Proactively monitoring emerging risks through the analysis of occurrence and investigation data: Techniques used by the Australian Investigator

Stuart Godley
Manager Research Investigations and Data Analysis
Australian Transport Safety Bureau

Stuart is currently managing the Research Investigations and Data Analysis branch at the ATSB. He was initially employed as a Senior Transport Safety Investigator in November 2004 in the aviation branch specialising in human performance, organisational, and cabin safety perspectives. He was previously employed in the safety department of Qantas in safety analysis and human factors roles. Prior to Qantas, Stuart worked in applied experimental psychology at Monash University and the University of Sydney, focussing on road safety. He has a PhD from Monash University, and a Bachelor of Science with a psychology major from the University of New South Wales.

Abstract

As the national transport safety investigator, the ATSB has recently made several developments to move towards a predictive approach to safety in support of its core investigations function. Through greater use of reported occurrence data and the analysis of data generated from investigations themselves, the ATSB is aiming to provide a greater level of assurance to the Government, industry and the public about the level of aviation safety in Australia. In doing so, the ATSB is also maximising the chance that emerging and existing safety risks are actually uncovered before they lead to dire consequences.

This paper outlines the advantages and processes of a number of proactive safety initiatives currently being undertaken or in development at the ATSB. They include the regular systematic trend monitoring of reported occurrence data, the development and use of event risk classification for reported occurrences, the analysis of investigation findings across investigations, processes for ensuring safety issue risks have been addressed, identifying safety issues across the industry, and the use of investigations of groups of occurrences as an alternative or supplement to individual occurrence investigation.

Introduction

A holistic approach to aviation safety, either at an airline level or at a national level, can never be totally reactive. Like the safety department of an airline, national investigation agencies have a role in assuring their government that the nation's aviation industry is either safe or, when there is evidence that this is not the case, where industry and others need to focus to their attention. Like airlines, nations rely on the collective wisdom of a combination of proactive initiatives in addition to reactive investigations to do this.

This paper will explore recent developments within the Australian Transport Safety Bureau (ATSB) to integrate systematic predictive research and analysis into the nation's investigator. A key driver of these developments is the obligation to raise awareness of all potential safety issues within government and industry. However, it is also becoming an integral way of conducting business as an investigator by guiding the best use of limited investigative resources.

Importantly, however, the biggest driver towards the proactive safety initiatives in Australia's national safety investigator is the prospect of discovering what so many investigations of airline accidents have found. That the bureau, like an airline that collects safety occurrence data, conducts investigations, records findings and tracks safety actions, is sitting on information that could have highlighted an emerging safety risk *before* it led to a catastrophic accident.

The initiatives described below not only include the use of reported occurrence data, of which the ATSB is the national dataset holder, but also data generated through investigations themselves, including data generated from investigation findings and resulting safety actions. These approaches will hopefully challenge the notion that national aviation investigation authorities passively wait for accidents to happen so that an investigation can be launched and recommendations made, before the next accident occurs and the cycle continues.

Trend analysis of reported occurrence data

In addition to being Australia's independent no-blame investigator for aviation accidents and incidents, all aviation accidents and incidents are required by law (1) to be notified to the ATSB and the ATSB maintains Australia's official aviation occurrence database. The ATSB receives notifications from industry of over 8,500 occurrences each year, but conducts investigations on only about 100 of them. Although occurrences that have been investigated result in a rich data set for analysis, the addition of the non-investigated occurrences provides a window to the wider world of aviation safety that is not always available with the smaller subset of investigated occurrences.

Through a routine and systematic analysis of reported occurrence data, the ATSB can routinely observe the health of aviation across the country in the same way that a safety department does for an airline. Rather than looking for specific issues, the process is data driven, looking at everything in the database to see if there are subtle changes that may point to a larger issue. In this way, potential issues can be monitored, by the ATSB, other government agencies and industry, so that there is early identification of the potential for them to evolve into a significant and/or systemic problem. In addition, these flagged potential issues may point to the need for action by the ATSB, such as through targeting specific types of occurrences for investigation or initiating a broader systemic investigation of a particular transport safety matter.

