McKelvey 2005*

Adriano and McKelvey (2005) described the physical characteristics between a normal distribution or a power law distribution.

> **A normal distribution** is a distribution where the data points are influenced by one or more factors (or random variables) that move individual population measurements, but the factors are independent from each other. One factor cannot influence another factor. As the factors are independent, their combined influence is the sum of their influences, and the data points are assumed to be *independent-additive.*

I consider that the distribution of readings on fuel quantity gauges across a serviceable fleet of an aircraft type, each with the fuel tanks half full, would hopefully tend toward a **normal distribution.** The readings will (hopefully) be similar, and influenced by independent variables: maybe one aircraft has a low battery voltage, and another aircraft

has an abnormally tall pilot who has to lean over to read the gauge, leading to parallax errors. The variables are mutually independent, and randomly distributed.

> **A power law distribution** is a distribution where the data points are influenced by one or more factors that that can move individual population measurements, but the factors are not independent from each other. The factors themselves may influence each other, and the mutual influences means that the variables cannot be random. The presence or existence of one factor may influence the power of another factor to influence the population measurement. As the complexity of the *independent-multiplicative,* or *interdependent*, factors increases, the frequency distribution more closely reflects a power law distribution (West and Deering, 1995).

I consider that the runway required for take off for a fleet of an aircraft type will more closely follow a **power law distribution**. Although the performance manual allows headwind component, density altitude, runway surface, runway slope and takeoff weight to be factored into the required take-off distance, there are also assumptions that engine power produced, airframe cleanliness and configuration, weather turbulence and pilot ability will all meet a consistent standard. Most of these variables will influence at least one other variable: an underperforming powerplant will have an increased influence on an up-sloping runway, a turbulent day may tempt an inexperienced pilot to add an extra speed margin for take off, compared with an experienced pilot.

The difference between the two curves' tails

The frequency curves for a power law distribution and a normal law distribution are very similar where the frequency distribution density is at the greatest, where the curves are at their highest point (Figure 3). However, the further the graphs move away from this point, the greater the difference between the two graphs. The likelihood of a 'rare event' or an 'outlier' becomes greater under a power law distribution compared with a normal law distribution, as the graphs move further into the rare event area of the graph. The consequences of forecasting risk at the tail-end of the wrong distribution curve have been described well by Taleb in *The Black Swan*. An example is provided to determine the probability of being extremely rich: the normal distribution predicts 1 in 6,400,000, whereas the power law distribution predicts 1 in 20,000 for the same level of unusually high personal wealth. Fortunately, Taleb justifies the lower odds!

This difference was noted by Mandelbrot and Hudson, (2004) who observed that if the normal law was applied to the movements of the Dow index, there should have been fifty-eight days when the index moved more than 3.4 per cent in a day between 1916 and 2003, but in fact there had been 1001 days when this had happened. This difference may suggest that the frequency distribution more closely reflected many *interdependent* influencing factors that may have influenced movements in the index, compared with the influence of mutually *independent* factors. Indeed, stock market behaviour is renowned for one factor amplifying the influence of another factor, sometimes leading to a domino effect with little immediately apparent reason.

Consider the many limiting parameters that must be applied to ensure that a flight will be safe. The International Civil Aviation Organisation Continuing Airworthiness Manual[1] '... assumes, arbitrarily, that there are about 100 potential failure conditions that would prevent continued safe flight and landing.' (The number is approximate because of the many variables between different types of flight, however it provides a reasonable starting assumption.)

In the associated presentation, limits to weight, balance, weather, design component life, aerodynamic limits, maintenance, operating range, mechanical reliability, operating speeds, runway required and pilot capability have been used as examples of safety-critical parameter limits. If you select pairs of these limits, one can find many cases where one limit may influence another limit. While mechanical reliability may not affect the weather, the opposite may be true: an aircraft that flies through thunderstorms is less likely to be mechanically reliable, (all other things being equal). In the same way, pilot handling capability has the potential to greatly influence mechanical reliability.

On this basis, I argue that many of the operating parameter limits that must exist for safe aviation are interdependent. As such, the distribution of operating parameter variations should follow the power law distribution curve more closely than the normal law distribution curve.

What is the importance of the difference between these two distributions when applying them to manage aviation safety? The two distribution types are accepted to have several important differences.

- **Usefulness of the mean:** In a normal distribution, a stable mean for the population can be identified. However, under a power law distribution, what appears to be a 'mean' will be strongly and idiosyncratically influenced by extreme events. Extreme events will move the position of the mean, making it unstable and of little use as a statistical indicator.

- **The influence of population size on stability:** In a normal distribution, the larger the population size, (or the sample size in relation to the population size), the greater confidence you may have in the mean and the variance of your distribution. In contrast, a power law distribution has a very large variance (influenced by extreme events), and confidence intervals will vary with the occurrence of each new extreme. In a power law distribution, the larger the population, the greater the variability; which is the opposite of a normal distribution. The important part of a power law distribution is in the tails.

- **Scale-free fractal structure:** This argument about the significance of extremes in larger

[1] The International Civil Aviation Organization (ICAO) Continuing Airworthiness Manual 1995 (Doc 9642-AN/941)

– populations is demonstrated in the fractal patterns seen in snowflakes or coastlines. Any given line or shape in a fractal pattern will change shape or direction to the same degree, irrespective of the size of the line or shape. If you look at a small fractal image, then the most extreme change will be relatively small (and hence manageable), but the larger version of the same image will produce a correspondingly larger change which may be harder to manage. The larger viewpoint (or population under study) has a greater chance of a correspondingly larger variation. Extremes influence power law distributions and amplify the variability, whereas a normal distribution tends to compress the distribution of data points toward the mean as the population gets larger (Figure 4). In a normal distribution, outliers are normally ignored and the assumption of mutual independence supports any justification to restrict predictions to within two or three standard deviations from the mean.

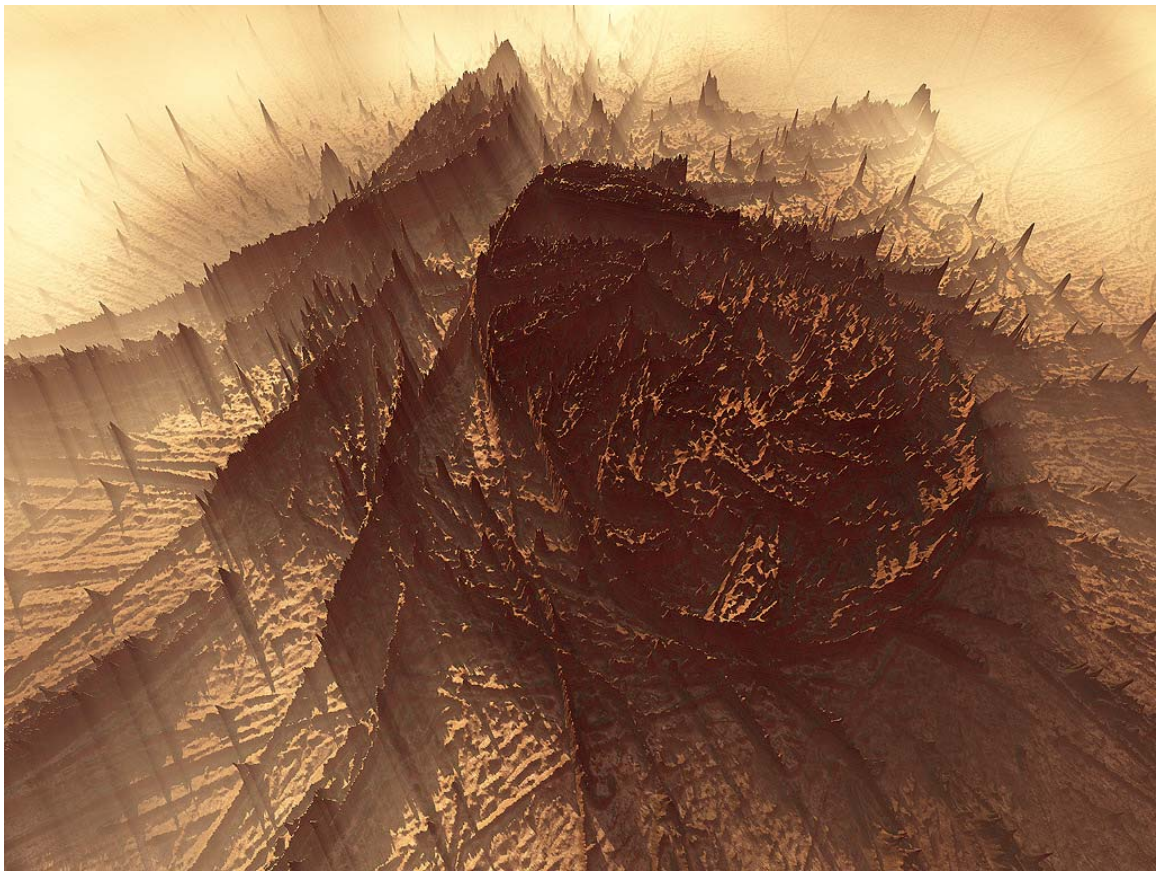

Figure 4: Is this snow fallen over a burnt forest in Siberia, or a microscopic photo of a bacterial infection?

