Risks People Take and Games People Play

Peter Bernard Ladkin

University of Bielefeld and Causalis Limited

Bielefeld, Germany and London, UK

Abstract *In July 2014, a commercial transport aircraft, Malaysia Airlines Flight 17, in cruise flight over Ukraine, had its flight abruptly terminated through* "impacts from a large number of high-energy objects from outside the aircraft"*, The suspicion is that it was shot down. Three other commercial aircraft on international flights were in the same control sector at the time; other airlines had chosen to avoid the area. I argue that the kind of risk analysis one must perform to assess such possible security threats cannot be of the IEC 61508 type. I propose Meta-Game Theoretic Analysis, MGTA.*

1 What is Risk Assessment? An International Standard or Two

Safety assessment of critical systems in commercial aviation has been based for a long time on risk assessment. Since the late 1990's, the international standard for functional safety of electrotechnical systems IEC 61508 has also propagated an approach based on assessing risk (IEC 2010). Indeed, there is a general guide for electrotechnical standards which incorporate safety aspects, prepared by the Advisory Committee on Safety of the IEC, the international electrotechnical standardisation body, which requires that all such standards incorporate a risk assessment (ISO/IEC 2014).

A risk assessment according to the 2014 edition of Guide 51 (op. cit.) proceeds as follows:

1. You identify hazards;

   Loop:

2. You estimate risk;
3. You evaluate risk;

4. You reduce risk where intolerable;

   until <residual risk is tolerable>

5. You validate and document your reasoning along with the evidence.

In evaluating whether residual risk is tolerable, a nod is given to ALARP, and to social conventions concerning tolerability as well as other factors.

The use of technical terms here is as follows. A *risk analysis* comprises Steps 1 and 2, and is said to be a *systematic use of available information to identify hazards and to estimate the risk*. *A risk assessment* is a risk analysis followed by a risk evaluation and comprises Steps 1-4 above.

It's worth saying a couple more words about the underlying technical vocabulary, because it coheres with that of IEC 61508, which is not at time of writing incorporated into the International Electrotechnical Vocabulary (IEC various) and it varies from other, more common and maybe more intuitive, vocabulary.

Harm is what you think it is. It used to be restricted to persons, but in the last decade or so has expanded to include damage to infrastructure and environment, other words almost any kind of loss. A *hazard* is a potential source of harm; a *hazardous event* is an event that can cause harm. What is meant here is that a hazard is a state, or an event, or a combination, from which harm may result, and a hazardous event is something that happens which may, but must not, result in harm, and the harm, if any, resulting from a hazardous event is variable. So a hazard can be a sharp bend in the road; or a sharp bend in the road without a speed restriction; or a sharp bend in the road without a speed restriction and a car coming towards it faster than it can negotiate the corner. A hazardous event can be a car coming towards the sharp bend faster than it can negotiate the corner (but presumably not if this is already considered part of the hazard); or it can be the car coming off the corner and hitting the wall. That may not result in harm if everyone is belted in and the airbags deploy; equally it will result in harm if neither is the case. And the harm that results is dependent on the speed of collision as well as other factors.

It is important to note that there is a choice of what to construe as a hazard. Such a choice is amongst other factors practically bound up with the possibilities for prophylaxis. A hazard identified earlier in a possible accident progression, and then avoided or mitigated, may be easier to document and handle. But such early intervention may exclude certain system behaviors that would have been OK, and one would thereby have taken unnecessary action. Leaving the identification of a hazard to later in a possible accident sequence, when it becomes clearer that something bad is about to happen, may avoid unnecessary earlier intervention, but may also require a more resource-intensive reaction to avoid or mitigate an accident.

There is much in this vocabulary to quibble with; my preferred vocabulary is published elsewhere (Ladkin 2008). But most of the necessary concepts are here somehow.

It is not defined what a risk estimation is, but *risk* is a *combination of the probability of occurrence of harm and the severity of that harm*; it is not said how they are to be combined. One option is that of de Moivre: "*The Risk of losing any sum is the reverse of Expectation; and the true measure of it is, the product of the Sum multiplied by the Probability of the Loss*" (de Moivre 1711), in modern terms the expected value of loss. Multiplication of risks associated with individual hazards, followed by a sum over all individual hazards, is a common way of deriving that expectation. A problem is that the enumerated hazards might not be probabilistically independent, but we shall let that be.

So a risk estimation is an estimate of risk. According to the de Moivre model, that would be an estimate of the expected value of harm. And if you have another combinator in mind, an estimate of the value of that combinator.

2 The Central Role of Probability

Notice the dependence of all this on notions of probability. You will need some theory about probability to fill all this out. The notion of probability has itself a wide variety of interpretations. Good explanations of the varying conceptions may be found in (Hacking 2001).We shall consider three.

There is the Laplacian interpretation, in which a probability is physically inherent in objects. A fair die, because of its careful construction, has an inherent probability of one-sixth of landing with any given face showing. The word "has" is possessive and here exactly right: the probability is a property of the die. A biased die has different probabilities for some faces; say a slightly-less-than-one-sixth probability of landing with 6 showing and a slightly-more-than-one-sixth chance of landing with 1 showing.