Process

Every 3 months, reported occurrence data is processed to look for increasing or decreasing trends in types of occurrences (based on coded occurrence types from a taxonomy of about 100 types, such as wirestrike, hard landing, depressurisation, stall warning, loss of separation, power loss, birdstrikes etc). This is done separately for three aircraft operation types (high capacity operations (2), low capacity commercial air transport - including both regular public transport and charter operations, and general aviation).

Frequency counts of occurrences are not used as the only occurrence types that will have large differences between one period and the preceding period are the more frequently reported occurrences (which also tend to be those with the lowest inherent risk).

Rather, for each occurrence type and operation type, the count of occurrences and the rate per 100,000 departures for the most recent period (last 3, 6, and 12 months) is compared to the mean for the last 5 years to establish a historical baseline. The comparison is made in units of standard deviation from the 5-year period.

If the number of occurrences for the current reporting period is more than 1.28 standard deviations from the historical mean of the last 5 years, it is deemed to be significantly different from normal and a basic alert is generated. At this point, there is only a 10 per cent chance that the observation is due to chance alone. At 2 standard deviations from the historical mean, the chance of error is only 2.5 per cent, and a higher alert is generated.

When an occurrence type is greater than 1.28 or 2 standard deviations, it raises an alert for follow up. Further analysis can show what aircraft models, operators, locations, etc, account for most of the difference, and whether this has been a long term trend or a recent ‘blip’. When a single operator accounts for most of the difference, they are approached with the trend for both their awareness and so that any obvious explanation can be passed onto the ATSB. Sometimes, increases are solely due to better reporting (by a single operator or across industry), sometimes because of changes to operations, aircraft, or regulations, and sometimes there is no apparent explanation.

For each significant alert, one or more action is assigned by ATSB investigation managers to provide a follow-up or review of why that alert occurred. These actions are either:

- monitor into the following reporting period to see whether the increase is sustained
- stop monitoring (when a one-off alert returns to within normal limits)
- closely monitor through a weekly report to an investigator of new occurrences of this type to provide the opportunity to gather more information
- contact operator or industry participants to provide them with some information on ATSB-identified trends that affect their aircraft or operations
- report to the regulator (Civil Aviation Safety Authority (CASA)) and/or the air traffic control provider (Airservices) for input into their surveillance and other safety monitoring processes
- target occurrences for new investigations (possibly short factual-only investigations) on the basis that the trend may be exhibiting a safety issue
- initiate an ATSB research/safety issue investigation on the basis that specific reasons for the trend are suspected or known.

The trend monitoring report and assigned actions is presented to the ATSB Commission each quarter. It is then shared with the aviation regulator and air traffic service provider, and in the future, it will be shared with operators and industry associations. Future developments are also planned and include expansion into trend monitoring of technical occurrences for aircraft models and airspace and operational occurrences for airports.

Event risk

A limitation of the process of trend monitoring of occurrence data described above is that it focuses on occurrence frequency, but does not necessarily show where the greatest *risk* lies. To this end, the ATSB has been introducing *event risk classification* (ERC) of each individual occurrence entered into the database. The ERC methodology is based on the Aviation Risk Management Solutions (ARMS) (3) methodology developed for large airlines. However, to apply this across the entire aviation industry has required several modifications.

Using ERC on occurrence data has several advantages. At an individual occurrence level, it provides a benchmark which the ATSB can consider when deciding if an incident should receive a formal investigation. Used in analysis of occurrences, ERC ratings provide a single measure that incorporates many and varied aspects of each occurrence so that occurrences can be compared in a way that helps predicts the most likely who/what/where of the next accident.

Process

The ERC methodology calls for all occurrences to be assigned a classification based on the most credible potential accident outcome that could have eventuated, and the effectiveness of the remaining defences that stood between the occurrence and the most credible accident.