– **Interdependent sequential influences in a dynamic situation:** Bak (1996) described how power laws can result from self organised criticality dynamics. He measured the extent to which events unfold from an initial instigating stimulus. Given mutual causal, positive feedback mechanisms, (as happened in his

avalanche sizes in critically stable sand-piles), the logic of preferential attachment generates a reinforcing trend, (a larger avalanche is more likely to develop into an even larger avalanche). This process extends and fattens the tails of power law distributions.

Implications for extremely safe flight

Staying within operational parameter limits

Flying an aircraft is a dynamic operation. Events and changes happen throughout a flight, and the events and changes can be mutually interdependent. To ensure safe flight, the interdependent safety-critical parameters must all remain within their limits. The probability that all the parameters will remain within their limits will therefore fall under a **power law distribution**. The fact that all the certificated[2], safety-critical parameters have to remain within their limits, then the total reliability should be calculated from the individual reliabilities in a series reliability system, total reliability is based on (reliability A) AND (reliability B) AND (reliability C) AND…(reliability n). This 'Bayesian' method is described in the associated presentation.

However, this method will not work so nicely because it is a **power law distribution**. Random outliers in the fat tails of a power law distribution have to be planned for, and managed. The problem is that one variable will influence another variable, and as the variables influence each other, there is an increasing chance that a parameter will be exceeded by the influence of other varying parameters, and certificated safety can no longer be assured.

If we cannot help but work under a power law distribution as we aim to remain within all the necessary parameters: the less that one parameter is varied, then the less the chance that variable will vary another parameter. The risk of parameter exceedences is reduced if there is as little operational change as possible. This makes intuitive sense, and supports the evidence from Boeing that the large majority of airliner catastrophic accidents do not happen in cruising flight, where parameters are changed less as the aircraft flies straight and level, with less parameter change.

If the fat tail, power law frequency distribution does succeed in creating a parameter exceedence, then the following section will apply.

Returning to operational parameter limits.

Systems exist to monitor safety-critical parameters during flight, enabling confidence that operating crew can become aware of parameter deviations in sufficient time to control the deviations. However, because aircraft fly in an imperfect world, and because the influence of parameter deviations on the safety of flight falls under a **power law**

[2] Certification is the process whereby a regulator exhaustively satisfies itself that a new aircraft type will operate predictably, safely and reliably across all possible combinations of defined parameter limits.

distribution, multiple defences are also maintained to identify and manage parameter exceedences when they occasionally, but inevitably happen.

All the parameter control systems have to work if safe flight is to be ensured, however only one defence has to work in order for a parameter exceedence to be controlled. Therefore, while the parameter control systems must work in series when calculating their total reliability, only one of multiple defences, that work in parallel, has to work in order for an exceeding parameter to be returned to within its limits. The total reliability of the defensive system can be calculated from the individual reliabilities in a parallel system. Total reliability is based on (reliability A) OR (reliability B) OR (reliability C) OR …(reliability n). This 'Bayesian' method is described in the associated presentation.

Calculating the probable success rate of a set of independent parallel defence systems is therefore different, compared with the probable success rate of a set of series parameter control systems. The principles behind maximising the probability of ensuring that all of a set of interdependent conditions always work are very different from the principles behind maximising the probability of ensuring that one of a set of interdependent conditions always works.

As the influences that affect each defence system's reliability is (hopefully) independent, then each defence's reliability is more likely to fall under a **normal distribution** curve, with less potential for outliers to compromise the defence system's reliability.
To calculate the reliability of all the independent, parallel defence systems together, one multiplies the individual unreliabilities to calculate the combined unreliability. An increase in the number of comparably effective, independent defences will provide a geometric improvement in overall reliability, so long as the assumption of mutual independence between the individual defences can be maintained.

If, however, the individual defences become interdependent, then the distribution of the overall defence reliability will tend toward a **power law distribution**, with an increased probability of uncaptured, undesired outlier events. Inadvertent interdependence will compromise the effectiveness of the total defence system.

One possible example of inadvertent interdependence would be if one defence was a predefined set of rules, and other defences existed to reassure compliance with those defensive rules. In that case, the defences would be interdependent and the system would be both less reliable and less easy to predict. The degree of interdependence might be reduced by expecting different defensive layers to develop their own defensive systems, independent of the expectations of the other independent systems. The workload behind this approach would be great, however, and the cost/ benefit would have to be assessed independently.

Examples

An aviation story

One of the most multi-faceted investigations undertaken by the Australian aviation safety investigator was the Investigation into Ansett Australia maintenance safety deficiencies and the control of continuing airworthiness of Class A Aircraft, reported to this conference in 2005.

This investigation did not look at an accident, or even an incident. It did, however, look at an event when an airline was astonished to find that it had lost confidence in its own ability to keep its aircraft safe. The investigation looked at the management of safety-critical knowledge that provided confidence in an aircraft's structural integrity: incomplete knowledge, led to no confidence.

Operating an airline is a dynamic process, with many stakeholders storing, transferring and using safety-critical data. Any possible inconsistencies in interpretation, transfer or use of that information would have an interdependent knock-on effect, for example in the international cycle of continuing airworthiness information. For many, individually plausible reasons, the defences in place to protect the use of this information were eroded, and when a relatively minor exceedence in the use of that information (one maintenance task was missed), the remaining defences in place could no longer provide assurances that all the necessary safety actions had happened.

The global system for managing the airliner fleet's safety-critical airworthiness information depended on sharing safety information among every operator, component manufacturer and associated regulator. The system was immensely complex, highly networked and quite interdependent. The influence of varying factors would therefore followed a **power law distribution**, so considering the principles described above for maintaining an extremely high level of safety, it would have been prudent to have maintained multiple, independent, parallel defences, to help maintain a high level of reliability.

The consequences of the gradual erosion of the defences meant that the overall safety capability had been eroded far more rapidly. The erosion of the defences was an arithmetical progression, but the consequential erosion in safety was a geometrical progression. The probability of capturing the loss of safety-critical information was dramatically eroded with each relatively small reduction in the defences. The incremental reduction in the defensive cheese slices had a far greater compounding effect on the safety performance of this high reliability organisation.

A recent non-aviation story

The magnitude 9.0 Tohoku-Taiheiyou-Oki earthquake at 2.46 pm on 11 March 2011 did considerable damage. It was a'1 in 1000 year' earthquake and the previous earthquake of a similar magnitude had happened in 986 AD[3]. Eleven nuclear reactors at four nuclear power plants in the region were operating at the time and all shut down automatically

[3] http://coastalcare.org/2011/03/nuclear-plant-and-tsunami-risk-3000-years-of-geological-history-disregarded/

when the quake hit. Power was available to run the cooling pumps at most of the pumps at most of the units, and they achieved cold shutdown in a few days[4].

The Fukushima reactors were exposed to horizontal accelerations at or below their design acceleration tolerance, and they were not damaged as a consequence of the earthquake-based accelerations. The shutdown procedure was continuing normally at Fukushima until the Tsunami submerged the diesel generators providing backup power to the cooling pumps, and the diesel generators stopped working. In the absence of pumped cooling, the reactors became thermally unstable.

The safety system for automatic reactor shutdown worked normally until the earthquake. The earthquake was the initiating mechanism for the automatic shutdown, but italso compromised the safety backup for the normal electricity supplies required to conduct the automatic shutdown. The earthquake initiated the sequential events through different mechanisms, one by shaking and one by submerging, so the shutdown of the reactor and the loss of electric power were interdependent. The earthquake also compromised other defences, including the normal electrical power supply, further enhancing the interdependence of the safety defences and making the probable reliability of disaster recovery fall more under the **power law distribution**.

Recent press reports indicate the reactor's operator was basing its earthquake risk analysis on presumptions associated with normal distributions of adverse event frequencies[5]. The two critical safety defences, the automatic shutdown and the backup power supply, were highly interdependent with a common initiating source for potential failure. This meant that a **power law distribution** may have been a better model than a normal distribution for planning and designing defences against extreme events. Furthermore, as the [admittedly uncontemplated] defences failed, there was a geometric risk increase of an unsafe outcome.