Then there is the frequentist interpretation, associated with Jerzy Neyman. Probability is associated with events, and is a statement of how often a specific type of event occurs. How frequently your bicycle tire punctures, say. If you go out on a ride and estimate the probability of a puncture as one in four, or one-quarter, depending on how you present probabilities, you are saying according to this interpretation that when you do a lot of these specified kinds of rides, ceteris paribus you'd experience a puncture on about a quarter of them.

The third kind of interpretation is the Bayesian, or subjectivist, interpretation, associated with de Finetti, Savage, and in Britain especially D.V. Lindley, after the Reverend Thomas Bayes and his theorem. This says that a probability is a statement of a degree of rational belief. Here, the word "rational" is normative: one is expected to form a belief on account of reasons and evidence, and update that estimate as evidence becomes available. You've seen one black swan and one white swan in your life. You know (somehow; by authority, or by painstaking genetic analysis) that a swan must be white or black and not both and not vaguely neither, so it is certain that any given swan is white or is black. You rationally assign the

probability of a swan being white as one-half; identical with the probability of a swan being black, based on your experience to date. Then you see a lot more white swans, and no more black swans. For each swan you see, you update your estimates of the probabilities of whiteness or blackness according to Bayes's rule (a process called Bayesian updating), subject to the a priori constraint that it is certain that any given swan is white or is black and not both.

For events which are repeatable and frequent, there are theorems of probability theory which entail that all these conceptions come up with more or less the same values of probability for classes of such events. However, people building safety-critical systems are concerned with harmful events, or more precisely harm-loaded events (those events in which it is happenstance, or independent of the event itself, whether harm is caused or not, such as the car hitting the wall at speed). And such events are neither desirably frequent nor desirably repeatable.

For civil transport aircraft, one speaks not of hazardous events and their consequences, like the IEC, but rather of events resulting in specific effects. Extremely improbable effects are those unlikely to arise in the life of the fleet (all aircraft of a given type); extremely remote effects maybe once or so; remote effects maybe once per aircraft life (and many times in the life of the fleet). The certification regulations require is that a single failure that results in a catastrophic effect must be extremely improbable. Certification requires the constructor to show that this is so. Other effect severities are hazardous, major and minor (not that this notion of "hazardous" is different from that in IEC 61508). A classic introduction to these conceptions is (Lloyd and Tye, 1982).

A Laplacian interpretation applying to, say, the wing of a modern airliner may be plausible, as follows. The structures are designed to have it break under a specific load distribution at just over "ultimate load", which is defined to be 1.5 times "limit load", which itself is a number fixed at design time and which is purported to represent the highest loads to which the structure could be subjected during anticipated operations. And wings do so break at or above "ultimate load", during the required destructive test. They are engineered to withstand the required load, but no more, and this seems to be well achieved. Then the wing (rather, its intact successors) goes on to fly in uncontrolled but moderately well understood aerial environments, which can be argued to have probabilistic aspects.

So the wing is like the die; the engineering structure is well understood, as are the general characteristics of a throw, respectively of the weather, but the precise characteristics – the actual motion of the hand during the throw; respectively the precise behavior of the atmosphere during the flight – remain not sufficiently determined to render a deterministic calculation plausible. But notice here the justification in terms of what is known. The Bayesian approach takes the phenomenon of known information more rigorously, and arguably leads to a better intellectual fit.

A frequentist interpretation of wings breaking seems nowadays inapplicable, even implausible – wings just don't break in commercial service (any more), just like the regulation requires them not to. So the frequency is zero. (There are ex-

ceptions to this, but not in commercial transport.) But suppose one were to break, sometime. Then what's the frequency? One in ... what? Were the ceteris paribus conditions satisfied on that one occasion? Or were there particular conditions? How do you decide whether causal conditions have a probabilistic nature or a particular, exceptional nature?

It seems we are best advised to be Laplacians or Bayesians. But the Laplacian construal is nowadays "denigrated", so I guess we would have to be Bayesians.

But say you are inspecting a newly-built homebuilt aircraft for airworthiness. You can't see any of the composite lay-up of the wings – it's all hidden. So you interview the owner and form a rational belief about hisher construction capabilities and the care taken. It all looks good; you declare the aircraft airworthy. The owner goes up on a test flight and promptly a wing breaks off. You calmly update your estimate as the Reverend Bayes said you should…

Surely, given that the design is in order, the chance of the wing breaking as it did depends, not on your beliefs, but on how the wing was built, objectively? The owner didn't take as much care as heshe said during building; heshe screwed up badly in one place and didn't realise it. It seems we're back to Laplace: it's the airplane that has been built well or badly and the – what shall I call it? - propensity to break, the greater or lesser chance of breakage, is inherent in the structure. It seems that the construction and its thereby inherent propensities to fail matter rather more concretely than an assessor's beliefs.

3 The Way It Is Done in Aeroplane Certification

The acceptable means of showing compliance with aviation regulations are codified and formulated explicitly by the main airworthiness certification agencies, the US FAA and the European EASA. FAA rules are in 14 CFR Part 25 (United States Government, various dates). EASA rules may be found in the EASA Certification Specification CS-25 (European Aviation Certification Authority, various). They specify what is called in other contexts a risk matrix, a discretisation of effects against occurrence likelihood:

- Catastrophic effects must be extremely improbable
- (EASA) Hazardous effects must be extremely remote and major effects remote; or
  (FAA) major effects must be remote/improbable
- Minor effects may be probable, or even frequent.