The ERC was developed for airlines that generally fly similar aircraft in similar operations. The ATSB, however, wanted to apply this methodology to all aviation occurrences, irrespective of the aircraft or operation type, from small single-engine privately operated aircraft up to Airbus A380 aircraft. As a result, it was important that an occurrence that could lead to a small number of fatalities was not given the same weight as one that could lead to 10 or 20 fatalities, or one that could lead to hundreds of fatalities. This was achieved through modifying the standard accident outcome rating table by adding an accident outcome for high capacity aircraft catastrophic accidents, and redefining a major accident as one leading to a maximum of six fatalities, as the latter corresponds to the maximum seating of most smaller and non-commercial aircraft (Table 1).

Table 1: Accident outcome ratings for the ATSB's ERC

Accident outcome	Outcome description
High capacity catastrophic accident	More than 38 fatalities
Catastrophic accident	Multiple fatalities (7 to 38)
Major accident	1 to 6 fatalities
Injury accident	1 or more injuries (no fatalities), minor damage to aircraft
No accident outcome	No potential for aircraft damage or injuries

The effectiveness rating was used as recommended in the ARMS methodology (3) as seen in Table 2.

Table 2: The ARMS ERC effectiveness ratings

Effectiveness rating	Definition
Effective	The safety margin was 'effective', typically consisting of several good barriers.
Limited	An abnormal situation, more demanding to manage, but with still a considerable remaining safety margin
Minimal	Some barrier(s) were still in place but their total effectiveness was 'minimal'
Not effective	An accident was not avoided, or the only thing separating the event from an accident was pure luck or exceptional skill, which is not trained nor required.

Combining the rating of accident outcome with the rating of the effectiveness of remaining defences results in a risk score for each occurrence and a risk level, as per Table 3.

Table 3: ATSB's ERC matrix

	Effective	Limited	Minimal	Not effective
High capacity catastrophic accident	250	503	2,503	12,500
Catastrophic accident	50	102	502	2,500
Major accident	10	21	101	500
Injury accident or minor aircraft damage	2	4	20	100
No accident outcome	1			

Automating ERC

The ARMS methodology calls for an individual assessment of each occurrence to assign an event risk according to the above steps. Although this process works at many airlines, the sheer number of reported occurrences has led the ATSB to take an alternative approach of having an automatically-applied event risk assigned through a series of rules. Apart from the lower daily time and effort required, this approach provides for greater consistency and allows for future adjustment of criteria without losing historical data.

However, a blanket risk for all occurrences of a certain occurrence type is not appropriate. For example, the credible accident outcome from a loss of separation occurrence is based in part on the capacity of the aircraft involved, and the effectiveness of barriers such as ATC intervention and on-board warning devices may vary from case to case. In some occurrences, the presence of other occurrence types (such as an airspace infringement, or a runway incursion) or circumstances (such as instrument conditions or the relative tracks of the aircraft involved) may influence the time or defences available to prevent an accident.

To overcome this, most occurrence types have a number of rules for the worst credible accident outcome and remaining barriers. When more than one rule is applicable to a single occurrence, the worst-case ERC is chosen. Of course, no set of rules can cover every possible occurrence, so an individual occurrence can also receive a manual coding.

A number of considerations are made when determining the ERC rules that make up the risk profile for each occurrence type. The rules take into account both exposure criteria and severity criteria, and use data already entered/coded into the database. Some of the most common considerations are:

- the inherent (minimum) risk for that occurrence type
- the aircraft operation type, and/or subtype (e.g. air transport, general aviation, agriculture, flying training) – this may have a bearing on the worst credible accident outcome (seating capacity) and the effective defences against this outcome (e.g. warning or other aircraft systems, pilot experience and training, system redundancy, provision of equipment or procedures for non-normal situations)
- the weight category of the aircraft involved (typically indicates seating capacity in combination with the aircraft operation type)
- the number of aircraft involved
- the number and type of engines fitted to the aircraft
- the type of aircraft involved (e.g. aeroplane, helicopter, glider, balloon)
- the presence of other occurrence types
- occurrence type-specific information (eg. for birdstrikes, information considered includes bird size, number of birds struck, and whether there was an engine ingestion)
- the airspace class where the occurrence happened (indicative of traffic density)
- the altitude and location of the occurrence, and whether the airspace was controlled or uncontrolled
- what aircraft parts or systems were damaged in the occurrence, and their effect on continued safe flight

Using ERC results

As per the ARMS methodology, the ERC results can be used in two ways. Firstly, on individual occurrences, the risk colour assigned can be used as a reality check for decisions to/not to investigate (Table 4). That is, if an early decision is taken to not investigate an occurrence, and that occurrence has a risk rating greater than low, then the decision may need to be reconsidered.