The preventive lessons from the fist example may well have been applicable to the second example cited here.

Summary

Frequency distribution
Plausible arguments exist to suggest that the frequency distributions of accidents in high reliability airline organisations will follow a power law distribution more closely than they will follow a normal distribution.

The power law distribution suggests that there will be a higher frequency of catastrophic events than might be predicted from assumptions made from assuming a normal frequency distribution of aircraft accidents, or from extrapolation of less complex incidents which may be more likely to follow a normal distribution. It is also harder to

[4] http://www.world-nuclear.org/info/inf18.html
[5] http://www.haaretz.com/news/international/study-japan-nuclear-plant-chiefs-downplayed-evidence-of-tsunami-risk-1.352189

predict accident frequency form power law distributions compared with normal distributions. Alternative models for assessing safety may therefore be beneficial when contemplating catastrophic accidents in complex, networked operations.

Maximising safety during normal operations
It is argued that safety will be better maintained if interdependent safety-critical operational parameters are changed less during flight, compared with the alternative of regularly changing operational parameters to keep them as far away from their operational limits as possible. Regular parameter change without good justification may increase the risk of knock-on effects from the change of one parameter, to one or more other parameter.

This argument supports the demonstrated high level of safety in airline operations that are designed to operate at maximum efficiency within their certified safe flight envelope. Airliners normally operate with maximum payload, minimum fuel, and close to their environmental operating limits. Despite this, they remain very safe, possibly because they immediately seek to fly near to their operational limits, and then change parameters very little while they carefully monitor their safety margins until it is time to come down again. Records indicate that the risk of a catastrophic accident increases during the first and last 15 minutes of flight, compared with the cruise.

Maximising safety during abnormal operations
Once a safety-critical operational parameter has been exceeded during flight, then the certified safe operational envelope can no longer provide an assurance of safety. This is the time when operational defences help to identify the parameter exceedence, and then return the parameter to safe limits. Defences are designed to be robust, yet because they are designed to manage the unpredictable, they cannot be assumed to be reliable. If all the defences fail, then the opportunity to recover back to safe flight is lost. As only one defence has to work, the reliability of the entire defensive system is best enhanced by adding more independent defences, compared with improving existing defences. Five reasonably good defences are probably safer than one excellent defence.

Multiple defences are largely compromised if they are mutually interdependent, so the reliability of the entire defensive system will be enhanced if the defences are interconnected to the minimum extent possible.

Auerbach, F. 1913: *Das Gesetz Der Bevolkerungskoncentration*. Petermanns Geographische Mitteilungen **59** 74–76.

Bak, P. 1996: *How Nature Works: The Science of Self-Organized Criticality*. Copernicus, New York.

Buchanan, M. 2000: *Ubiquity: The Science of History…or Why theWorld Is Simpler Than We Think*. Weidenfeld & Nicolson, London.

Krugman, P. 1996: *The Self-Organizing Economy*. Blackwell, Malden, MA.

Mandelbrot, B. B. 1975: *Les Objets Fractals: Forme, Hasard et Dimension.* Flammarion, Paris.

Mandelbrot, B. B. 1982: *The Fractal Geometry of Nature.* Freeman, New York.

Mandelbrot, B. B., Hudson R. L.. 2004: *The (Mis)Behavior of Markets: A Fractal View of Risk, Ruin and Reward*. Profile, London.

Pareto, V. 1897: *Cours d'Economie Politique*. Rouge, Paris.

Pierpaolo,A., McKelvey, W. 2005 *Beyond Gaussian Averages: Redirecting Organization Science Toward Extreme Events and power Laws*. http://www.dur.ac.uk/resources/dbs/faculty/working-papers/BeyondGaussianAverages19Jun06.pdf

Shewart, W.A. 1939: *Statistical method from the viewpoint of quality control*, The Graduate School, Department of Agriculture

Taleb N.N.: *The Black Swan*, Penguin, 2007 Ch 15

West, B. J., Deering. B.1995: *The Lure of Modern Science: Fractal Thinking.* World Scientific, Singapore.

Zipf, G. K. 1949: *Human Behavior and the Principle of Least Effort*. Hafner, New York.

Disclaimer: This is the presenter's personal opinion. It does not necessarily reflect any position held by the ATSB, and it is not endorsed by the ATSB.

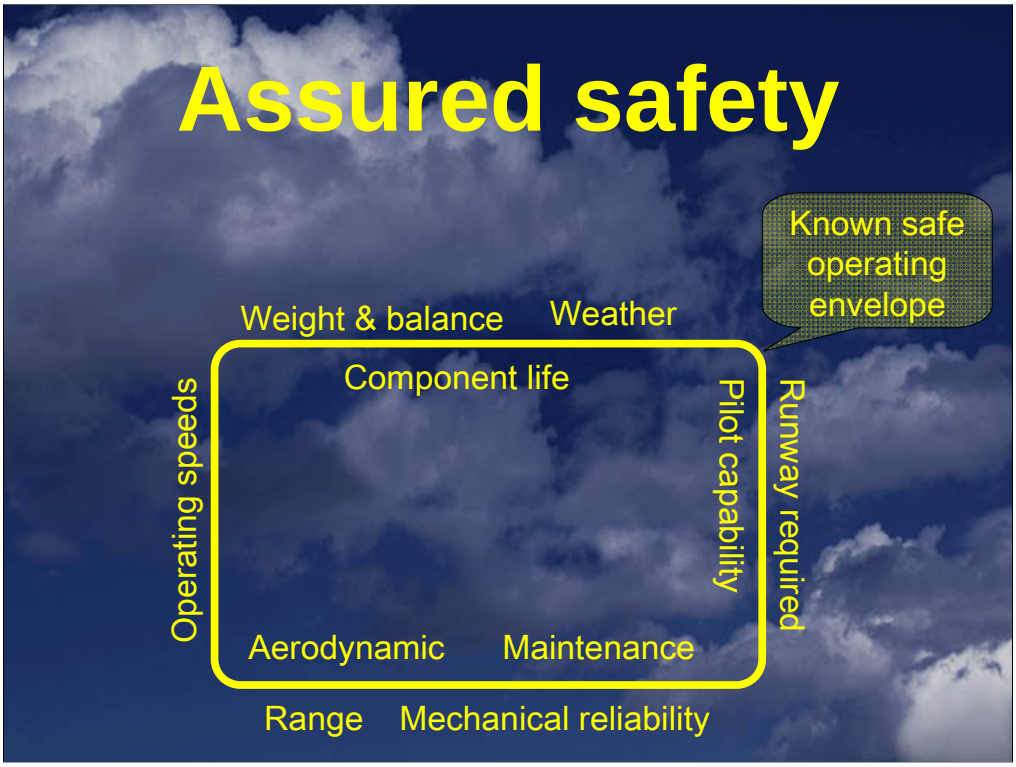

Concept of Certificated safe operating envelope. Certification assumes flight safety to exist, if all safety critical parameters remain within their defined and proven envelopes.

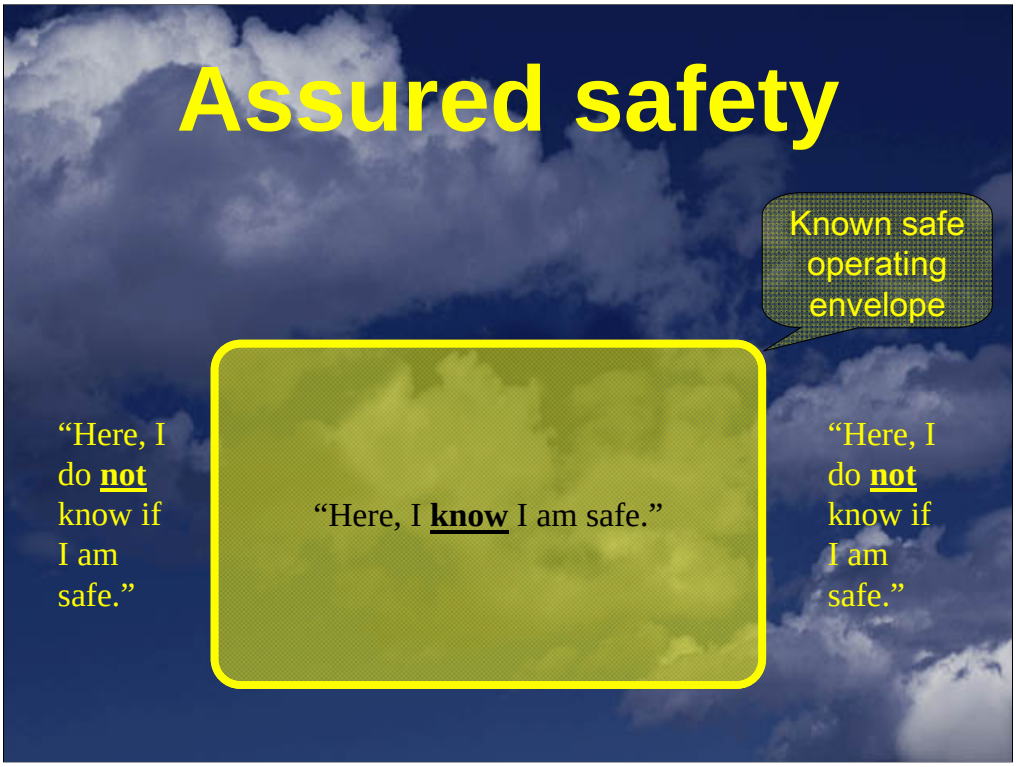

If you are outside the envelope, it doesn't mean that you are in a positon of danger. It means that you don't know that you are not in a position of danger.