Nowadays, a specific numerical probability per flight hour is associated with the qualitative probabilities.

But in fact what mostly happens is something rather different. Going back to the wing, recall that it must withstand ultimate load, defined to be limit load times

1.5, where limit load is an estimate of the highest loads to be plausibly experienced in service. A wing is built, and loaded until it breaks. And that should occur at equal to or higher than ultimate load. It is assumed (and checked and controlled!) that manufacturing-process quality, along with timely (checked and controlled!) in-service replacement of life-limited parts, ensures that all wings are interchangeable in terms of load withstood. That has everything to do with engineering and control and nothing at all to do with probability. You believe that it won't break because you built it that way and have enough experience to know that that suffices. And you test that understanding precisely once. (Actually, it turns out on a recent certification it was acceptable to have the wing break at slightly below ultimate load, then perform a redesign and show by means of extensive computer simulations that the strength of the wing had thereby been increased to withstand ultimate load, without destructively testing the redesign.) All this is taken to show that the possibility of a wing failing to fulfil its function in flight is extremely improbable; that is, it won't happen during the fleet lifetime, as far as anyone can tell. Note that there is no intellectual connection here with probabilistic criteria per se. Engineering design, simulation and deterministic test is deemed satisfactory to fulfil a criterion, itself expressed but not enforced in terms of likelihood.

Perceptive readers will note I have glossed over some of the subtleties in airworthiness certification, but I believe the story as I have told it suffices for my purpose here. In short, the notion of probability or likelihood is problematic when referring to very rare events. When possible in aerospace, we far prefer to have designs which we can plausibly argue on the basis of design and construction will withstand all occurrences of adverse events in their lifetimes.

Except of course when some other people have designed an object which is intended to cause your structure to fail, and is built according to similar principles as above to execute that function. Which we now consider.

4 Risk of a Different Variety: Security Risk

On 17 July 2014, a Boeing 777 operating as Malaysian Airlines Flight 17 between Amsterdam and Kuala Lumpur was destroyed in and over Eastern Ukraine. Witness reports and the fact that the wreckage was strewn over a very large area point unequivocally to in-flight disintegration. *"Damage observed on the forward fuselage and cockpit section of the aircraft appears to indicate that there were impacts from a large number of high-energy objects from outside the aircraft"* (Dutch Safety Board 2014). An admirably careful statement. Put another way, pieces of wreckage photographed by reliable observers show damage such as caused by shrapnel from the detonation of an explosive projectile with a proximity fuse. The Report also says there were no indications of any problems or malfunctions before the abrupt end of recording on the data recorders (op. cit. Section 3, Summary of Findings). In other words, it is almost certain that somebody shot the

flight down. There was and is an armed insurrection occurring in the area, with fighting between sovereign Ukrainian forces and heavily-armed "rebels" who appeared to be led by Russian citizens.

The Incident Aircraft, Boeing 777-200 9M-MRD
Photo by Alan Wilson
Licensed under Creative Commons

Ukraine is sovereign over the airspace in which MH 17 was flying. Many airlines had been using the airway, L980, and adjacent airways. Indeed, when destroyed, MH 17 was at Flight Level (FL) 330, a nominal 33,000 ft above mean sea level in a "standard atmosphere", in Section 4 of the CTA (Control Area) Dnipropetrovsk (known to aviators as Dnipro Control). In the same sector at that time were a same-direction Boeing 777 at FL 330 about 100km southwest on airway M70 heading towards waypoint PW, a same-direction Boeing 777 at FL 350 about 30km northwest, and an opposite-direction A330 at FL 400, 50km east-north-east on airway A102. (op. cit., Figure 2, p12). (Note: A report in the weekly journal Aviation Week and Space Technology from a week or two after the accident had MH 17 14 nautical miles or so in trail of a Singapore Airlines aircraft at FL 350, and about 8 nautical miles abeam of an opposite-direction Air India aircraft on another airway. (Schofield et al., 2014). The divergences between the two reports show again how difficult it is to establish facts about such events, even though the relevant information is ostensibly readily available from multiple sources.

At the time, there was a Temporary Restricted Area from the surface to FL 260, valid from July 1 through July 28. The existence of this area was distributed by NOTAM (Notice to Airmen, the international standard informational service). On 14 July, a further TRA existed from FL 260 up to FL 320, valid until 14 August,

covering the eastern part of the area covered by the first TRA. All the flights passing through Sector 4 of Dnipro Control were conforming with both NOTAMs, as indeed to be expected with commercial flights under positive control.

Some airlines had previously performed a "risk analysis" and had been avoiding overflying the area, such as, it was reported, Qantas and BA. Other airlines avoided the area afterwards.

A Sister Aircraft in Flight
Photo by neuwieser
Licensed under Creative Commons

MH 17 had filed a flight plan with requested FL 350 in the area, but when in contact with Dnipro Control at FL 330 was unable to transition to FL 350 and continued on FL 330 (op. cit.).

5 Security Risk Analysis: What's With Probability?

What kind of risk analysis can it have been which had been performed by those airlines avoiding the area? Could it have been one as described above? Let us try:

- Identify the hazards:
  - Getting shot down by a ground-based missile
  - Getting shot down by another aircraft
  - Getting shot down by ground-based artillery or flak

- Severity of all these events is the same: catastrophic, everyone on board dead, hull loss, damage on the ground, perhaps harm to people on the ground

- Estimate the risk: as defined, "combine" probability of each hazard with severity. So what is the probability of each hazard?