Table 4: Reality check matrix

Very high	→ Investigate and take action immediately
High	→ Investigate
Medium	→ Possibly investigate
Low	→ Database entry

The other use is through risk-based occurrence analyses of combining the ERC rating with occurrence frequency. Risk colours can be used to show whether increasing trends are due to more low risk events, possibly due to better reporting, or whether higher risk events are becoming more common, suggesting intervention may be needed. For example, Figure 1 shows the number of birdstrikes per 10,000 movements by occurrence risk. It suggests a better reporting culture across the 10 years as more no-accident-risk birdstrikes are being reported rather than there being a significant increase in the types of birdstrikes that could lead to severe accident outcomes.

Figure 1: Birdstrike rate per 10,000 aircraft movements

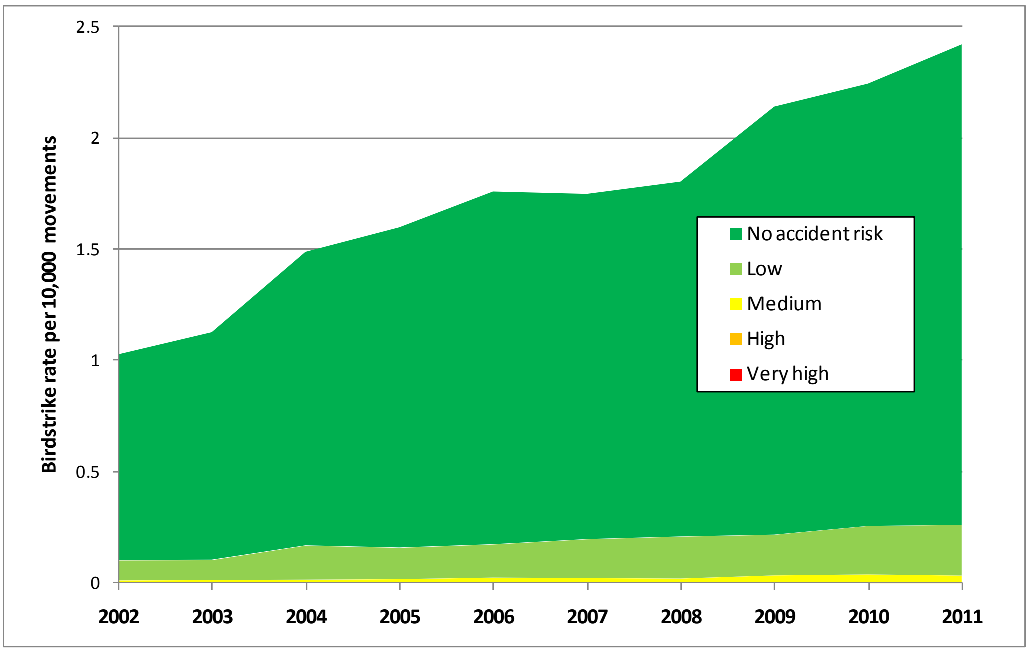

Analysis can also be done by summing the risk scores in Table 3. When the sum of risk scores is greater than the count of occurrences, it shows that there are risks greater than the no-accident-outcome level. Beyond this, however, the sum of the risk scores is meaningless by itself. Rather, it is used by comparing the relative risk between groups, where groups can be anything that is comparable such as months, years, aircraft types, aircraft operation types, airports, occurrence types.

Returning to birdstrikes, rather than looking for trends in frequency counts, trends in risk sums can be achieved. Looking at Figure 2, it is clear that birdstrike risk (orange bars) is not always a one-to-one relationship with the number of occurrences (blue line) as there are periods of higher risk which cannot be explained by frequency alone. In fact, the last 6 months show a higher risk even though the frequency count is lower.