You are playing Russian roulette, (probably with very few bullets in very many chambers,) but you are still no longer proven safe.

This normal, 'how it should happen', behaviour, similar to a model that many line production managers work with. Akin to page 200, Figure 7.4 of Reason (2000) Human Error.

Notice that a change parameter in one will affect others, e.g, runway required affected by pilot capability, aerodynamic, Weather, W&B, maintenance, showing the complex, interacting nature.

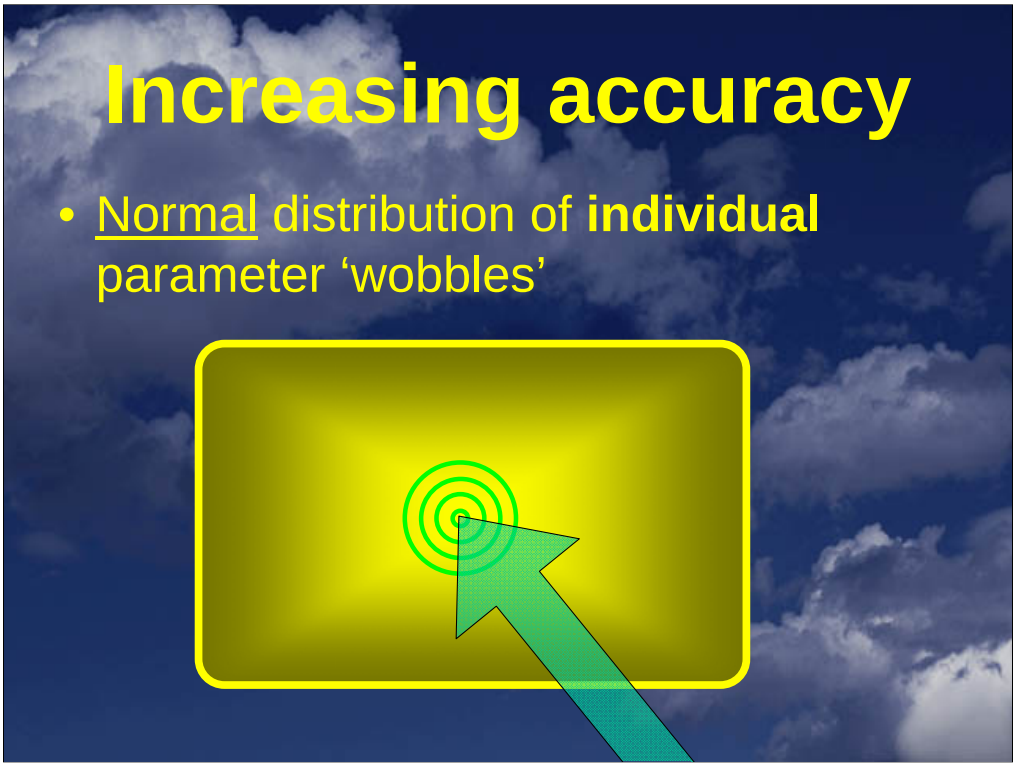

Consider just one, individual parameter.

Example, normal distribution of inaccuracy in measurement of fuel quantity, sensor accuracy, meter inaccuracy, fuel density gauge readibility, etc.

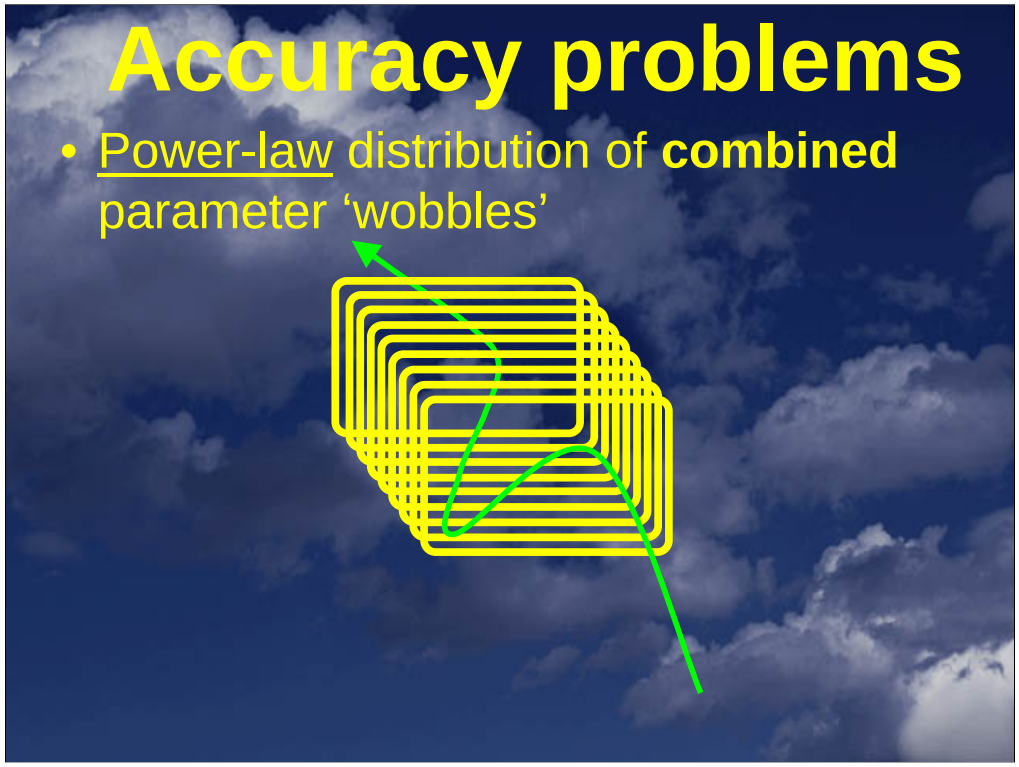

Parameter deviations are interdependent in a complex High reliability organisation (HRO) system. Reason (1990) says systems are now

•More automated

•More complex

•More dangerous

•Have more defences against failure

And most importantly [my thought]

•**More opaque**

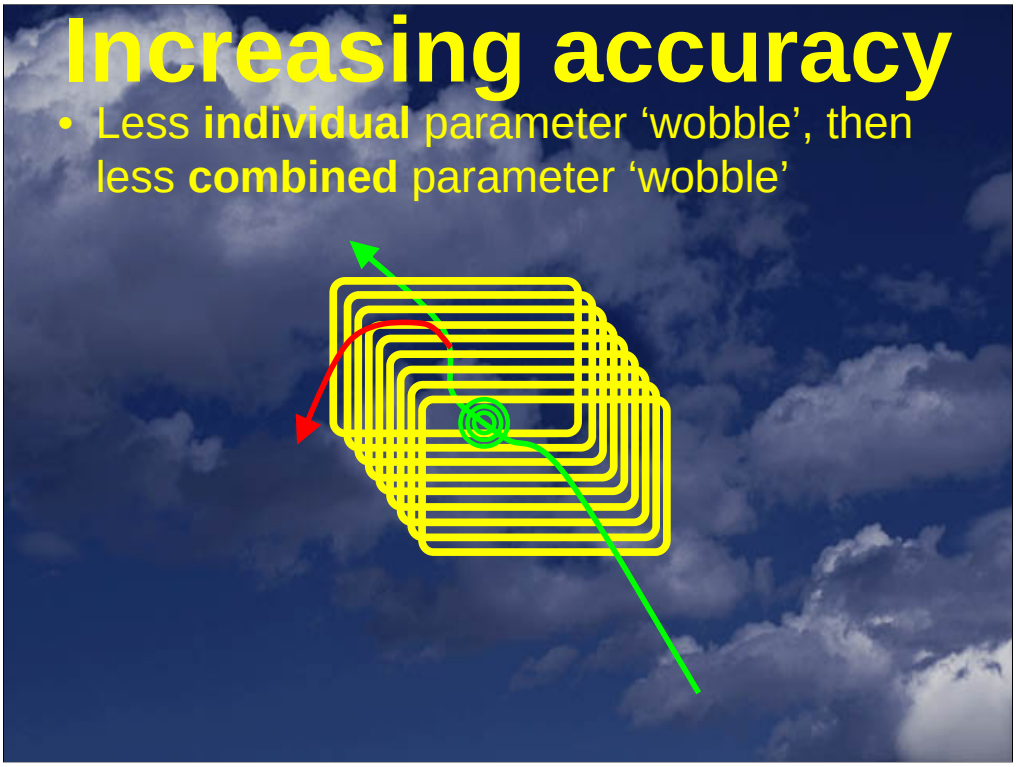

Deviations within acceptable parameters are modified to ensure they never go outside the normal parameters, these should be**,** 'fine tuning' behaviours, usually used to maximised productivity or efficiency.