- What is the probability of getting shot down by a ground-based missile? Zero if there aren't any in the area with the capability of reaching a target at FL 330. Someone explained to a journal that the commercial-aviation industry relied on sovereign militaries to control their assets -- does that mean zero probability if the only such missiles in the area are maintained by sovereign militaries? Well, not quite. Siberian Airlines (Sibir) Flight 1812 was shot down from FL 360 on 4 October 2001 over the Black Sea on its way to Novosibirsk from Tel Aviv (Aviation Safety Network, no date). The aggressive object was a missile operated by the Ukrainian military during military exercises, which locked on to the airliner rather than its intended target. OK; so the chance is not zero. What is, then, the probability? One in ... what? Can one possibly tell? What are the ceteris paribus conditions that say "a Flight 1812-type incident could occur here"?

- What is the probability of getting shot down by another aircraft? Ukrainian military aggressor aircraft, specifically Su-25 Frogfoots, use the airspace. But Frogfoots have an effective service ceiling some 10,000 ft lower and as far as we know can't "shoot up" (see, for example, (Sweetman, 2014), or details in (Locklin, 2014)). Besides, why would such an aircraft try such a thing? There are no "rebels" up there at FL 330. A Russian or Ukrainian fighter aircraft could be up there; indeed there were previous unconfirmed reports of unauthorised Ukrainian-airspace intrusions by Russian military aircraft. But what would aircraft under strict sovereign-state control possibly be doing up there shooting at traffic at FL 330? As far as anyone has seen or said so far, there were no such aircraft up there at FL 330 in Ukrainian airspace anywhere in the neighbourhood.

The precursor state to Russia, the Soviet Union, had shot down civilian airliners. The first was a Korean airliner violating Russian airspace, which refused an interception using internationally-recognised manoeuvres and was consequently shot at by an interceptor, on 20 April 1978 (Aviation Safety Network, no date). The fire killed two people. The aircraft was not destroyed, but landed relatively safely off-airport on a frozen lake. The second was also a Korean airliner, a Boeing 747 which had also violated Russian airspace, crucially in the neighborhood of and around the time of an important missile test. The aircraft was shot down by an

interceptor who had mistaken it for a US military intruder, a reconnaissance aircraft also built by Boeing, of Boeing 707 size, and believed it was manoeuvring to avoid interception. That was on 1 September, 1983 (Aviation Safety Network, no date).

An Su-25 Frogfoot Aircraft
Photo by Rob Schleiffert
Licensed under Creative Commons

However, ceteris paribus conditions are nowhere near satisfied. Neither of the shot-down airliners was on or indeed near internationally-recognised civil airways for which it had a clearance. Both were formally intercepted using internationally accepted protocols. One airliner refused the interception; the other airliner was honestly judged to be actively avoiding one on a dark and somewhat cloudy night. Both incidents occurred during the "Cold War", during which the Soviet Union was on one side and South Korea, considered by the Soviets as something of a protégé of the United States, definitively on the other. The Soviet Union believed itself, with reason, to be at times actively intruded upon, sometimes by civilian assets performing military tasks under subterfuge. (And indeed vice versa.)

In stark contrast with these circumstances, MH 17 was following a recognised airway at a cleared Flight Level on a filed flight plan and was not violating, or about to violate, anyone's sovereign airspace without clearance. Neither is it plausible to imagine it was trying to perform military tasks by subterfuge. Neither was Malaysia on one side of a "Cold War" with Russia on the other.

At time of writing, Russia has in fact claimed that MH 17 was shot down by a Ukrainian Frogfoot. Russia has published radar data they claim is proof, which has been assessed by reliable third parties who are less than convinced by it. The United States claims to have proof that MH 17 was shot down by a surface-to-air

missile launched from Eastern Ukraine. The United States is known to have assets which could establish this beyond reasonable doubt, but at time of writing this information has not been published and independently verified.

A Buk-M1-2 Launcher
Photo by .:Ajvol in the public domain

A Complete Buk-M1-2 System, Comprising Multiple Vehicles
Photo by Vitali V. Kuzmin in the public domain

On 20th October, 2014, I discovered through my local newspaper that the head of the German Federal Intelligence Service, BND (Bundesnachrichtendienst) told its parliamentary oversight committee on October 8 that MH 17 had been shot down by separatists using a Buk system which they had obtained through plundering a Ukrainian military base. It is said that convincing evidence was presented (Gude and Schmid 2014).

- What is the probability of getting shot down by ground-based artillery? Nobody thinks that anyone has any artillery assets in the area that can reach up to FL 330. Even if there were, people estimate chances of getting a ballistic hit at close to zero. Ballistic projectiles are intended for buildings and very slow-moving objects such as battleships, not for high-performance aircraft.

- What is the probability of getting shot down by flak? Up there at FL 330, almost zero. Besides, as far as anybody knows there are no flak delivery assets in the area.

So where is here the probability value? As far as I can see, and I am suggesting as far as the reader can see also, there isn't one. A Guide 51-type or IEC 61508-type risk analysis is not what is being performed when analysing such risks.