Figure 2: Birdstrike rate risk per 3 month period for low capacity air transport operations

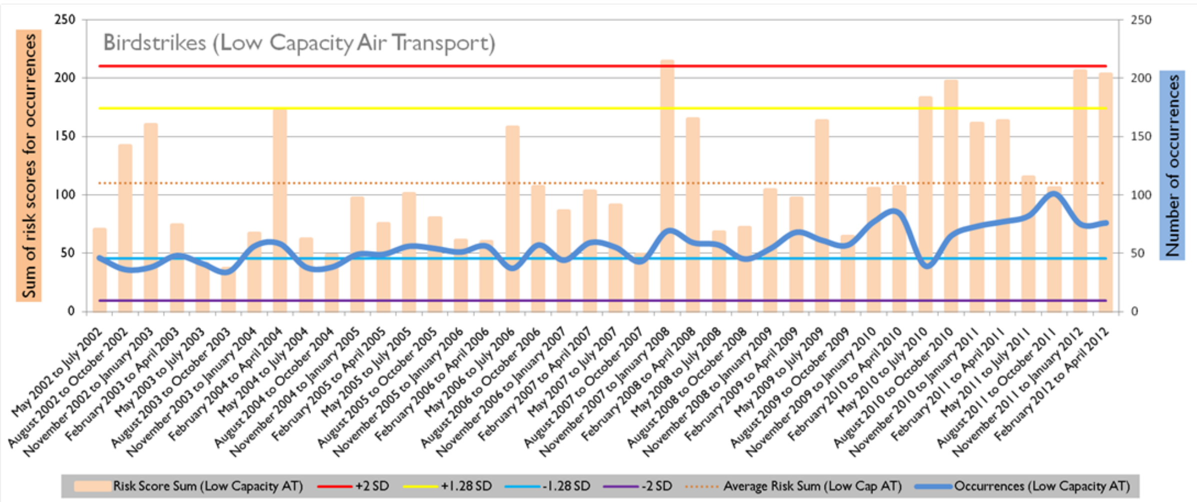

Another obvious comparison for birdstrikes is airports. Figure 3 shows those airports with the highest birdstrike rate per 10,000 aircraft movements (blue solid bars) are not necessarily the ones showing the highest risk (red outline bars). The higher risk airports are a result of a relatively high number of birdstrikes involving large and medium-sized birds, and additionally in Townsville, the number of multiple-bird birdstrikes involving large or medium-sized birds.

Figure 3: Birdstrike rate per 10,000 aircraft movements and risk per movement (2010-2011)