However, the more that one parameter does not deviate (or deviates less), the greater the chance that other interdependent parameters will not be amplified into a subsequent parameter exceedence.

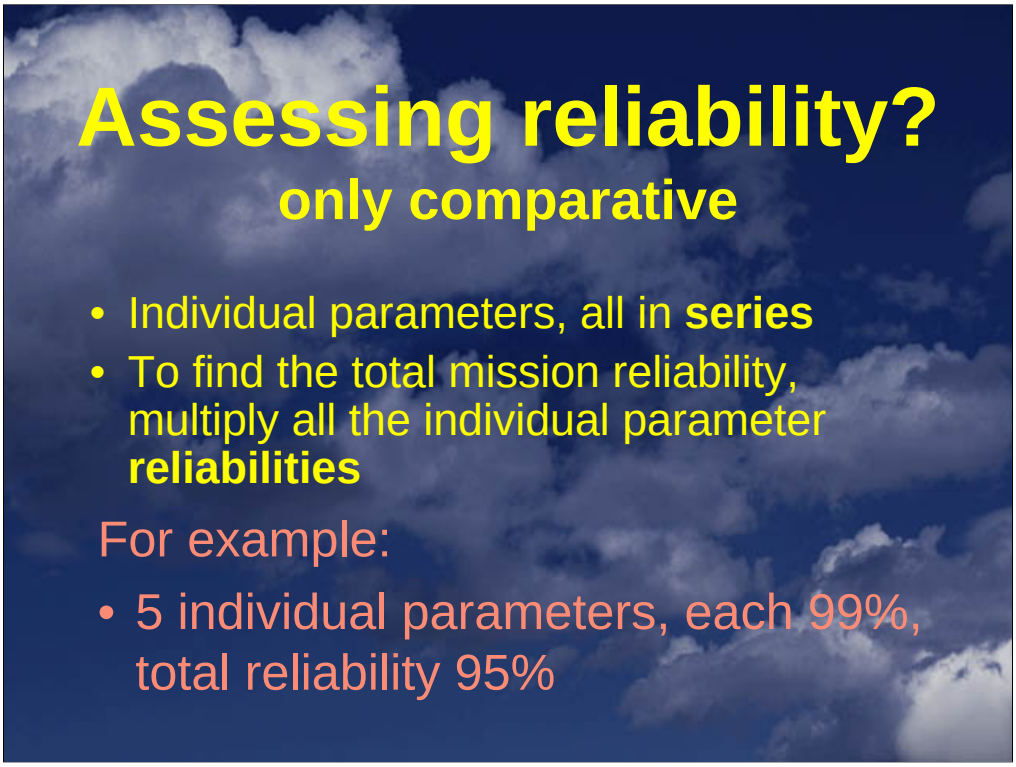

In series, the total reliability is the product of the individual reliabilities, (i.e for a pretty reliable system.)

If the distribution follows a power law, then a mean should not be used as a predictor of future events.

It can, however be used to compare against other calculated probabilities, to assess which is the least risky, but it cannot tell how risky any particular option is.

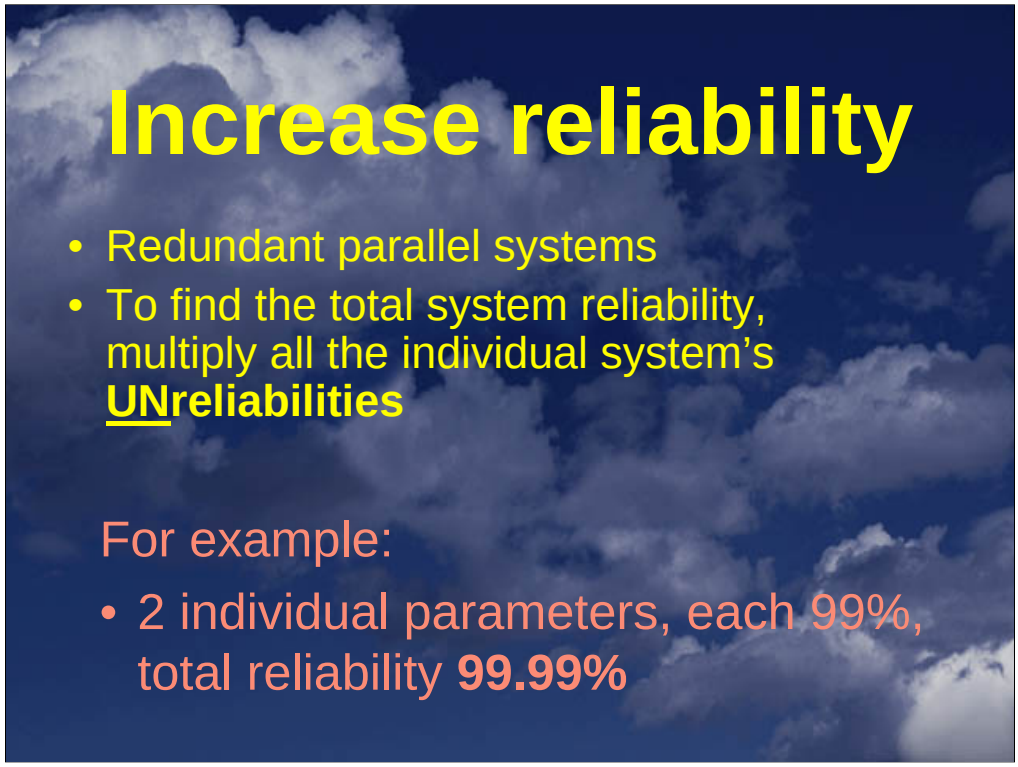

In parallel, the total unreliability is the product of the individual reliabilities (i.e. much more for a pretty reliable system)

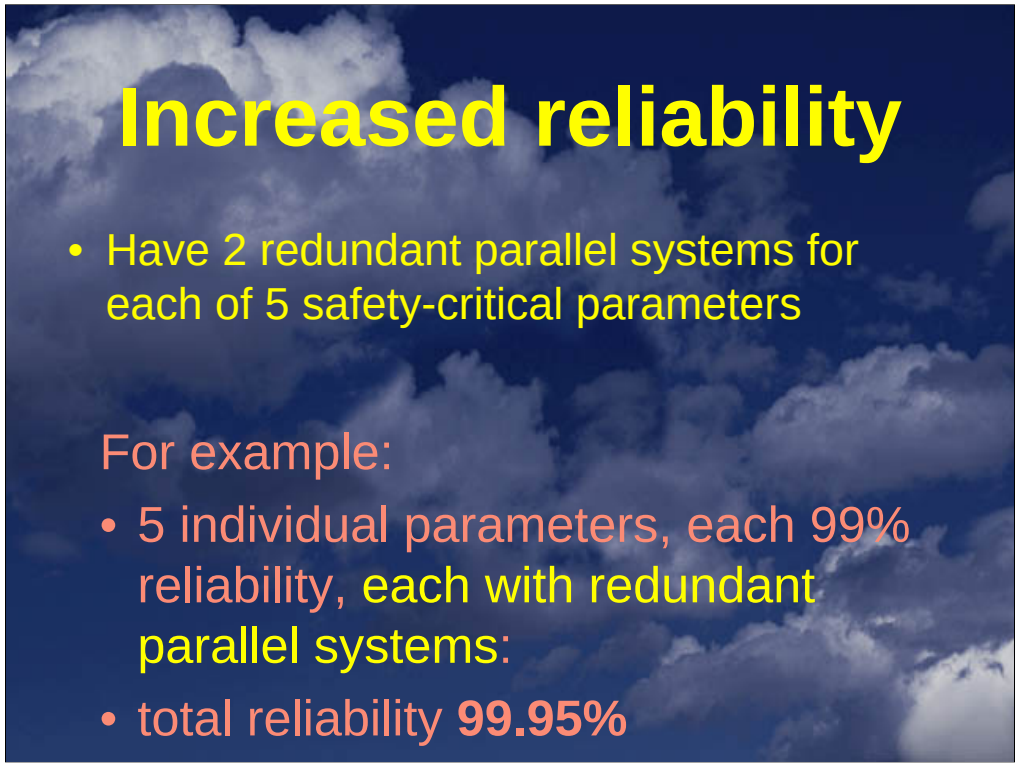

Built-in parallel independent subsystems incorporated into a series system generate a dramatic improvement in the overall reliability, compared with a series system alone.