6 What's Really Going On

So what reasoning is being used here? I have just performed something like the following:

1. It is observed that hostile military engagements are taking place in the area.
2. The area in which those engagements are taking place, or to which they could plausibly spread, is circumscribed.
3. A hoped-complete list of hazardous events occurring through hostile military acts to commercial aviation flying in open civil airspace is enumerated.
4. Scenarios leading to those hazardous events are constructed.
5. The plausibility of each scenario is assessed.
6. Plausibilities are ranked. First, plausible-implausible. Then, more plausible-less plausible.
7. A discrete decision is made based on those plausibilities: use the airspace/don't use the airspace.

Up to Step 3, that is what the IEC documents on engineering risk say to do under hazard identification. But then the method diverges. Scenarios are not neces-

sarily considered in IEC methodology. Some may consider Fault Tree Analysis followed by Event Tree Analysis to be a form of scenario construction, but I suggest that the current type of scenario construction is significantly more detailed than what occurs in a typical FTA/ETA. One could, perhaps, consider a qualitative fault tree as some kind of enumeration of scenarios, or at least as a structure which yields such an enumeration. But the scenarios considered here are not possible causes hierarchically ordered in subsystems, as in an FTA. Neither are they abstract possible futures as in an ETA. They are temporal scenarios with actors performing actions according to motivations and reasons and other human characteristics. Then, some decision is made on the basis of that analysis: do or don't.

What is most important about that decision is that it is the Right One: don't fly there if somebody's maybe going to get shot at in any place where you are going to be.

In a probability-based analysis, one could make all the rational decisions based on probabilities and still get stung on your first outing. Your analysis is valid according to the IEC conception. You took a risk and then you lost the bet. So go ahead, do it again! Your analysis is still valid. Toss the die!

Contrast this with commercial transport aircraft certification. The rules say: your airplane will do this-and-this. And furthermore the evidence will be documented. The evidence deemed acceptable may be probabilistic and is retained and available. So rational decisions were made based on evidence couched in terms of probabilities, as in the IEC approach. Say you go out and get stung on your first outing. The judgement is different: your airplane is not airworthy; make it airworthy and you can go fly it again (this is accomplished by means of instruments called Airworthiness Directives, which are remedies mandated for all operators of the aircraft type to restore and/or maintain airworthiness of their aircraft. If you don't fulfil an AD, your aircraft is not airworthy and may not be flown.) This outcome is different from the IEC outcome of a critical failure. You can't just go ahead and do it again; you must remedy.

The current airspace-use situation we are considering is comparable with the aircraft airworthiness procedures in that immediate remedy is required: the airspace is closed to civil traffic, and even if it weren't it would be doubtful if anyone would be using it. But it diverges in that it is called a risk analysis; aircraft certification is not called "risk analysis" by anyone, and the process is not treated as if it were. Testing a wing to destruction is not analysing risks; it is assuring ourselves that the engineering is sound and a wing will not break in service because it is functionally identical (through process and quality control) to the successful-test object. And, conversely, a decision to use airspace is not called "traversal-worthiness certification" and neither will it be.

In truth, the probabilistic risk of getting shot down over Eastern Ukraine was low, even under a reasoned belief that there were high-altitude surface-to-air missiles (SAMs) in the hands of unreliable combatants. Troop and equipment movements had been seen at the weekend, 12-13 July (but it is not known at time of writing what the contemporary analysis had concluded), and a Ukrainian military

transport had been shot down at FL 210 on Monday 14 July. Hundreds of airplanes had flown the routes over Eastern Ukraine in the meantime; four were flying it at the time of shootdown. And only one of those aircraft was shot down. An $O(10^{-2})$ risk is high compared with other estimates of risks in aviation, but in objective terms one might question whether such an event is likely.

It is also plausible to think that right after MH 17 was shot down, if it was shot down by a SAM then the chances another aircraft would be shot down in the region had plummeted to near zero.

That conclusion is also based on scenario analysis. Such assets were widely assumed to belong to the Russian military. It is true that "rebels" had boasted of capturing some Ukrainian Buk SAMs in June but this had remained unverified and it would have been unlikely they could operate them effectively without having some sort of rudimentary training which would not have been available. So if Buk SAMs were available to rebels, it is likely they would have been Russian assets and thus recommandeered immediately after the shootdown for many reasons; and it is not regarded as plausible that Russian military assets under direct Russian control, as re-commandeered devices would have been, would be used to shoot down civil aircraft. (But contrast this reasoning with the reported claim by the Head of the BND in camera, noted above.)

It is not regarded as plausible, but it could happen. Some odd soldier of almost any army could get drunk or suicidal or both, and fantasise about going out in a blaze of notoriety, like the 9/11 terrorists. And succeed, as two out of four cohorts of the 9/11 perpetrators did. This possibility appears not to be sufficient reason for any of the world's airlines to avoid Russian airspace. Neither did the shootdown of Siberian 1812 in 2001 cause Russian or any other airlines to avoid Ukrainian airspace; a repeat was not regarded as plausible.