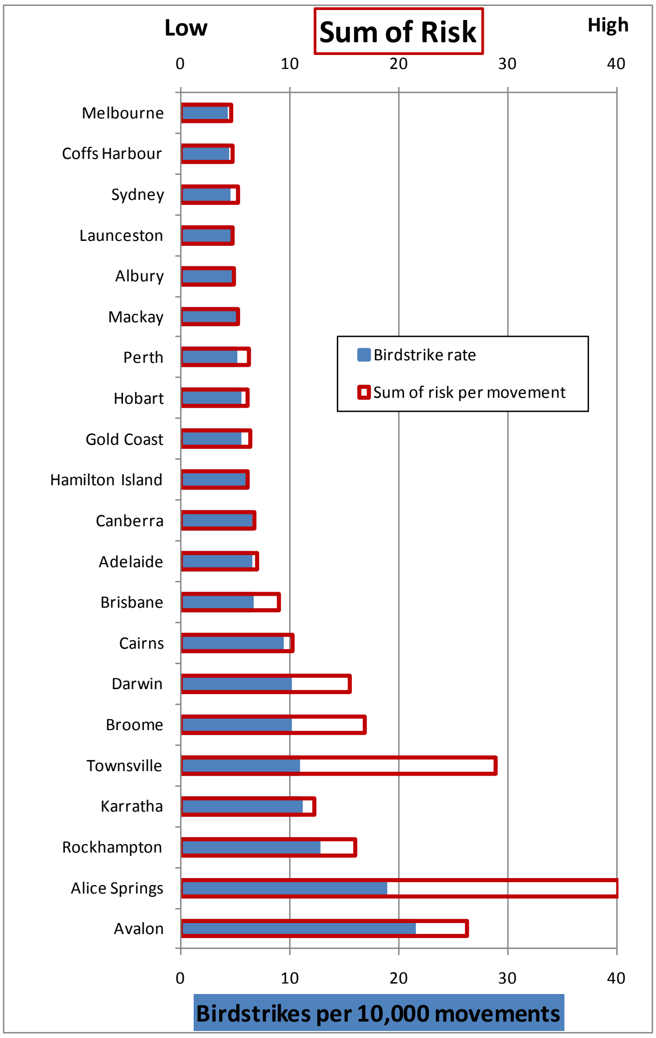

Analysis of safety factors from investigation findings

Findings published in ATSB occurrence investigation reports contain both contributing and other *safety factors*: events and conditions that increase risk. Individually they are important for individual investigations, but collectively, they can provide valuable evidence for why some accidents continue to occur. With such evidence, interventions and education can occur for whole sectors of the aviation industry rather than just single operators.

Research investigations

Safety factors from investigation findings are coded by investigators in the ATSB database using a taxonomy based on James Reason's model of accident causation (4). Each of the five levels of the taxonomy (individual actions, technical failures of equipment, local conditions, failed or absent risk controls, and organisational influences) have two levels of factor codes, and more than one safety factor code can be assigned to each investigation finding. Non-investigated occurrences can also have safety factor coding applied when there is sufficient information reported.

The analysis of safety factors have been used in research investigations, typically to understand why specific types of occurrences have happened, and can even been used to investigate why accidents in general are occurring. For example, a 2010 ATSB Research Investigation report analysed what was contributing to fatal accidents in private operations (5), and how these differed from non-fatal accidents.

Problems with pilots' assessing and planning were identified as contributing factors in about half of fatal accidents in private operations, and about a quarter involved problems with aircraft handling (Figure 4). Other contributing factors were visibility, turbulence, pilot motivation and attitude, spatial disorientation, and monitoring and checking.

Non-fatal accidents were just as likely to involve aircraft handling problems, but had fewer contributing factors than fatal accidents.

Figure 4: Contributing safety factors for accidents in private operation

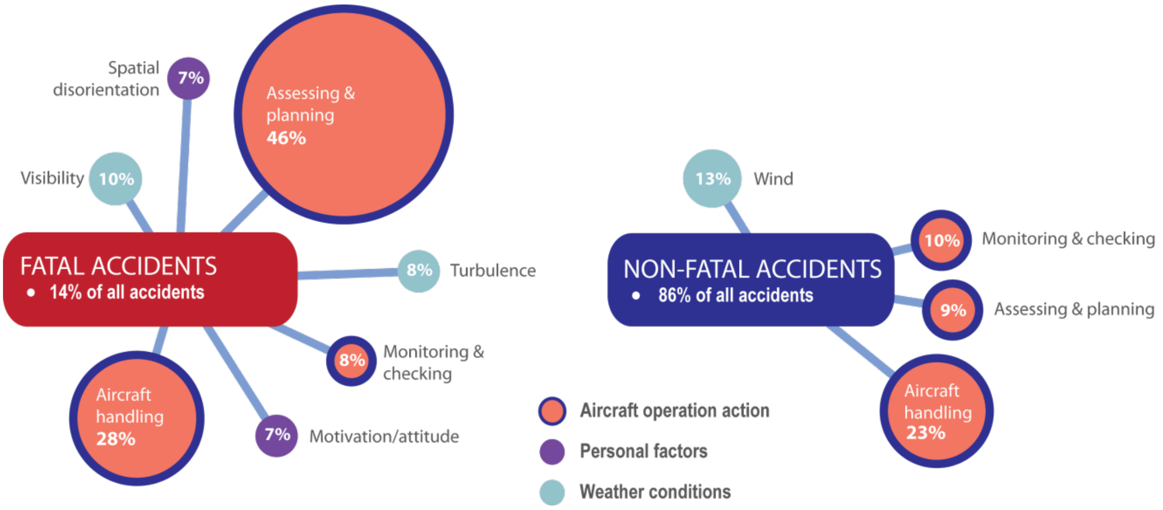

The large difference in the proportion of assessing and planning errors between fatal and non-fatal accidents shows their importance in avoiding serious accident outcomes and has led to training and education campaigns by the Australian aviation regulator, CASA.