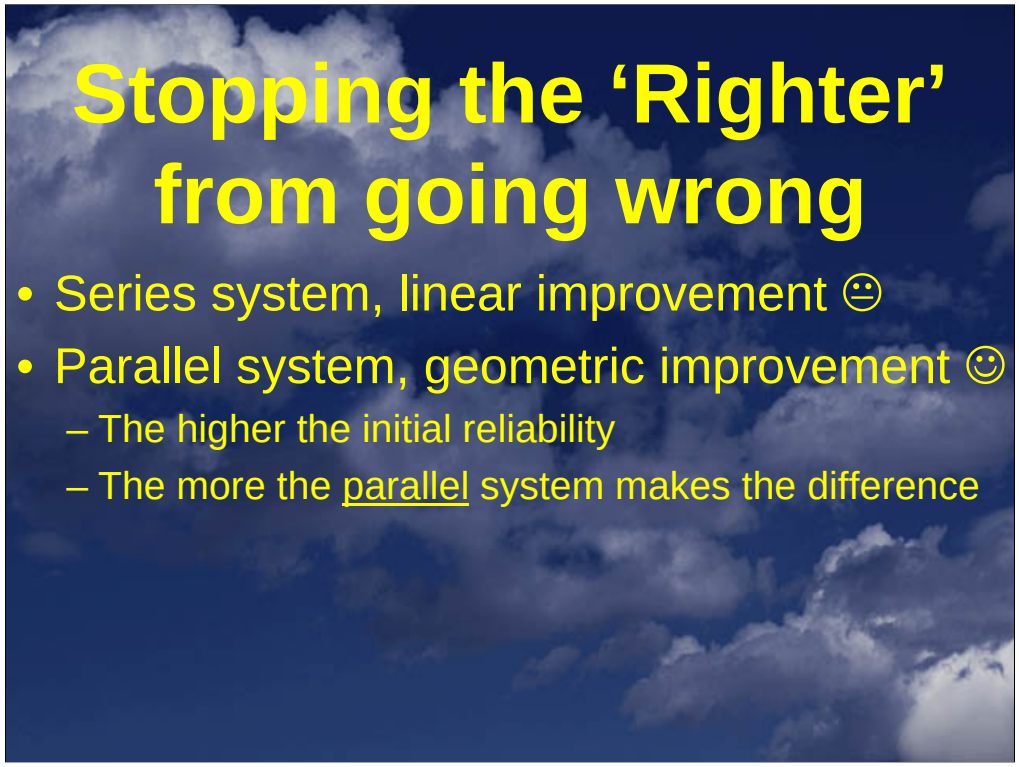

Series system, think of Shewart, and Seiko vs. Timex watches. Or Leyland vs Toyota reliability.

The Japanese 30-40 years ago made things to a much higher tolerance than required under Shewart's presumed normal distribution system, and found that the up-front capital investment created a paradigm change in improved reliability and customer expectation.

Stopping the 'Righter' from going wrong

- Assumptions:
  - Redundancy system is **reliable-** the remaining system will always take over
  - Redundancy system is **independent-** one system's failure mechanism will not affect the surviving system.

Reliable take-over; e.g. independent nav systems, runways, flight crew.

Independent: consider engine failures in these 4 engine aircraft.

The 707 was a 4 engine aircraft.

The Vulcan was a 4-engined twin without the 2-engine performance capability. If one engine spat a blade out the front ; it was usually hoovered by its adjacent engine.

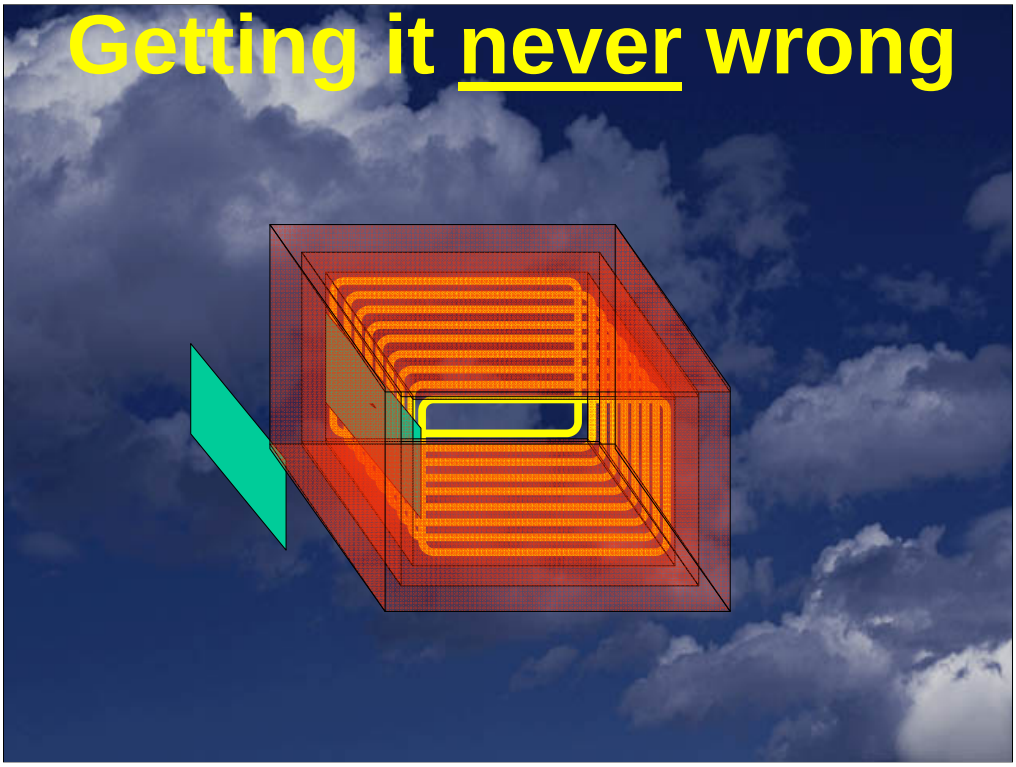

At the end of this slide explain that the next slide is looking from the outside to the left, through the 2 green squares,

Getting it never wrong

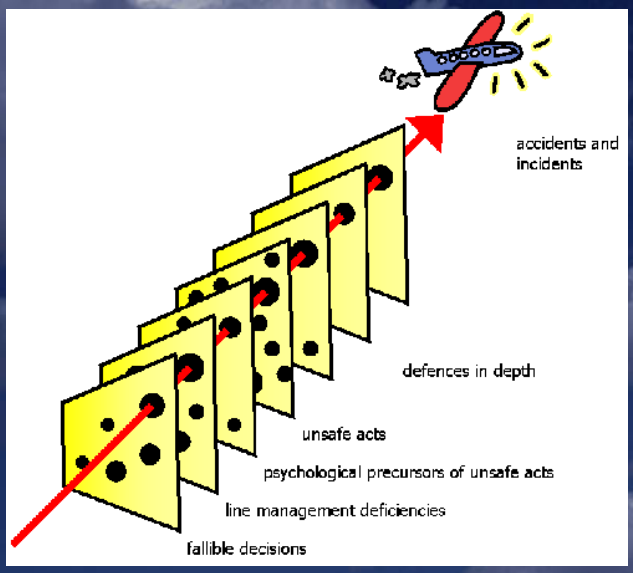

Seen that before somewhere!!, yes, it can be found through Google.