Why not? I believe it has to do with people, motivational and goal analysis, and assessments of capabilities at fulfilling goals. Put crudely, the Soviet Union had screwed up badly with KAL 007 in 1983; that was never going to happen again. Ukraine had screwed up badly with Siberian 1812 in 2001; that was never going to happen again. Controls were already in place and must be followed more precisely. Whereas two of four cohorts of 9/11 aggressors had achieved what appeared to be explicit goals. Few controls were in place and it was unknown whether others with similar goals were still out there. World civil air traffic stopped, and restarted slowly with considerably more assessment and control, including previously unthinkable measures such as giving the USAF rules of engagement to shoot down civilian transport aircraft.

When we are in the realm of personal and organisational goals, motivations, means and so forth, we are no longer in the realm of probabilistic assessment. Probabilistic assessment is based ultimately on a notion of a random variable and while goal, motivational and strategic analysis may rely somewhat on uncertainties, as in "taking a chance", it hardly relies on any notion of randomness. "Purposeful" behavior is indeed the contrary of "random" behavior.

Consider another example, from a different realm. The chances that your WWW server suffers a surfeit of incompletely-formed TCP handshaking packets are low; but they are very high to almost certain if your server is the target of a DDoS attack. The difference between the two situations is not probabilistic, or generally in any way related to chance, but rather concerned with some specific agent's purpose and means at that point in time. Analysis is concerned, not with bursty behavior on communications networks, but with whether there is an agent who had reason and means to elicit the behavior and why. Far from being a probabilistic random variable, it is more like an almost-Boolean environmental variable: are you currently subject to DDoS attack, or not?

Furthermore, there are no uniform assumptions one may make about chances in the background. If your civil aircraft has been subject to rocket attack in Eastern Ukraine, I have just argued that it is very unlikely you or anyone will be subject to further attack. Whereas if you have just survived a DDoS attack, the chances rise that you will be subject to another one soon. Or, to compare like with like, a ManPAD attack on a civilian cargo aircraft deploying around Bagram Air Base in Iraq might be seen to increase the chances that another such aircraft will be so attacked soon. The differences are not to be found in any quasi-objective analysis of inanimate situations; they are to be found in the goals, motivations and means of some of the players. (Or may be all of them. A second ManPAD attack may be canonically thwarted by grounding and guarding all aircraft, as happened for similar reasons immediately after 9/11. The goals, means and motivations of all the participants should likely be considered.)

7 Meta Game Theoretic Analysis (MGTA)

We are in the realm of game theory. Indeed, the classical game theory of non-cooperative games (so-called "game theory" is usually the study of non-cooperative games, contrasted with cooperative game theory, or coordination games, studied by the philosopher David Lewis (Lewis, 1969) as well as the political theorist Thomas Schelling (Schelling, 1960)). But this is not pure game theory, as studied by economists. It is more like a meta-theory of games. First there are methods to choose a game from amongst a variety of possibilities (the "meta" part), then follow methods to choose actions within the chosen game (game theory proper).

There are situation variables, which in some sense set the game being played. Am I currently subject to a DDoS attack? If not, I am administering a server in an unreliable bursty environment and there are lots of things I can choose to try to ameliorate the situation compatible with my goals. If so, then nothing I do for a while will change the environment and I have only two actions available to me: let my server be overwhelmed and clean up whatever mess results; or disconnect my server from its channels (most likely is that there are only a few channels to which I am connected). Are there currently high-altitude SAMs or high-altitude ag-

gressor aircraft available to unreliable players engaged in hostilities in Eastern Ukraine? If yes, my airliner might be shot at/down and I have to think of what I do. If the answer is no, the high-altitude airspace is just like any other airspace anywhere else in the world; free from worry (if I have reliable collision avoidance!). In that case I am in the null or trivial game: payoff is the same whatever I do and whatever the "opposition" does and is exactly the unit value.

In this first step, chances may reappear, as do the dilemmas associated with the interpretation of probabilities which I considered earlier. What is important for my decision making is what I know or have reason to believe. How likely do I think it is that unreliable players in Ukrainian hostilities have SAMs? Say I think there is an 60% chance. It then follows that there is a 40% chance I am playing the null game and a 60% chance I am playing another, more complicated game. Or so I reckon. Whereas the reality is either that unreliable players have SAMs, in which case I am truly in the complicated game and would do best by deciding my actions according to that; or that unreliable players do not have SAMs, in which case I am truly in the null game and can return to my Sudoku without further ado. Thus this reckoning of chance is an assessment of my uncertainty. We are unequivocally in the realm of Bayesian probability.

Can I collapse this twofold structure into a single structure, say Decision Theory? I don't believe so, for the reasons in the last paragraph. Best is to know what game you are actually, objectively in, and to choose your actions according to that game. In the case of a DDoS attack, I know and can choose. In the case of unreliable players with or without SAMs, I don't know, but it were best if I did. If the reality is the null game, I am optimally well off by deciding this correctly. And as MH 17 shows (if the most plausible scenario at time of writing is correct), I am not necessarily well off by deciding this incorrectly.

Can I, then, assume one game or the other? The perils of assuming the null game (no SAMs) are by now obvious. What of the other choice, that in the absence of knowledge I assume the "worst" game from amongst the possibilities? People have done the work for us – if we are to assume the worst and avoid all areas of hostilities in which the weaponry is not publicly known, I would have to get to anywhere east of Kiev more or less by flying around the Cape of Good Hope (see, for example, the graphic (Times 2014)). If an individual airline were to choose to follow this option, I would lose custom and fold quickly. If the world's airlines were collectively to choose to follow this option, then international business would instantly suffer a step change for the worse: all personal dealing suddenly becomes far more expensive in both time and money and costs of international business suddenly rise. Except for suppliers of aviation fuel, who are laughing all the way to the bank. Neither of those seem particularly attractive, let alone ideal, options.