Analysis of safety issues from investigation findings

Investigation findings for occurrence investigations also identify and risk-assess *safety issues*: safety factors that have an on-going influence. They include failed or absent risk controls to do with procedures, equipment, infrastructure, training, people management, and technical failures issues to do with design, manufacture, maintenance or operation, along with organisational or regulatory influences.

Safety issues are the systemic issues that we all aim to identify in investigations as it is only through fixing these that we can improve transport safety. The challenge, however, is how can we demonstrate any improvement in safety post-investigation so that safety issues uncovered in previous investigations do not remain as significant risks and lead to the next big accident?

The ATSB has approached this challenge by changing the focus from how many safety actions have been undertaken, whether these be from ATSB recommendations or proactively taken by industry, to whether identified safety issues have been eliminated or reduced in risk to an acceptable level.

Process

All safety issues identified in occurrence investigations are risk-assessed as they were at the time of the incident or accident. When this risk is considered to be significant (or in extreme cases, critical), the progression and resolution of the safety issue is monitored within investigation reports, quarterly reports to the ATSB Commission, and in a new development, through a listing on the ATSB internet site.

When the *owners* of safety issues, such as manufacturers, operators, regulators etc, report safety action undertaken, the ATSB conducts another risk assessment to determine whether the residual risk after that safety action has been reduced to an acceptable level or as low as reasonably practicable. When the safety issue risk is still significant, even if ATSB recommendations have been addressed or other safety actions have been undertaken, the ATSB Commission has clear visibility of the risk remaining and can continue to encourage for safety action.

Analysis

The recording and coding of safety issues has allowed for the systematic analysis of common types of issues across the commercial air transport and general aviation industries, along with the common types of safety actions resulting from ATSB investigations. This has been published for the each financial year, and is being developed into a multiple year analysis to look for changes across time.

As can be seen in Figure 5, as each safety issue is risk assessed as it was at the time of the occurrence, the use of a risk-based analysis of safety issues quickly identifies areas with a need for greater focus. For example, although only a few safety issues concerning equipment were uncovered through the 2010-11 financial year, a relatively high number of 'significant risk' safety issues were to do with equipment.

Figure 5: Safety issues identified in investigations completed in the 2010-2011 financial year