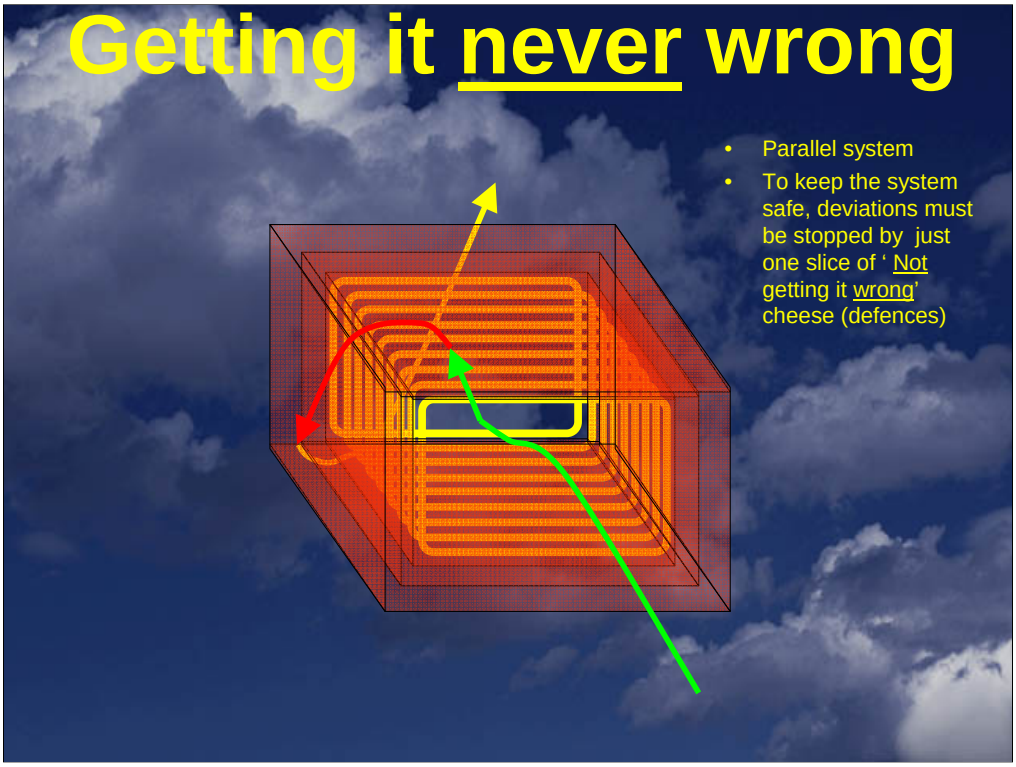

So where does this take us? What is the difference between series and parallel layers of cheese, that must, or must not be gone through?

Only one effective defence has to work, so this is a parallel system to return to certified safety parameters.

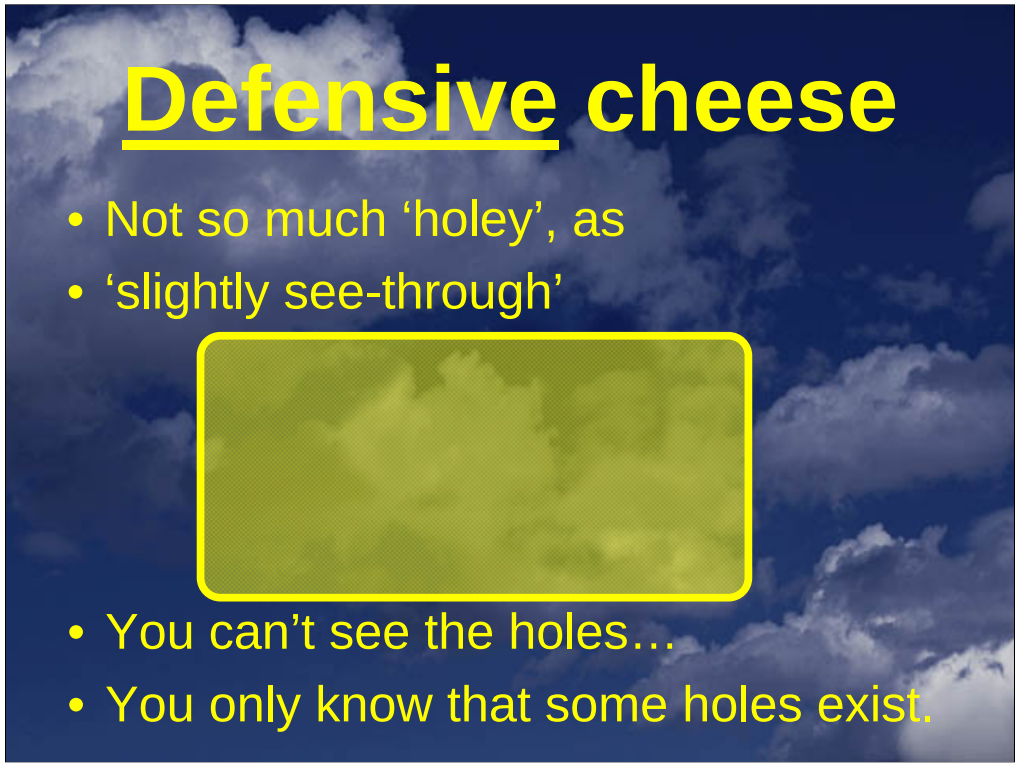

The holes are truly independent because they are not visible to the user, (only to the investigator).

If they WERE visible, then a properly functioning operating system would have engineered an avoidance or a closure of the hole. (Surely…)

You hope that any exceedence will be trapped and managed by a defence, but you can only hope in a well-managed HRO (is this overly contentious?)

This is why a well-managed HRO expends resources and authority in actively seeking not-yet-discovered holes.

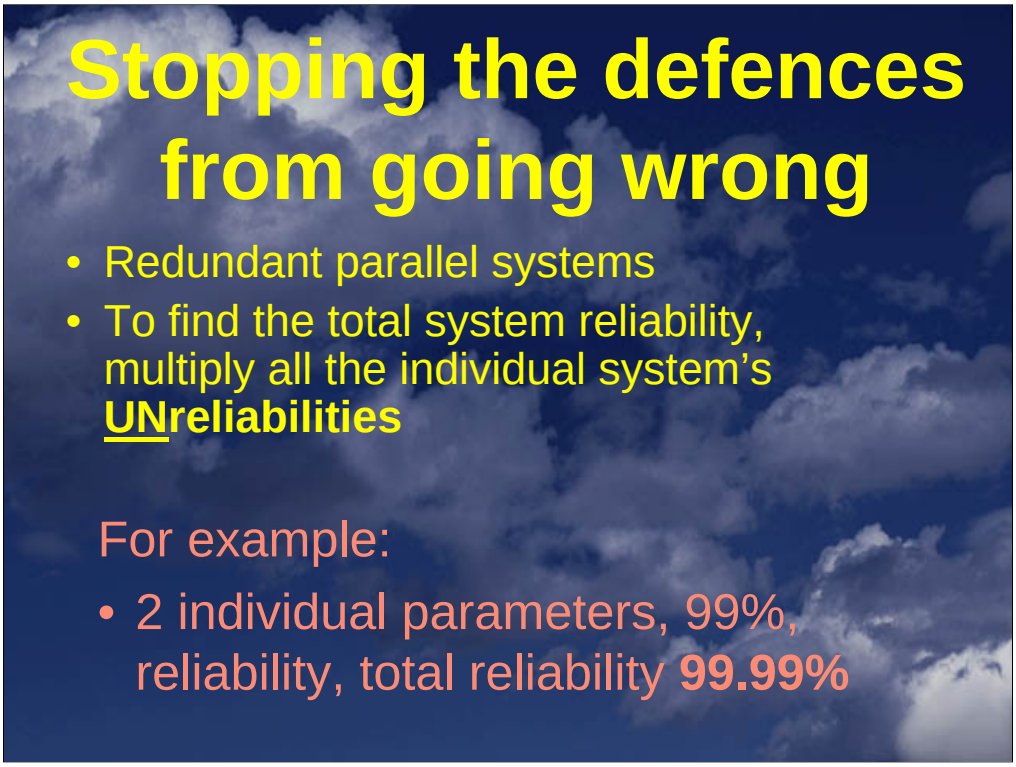

This is the same calculation principlee as for parallel systems used when operating within normal operating parameters.