If I am in the situation, though, in which I have two choices of game and one of those happens to be the null game, then there is a way I can "play both". The null game says always: do nothing, and ye shall neither suffer nor gain. So I can play the other, non-null game but weight my payoffs by my assessed chances that I am

in that game. Then I can consider that I am somehow in both games simultaneously: I am in the one game to 60%, by obtaining a 60% payoff for my actions, and also in the null game to 40%, by obtaining 40% unit payoff.

So I can solve – let us call it - the "Ukraine" problem simply by playing the "SAMs-yes" game to the weighted value of my belief that SAMs-yes.

But consider the following situation. On the ground, there are two combatants and one SAM base whose operators are effectively commanded by whichever combatant has control at any one time (the operators behave neutrally in order to, they hope, "save their skins"); and control changes hands regularly and, let us suppose, evenly (50% each). Suppose the one combatant dislikes airlines whose names start with A, C, E, G, I… and the other combatant dislikes airlines whose names start with B, D, F, H, J… As an airline, I am either in the null game or the firing line. I am pretty much forced to choose my game; if I choose a weighted average then half my planes are shot down and I lose custom and fold (not to speak of the distress caused my shot-down passengers and their relatives, for which I and my insurers are also liable). In this case it seems I cannot avoid choosing my game explicitly.

So in general we should expect that an explicit choice of game should be made. The weighted-belief approach works for the specific case in which one game is the null game, but not in general.

It follows that, in general, there are two separate, non-conjoinable steps to the analysis.

1. Choose your game;
2. Choose your actions in the selected game.

Airlines performing risk assessment on the "Ukraine" problem may have been able to use the risk analysis afforded when there are just two games, one of which is the null game, but in general this is not possible. Generally: choose your game; then choose your action.

It follows that security-risk analysis is fundamentally different from IEC-type safety-risk analysis. Game definition and choice, then action choice; respectively situation probability assessment.

8 MH 17 and Consequences in Light of MGTA

This construal of security-risk analysis already yields some results which run counter to the "prevailing wisdom". Many aviation professionals have been proposing that ICAO Do Something about the analysis of airspace usage risks apropos MH 17 (Schofield et al., op. cit.). First, there is arguably a misconstrual of ICAO's structure; second, there is a false expectation of effect, according to our analysis above.

First, to ICAO's responsibilities and capabilities. ICAO cannot Do Anything in the sense intended. ICAO is not a sovereign entity, it is a talking shop for sovereign entities in which things only happen if they are agreed amongst all sovereign members. Almost every nation belongs to ICAO. It advertises and propagates matters on which there is universal consensus. ICAO cannot issue recommendations not to fly over Ukraine (and recommendations are all it can issue) unless almost every member nation besides Ukraine decides it is not a good idea to fly over Ukraine (of course if Ukraine itself decides so, it may enforce the measure without any consultation). And if almost every member has so decided, then all airlines have already been informed of that advice by their sovereign and are following it for reasons of due diligence, not to speak of their insurance contracts. So such an ICAO recommendation would be a no-op; it would already be a done deal.

If people really yearn for an ICAO determination, this is no argument against that; let them have one by all means. It is only an argument that an ICAO determination would change nothing "on the ground" (that is, in the air).

More problematic is that such a determination would be insidious, in that it determines the game to be played. No matter what agreement might be reached in ICAO on a way to assign usage risk, there is a new game to be played based on the determined risk, and that new game is riskier for airlines.

For along with a determination of risk will come inevitably an assignment of responsibility. If there is a risk, say of 5%, of being shot down in Sovereign Airspace Q, then Q's sovereign, and/or the canonical risk assessors, and/or the airline which proceeded across Q's airspace according to the canonical risk assessment, will be held liable to some mathematical formula. This will happen because they are regarded as the pertinent actors in the regrettable decision to fly across – and thereby get shot down.

Let us suppose there is some "standard" assignment of airspace-usage risk. Say, as determined by ICAO for those entities who wish for this, but for the purposes of this argument determined by any means. Along with this assignment will come the liabilities associated with this assignment, as above.

Suppose further that Sovereigns A and B are at loggerheads. Sovereign B knows how Sovereign A calculates airspace usage risk; namely, according to the "standard". Sovereign B unattributably infiltrates A and brings down a commercial airplane, in order that A will be actively internationally criticised for screwing up the risk analysis, thus causing airlines to avoid A's airspace and diminishing that source of revenue for A, as well as possibly causing Sovereign A to pay compensation for the shootdown on the basis of the assignment of responsibility recounted above, which can nowadays run into ten-digit dollar sums.

That's a great win for Sovereign B at expense of Sovereign A. It would be appropriate to consider Sovereign A a victim (an undeserving loser) of that game. And the game can only be played if there is a "standard" risk assessment of airspace usage, as wished by those who want ICAO to establish a standard.

So, beware of what you wish for!