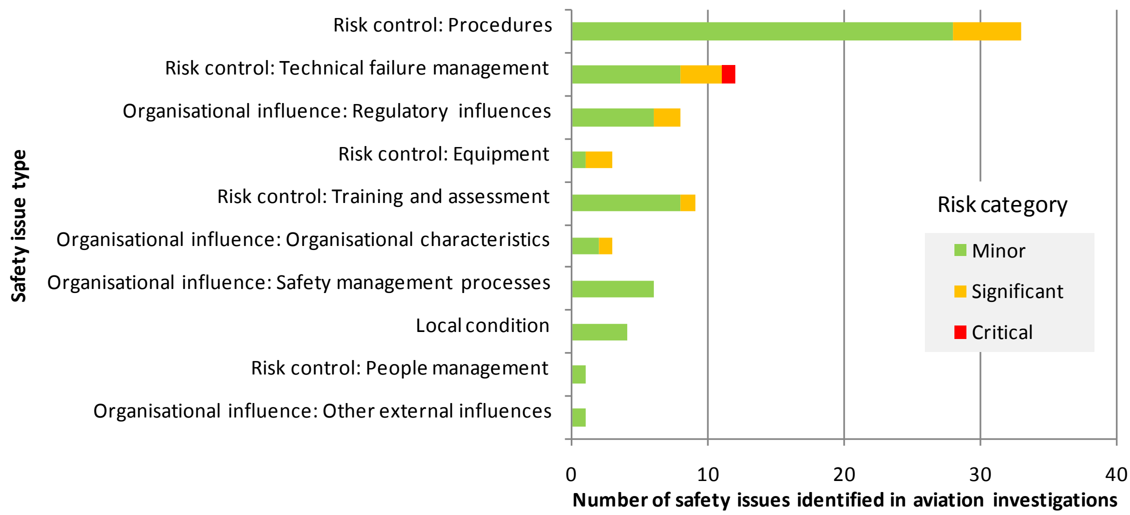

Longer term analysis is planned for safety issues across 3 or more years when the data is available. Combined with data about the post-safety action residual risk (an indicator of current risk), such analysis should point to the major unresolved issues facing aviation today.

Such analyses can feed directly into 'Top 10 most wanted lists' published by some national investigators.

Investigations of groups of occurrences

Although the investigation of some individual incidents and accidents can make safety inroads, sometimes a bigger picture only emerges when the same type of occurrence is looked at collectively. This is often the best option when similar occurrences become more frequent (as noticed from trend monitoring) or *appear* to be increasing (either through increased focus for investigation or through media interest).

Some types of incidents never by themselves pose enough immediate safety risk to launch an occurrence investigation, so the only way to investigate is through research-type investigations of groups of occurrences (such as all occurrences of that type across 3 years or 10 years). Apart from the above, such investigations can result from intelligence from industry, other government agencies, and also from issues raised within occurrence investigations.

Recently, the ATSB has initiated a research investigation into loss of separation (LOS) incidents. The study will examine these occurrences through: 1. analysis of all LOS occurrences across 4 years; 2. Identification of common findings and issues identified across a number of investigations; and 3. examination of confidential reports from employees of air traffic services providers. Through this approach, a holistic examination can be applied at a level not always possible for incident investigations, and issues such as whether these occurrences are becoming more common and/or riskier can be addressed.

Conclusion

For an investigation agency to improve safety, it must uncover safety issues and other emerging safety risks. This can be done, in part, through good investigation techniques. However, it also relies on the active search for emerging and existing safety risks across reported occurrences and across investigation findings. Such activities help guide decisions about which occurrences should be investigated (as individual occurrences or groups of occurrences), and therefore, affect the chance of discovering significant safety risks through investigations. Showing that similar findings are concluded across several investigations may point to larger or more widespread risks than the individual findings themselves suggest. Finally, improved safety can only be claimed when it is known that significant safety issues are not just identified, but have been reduced in risk.

By itself, however, there is also a very persuasive motivation for moving beyond reactive investigation into more holistic predictive safety activities. That is, are we doing everything we can to ensure that information we hold is not suggesting safety risks that, if uncovered and addressed, may have prevented a catastrophic accident.

References

(1) Aviation occurrences are required to be reported to the ATSB under the Transport Safety Investigation Act 2003 and Transport Safety Regulations 2003. These are available at http://www.atsb.gov.au/about_atsb/legislation.aspx

(2) High capacity refers to more than 38 seats. Low capacity refers to 38 seats or less.

(3) The methodology is from the report *The ARMS Methodology for Operational Risk Assessment in Aviation Organisations* (version 4.1, March 2010). ARMS is an industry working group set up 2007 in order to

develop a new and better methodology for Operational Risk Assessments. It is a non-political, non-profit working group, with a mission to produce a good risk assessment methodology for the industry. The results are freely available to the whole industry and to anyone else interested in the concept.

(4) See James Reason's two books: Reason, J. T. (1990). *Human error*. Cambridge, UK: Cambridge University Press; and Reason, J. T. (1997). *Managing the risks of organisational accidents*. Aldershot, UK: Ashgate.

(5) ATSB (2010). Improving the odds: Trends in fatal and non-fatal accidents in private flying operations. (Aviation Research and Analysis Report AR-2008-045). Canberra: ATSB