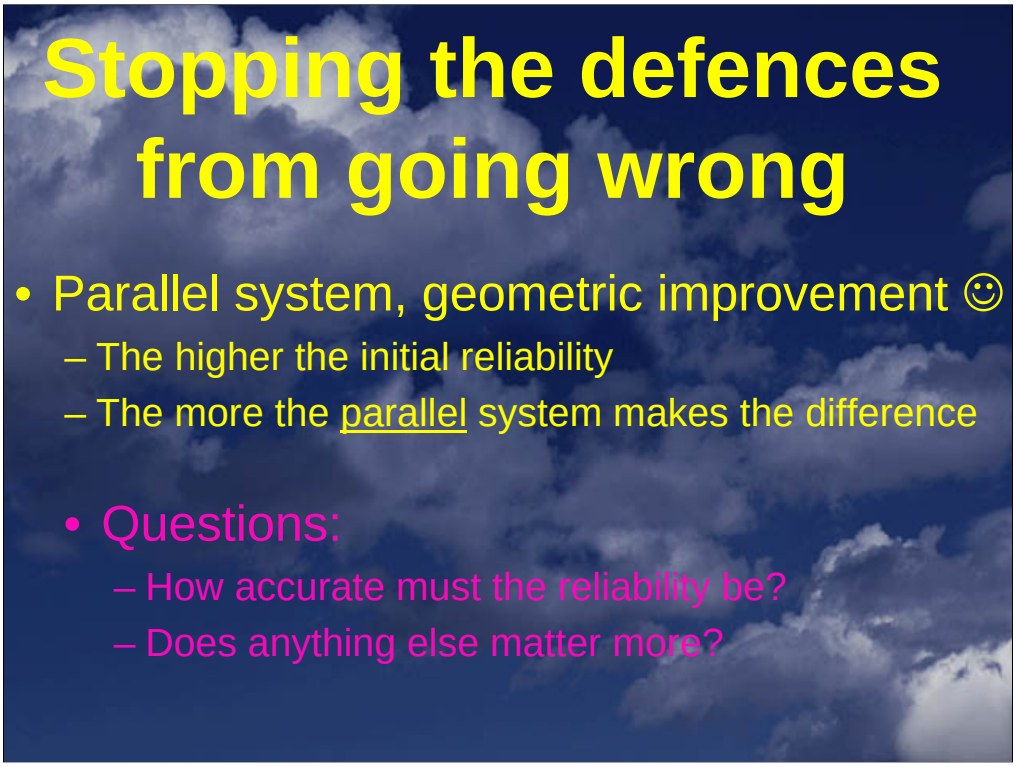

A good multiple defence system should only operate in parallel – that's the way they are designed.

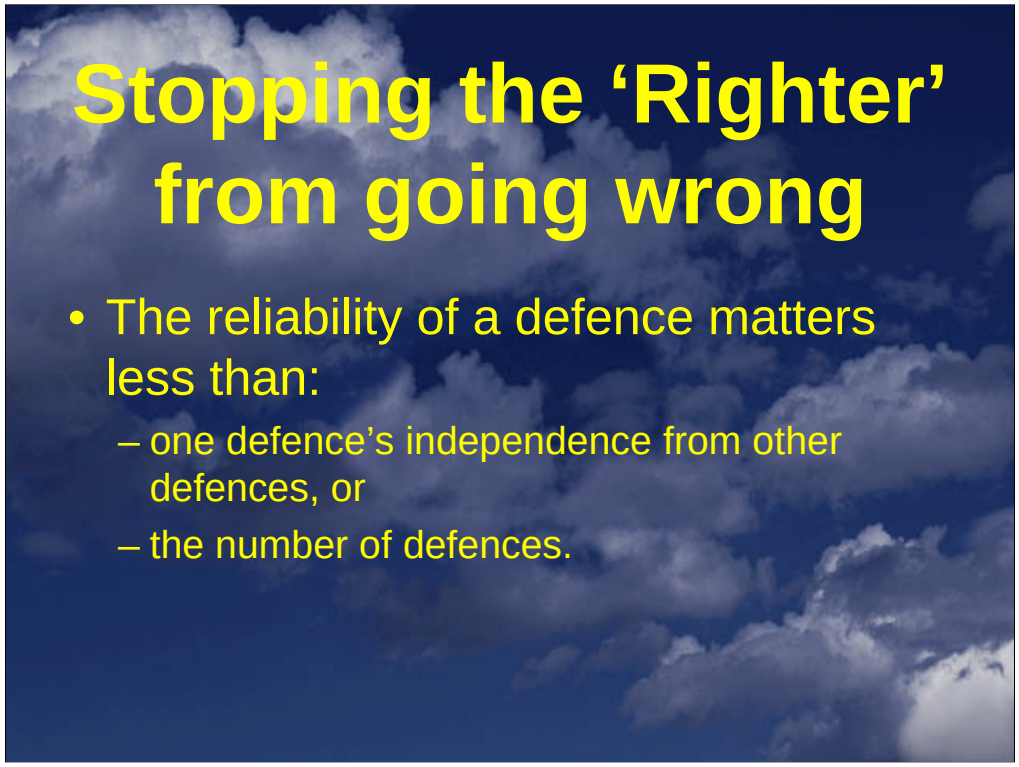

Independence:

For defences to work under a normal distribution and not a power-law distribution, they must be independent, not interdependent. The higher the reliability of the defensive system, the more this matters as you move further out along the tails of the different frequency distribution graphs. (See attached paper)

Number of defences:

In a parqllel system, the total unreliability is the product of all individual reliabilities. For example, if you add one more 90% effective system, you have increased the total reliability ten-fold (in a perfect world…)

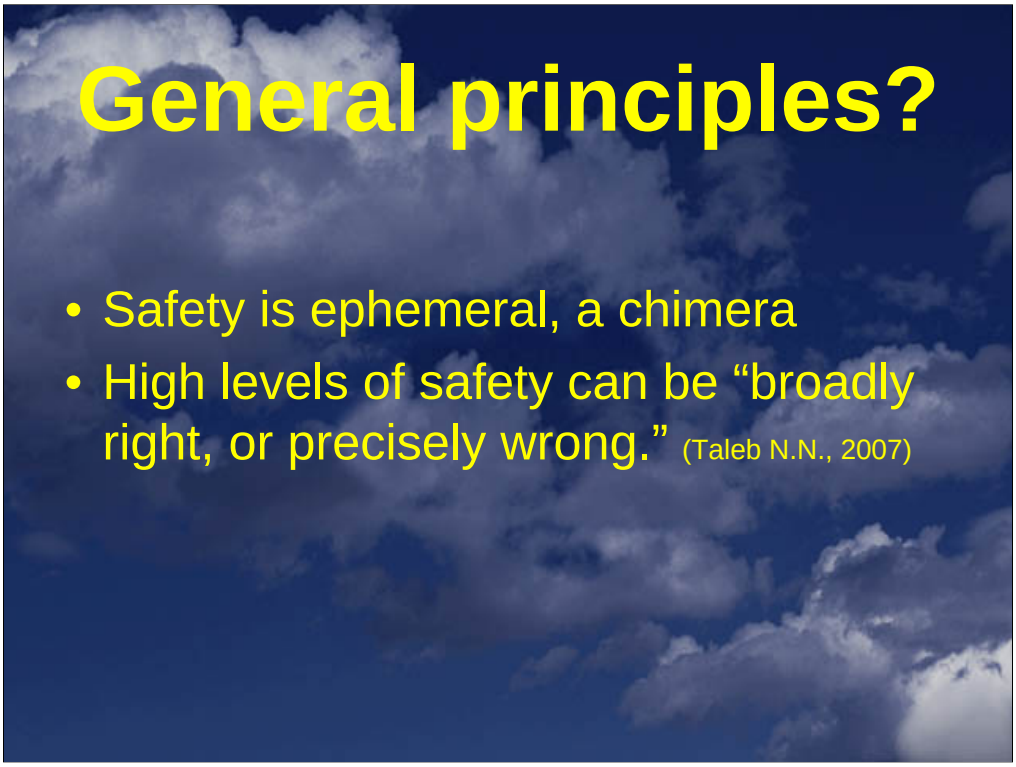

Definition of *chimera* (noun) impossible fancy; wild illusion. Examples Perhaps he saw a flying saucer, but perhaps it was only a *chimera*. **...**

•NOT (in Greek mythology) A fire-breathing female monster with a lion's head, a goat's body, and a serpent's tail. !!!

Definition of ephemeral: Lasting for a markedly brief time: "There remain some truths too *ephemeral* to be captured in the cold pages of a court **...**

Reason (1990 and 1997) said you cannot use accident frequency as a valid measurement tool in a HRO, as the numbers are statistically insignificant: he is so right. You cannot count danger measures because there are not enough of them, you must instead count (or manage) safety management measures. The statistical significance of good individual safety measures is nigh-on impossible to measure with any useful accuracy. However, the interrelationship between the individual measures (whether technical, organisational or individual), matters so much more than each measure in its own right in a good HRO.

It not the effectiveness of each measure that matters so much as what you do with it, in the context of the entire system’s safety management system.

Thank you

Mike Watson
+61 2 6161 6752
mike@tiffwat.com