9 Some Other Examples

I have applied MGTA to the phenomenon of ATM phantom withdrawals from demand-deposit bank accounts, and how they are handled by customers and banks. I have also applied it to the phenomenon of "security theatre" with respect to implanted/implantable digital medical devices, in which security "researchers" graphically demonstrate "vulnerabilities" with devices, such as hackers on the street reprogramming a heart defibrillator or an implanted insulin pump remotely. There is no space here, though, to recount those studies.

Acknowledgements Thank you to Mike Parsons, for inviting me to give a Keynote talk at the 2015 SCSC Safety-Critical Systems Symposium and expressing the wish for me to say something non-trivial about MH 17. I hope I succeeded. I thank Harold Thimbleby and Ken Hoyme for initial reviews of this work (including the phantom-withdrawal and medical-implant examples). Ken chairs the AAMI committee on medical device security, and gave me extensive references to what is being done. He also provided incisive comments way beyond the usual scope of review; the paper has been thereby improved, and I am particularly grateful.

References

Aviation Safety Network (no date) Sibir 2001 incident, Entry http://aviation-safety.net/database/record.php?id=20011004-0 , accessed 2014-08-08.

Aviation Safety Network (no date) KAL 1978 incident, Entry http://aviation-safety.net/database/record.php?id=19780420-1 , accessed 2014-11-18.

Aviation Safety Network (no date) KAL 1983 incident, Entry http://aviation-safety.net/database/record.php?id=19830901-0 , accessed 2014-11-18.

de Moivre, A. (1711) de Mensura Sortis, Phil. Trans. Roy. Soc, 1711. Reprinted with commentary in (de Moivre and Hald, 1984).

de Moivre, A. and Hald, A. (1984) A. de Moivre: 'De Mensura Sortis' or 'On the Measurement of Chance', A. Hald, Abraham de Moivre, Bruce McClintock, International Statistical Review / Revue Internationale de Statistique 52(3):229-262, Dec.1984.

Dutch Safety Board (2014) Preliminary Report, Crash involving Malaysia Airlines Boeing 777-200 flight MH17, Hrabove, Ukraine, 17 July 2014, Dutch Safety Board, The Hague, September 2014, available from http://www.onderzoeksraad.nl/uploads/phase-docs/701/b3923acad0-ceprem-rapport-mh-17-en-interactief.pdf

European Aviation Certification Authority (various) EASA Certification Specification CS-25, the initial version of which is available at http://easa.europa.eu/system/files/dfu/decision_ED_2003_02_RM.pdf

Gude, H., and Schmid, F. (2014) Deadly Ukraine Crash: German Intelligence Claims Pro-Russian Separatists Downed MH17, der Spiegel on-line, October 19, 2014, http://www.spiegel.de/international/europe/german-intelligence-blames-pro-russian-separatists-for-mh17-downing-a-997972.html

Hacking, I. (2001) An Introduction to Probability and Inductive Logic, Cambridge University Press.)

IEC (2010) International Electrotechnical Commission IEC 61508:2010 Functional Safety of electrical/electronic/programmable electronic safety-related systems. Seven parts, available from http://www.iec.ch/functionalsafety/standards/page2.htm

IEC (various) International Electrotechnical Commission IEC 60050, International Electrotechnical Vocabulary, many dates. Available on-line with some delay at http://www.electropedia.org

ISO/IEC (2014) International Organisation for Standardization/International Electrotechnical Commision ISO/IEC Guide 51:2014. Safety aspects – Guidelines for their inclusion in standards, ISO & IEC, 2014. Available from http://www.iso.org/iso/catalogue_detail.htm?csnumber=53940).

Ladkin, P. B. (2008) Definitions for Safety Engineering, Causalis Limited. Available at http://www.causalis.com/90-publications/DefinitionsForSafetyEngineering.pdf

Lewis, D. (1969) Convention: A Philosophical Study, Harvard University Press.

Lloyd, E. and Tye, W.(1982) Systematic Safety, CAA Publications, London.

Locklin, S. (2014) Can the Su-25 intercept and shoot down a 777?, Locklin on Science blog, 2014-07-21. Available at http://scottlocklin.wordpress.com/2014/07/21/can-the-su-25-intercept-and-shoot-down-a-777/

Schelling, T. (1960) The Strategy of Conflict, Harvard University Press.

Schofield et al. (2014) Schofield, A., Flottau, J., Buyck, C., Broderick, S., and Croft, J., Ukraine Shootdown May Spur Risk-Assessment Reform, Aviation Week & Space Technology, 2014-07-28. Available at http://aviationweek.com/commercial-aviation/ukraine-shootdown-may-spur-risk-assessment-reform

Sweetman, W. (2014) How An Su-25 Can Shoot Down a Faster, Higher-flying Aircraft, Ares Blog, Aviation Week and Space Technology, 2014-07-24. Available at http://aviationweek.com/blog/how-su-25-can-shoot-down-faster-higher-flying-aircraft

Times (2014), The Times, 2014-08-08. The article is behind a paywall; the graphic at http://www.thetimes.co.uk/tto/multimedia/archive/00740/inline_937c3442-15d_740840a.jpg is not.

United States Government (various) 14 CFR Part 25, available at for example http://www.ecfr.gov/cgi-bin/text-idx?tpl=/ecfrbrowse/Title14/14cfr25_main_02.tpl