

Can We Save Money with Safety?

A Novel Approach for Assessing Benefits of Safety in the Aviation Industry

Alfred L.C. Roelen¹, Mauro Pedrali²

¹National Aerospace Laboratory NLR, Air Transport Division, Anthony Fokkerweg 2
1059 CM Amsterdam, The Netherlands, roelen@nlr.nl

²D'Appolonia S.p.A., Transport Division, Via San Nazaro 19
16145 Genova, Italy, mauro.pedrali@dappolonia.it

Abstract

In this paper we describe a new approach for assessing the benefits of safety measures in the aviation industry. This approach spans from the identification of target safety levels for the aviation industry to the development of a model of the air transport system. The most contributory factors to direct and indirect costs of accidents/incidents are also considered in order to provide all the elements for carrying out cost benefit analyses.

Introduction

It holds for any system that in order to have good performance with regard to a particular property, this property must be part of the design requirements rather than being a mere spin-off of the development process of the system. For this reason, aircraft speed and fuel-efficiency are design requirements instead of the coincidental outcome of the design. The same "design approach" should be followed with regard to air transport safety in order to achieve a significant improvement in relative safety.

Usually there are several ways to achieve a particular safety improvement. Moreover, particular safety improvements in some parts of the air transport system may be more easily achieved than others and at lower costs. For this reason, since the resources available for the improvement of safety are limited, it is imperative to establish more effective ways of spending them. In order to achieve this objective, it is our opinion that at least the following aspects should be taken into account:

- Settlement of Target Levels of Safety (TLS) in the Air Transport System (ATS).
- Development of approaches both capable of establishing quantitatively the current level of safety in the air transport system, and estimating safety effects in terms of probability reduction in accidents and incidents that can be expected from the implementation of safety measures, including regulation.

These two elements can already enable the assessments of which gaps between actual safety levels and target safety levels are to be bridged by safety improvement measures. Nevertheless, the identification of costs due to unsafety, i.e., direct and indirect costs associated with accidents and incidents, and the description of costs derived from the implementation of safety measures and regulation is also necessary.

If we combine these data in a quantitative method for predicting the effect of safety measures, both in terms of safety benefits (i.e., reduction of the costs of unsafety) and in terms of implementation costs, then it will be possible to perform cost-benefit analysis of proposed safety improvement measures. The approach described in this paper encompasses all these aspects.

The Approach

The approach that we propose has its foundation in almost 5 years of research conducted under the frame of the IV and V Framework Programme (FP) of the European Commission of the European Union. In particular two international projects involving many stakeholders of the air transport system (e.g., manufacturers, operators and regulators) contributed to the development of such an approach for achieving an adequately large improvement in relative safety. These projects are DESIRE (Method for Air Transport Safety Improvement through Quantitative Risk Assessment) and ASTER (Aviation Safety Targets for Effective Regulation).

DESIRE was funded under Aeronautics Technologies of the BRITE/EURAM Programme (IV FP). The main objective of this three-years-project (1998-2000) was to develop and evaluate a quantitative risk assessment model of air transport safety that allows cost-benefit analysis of safety measures. The model is intended to be used by all actors in the aviation system as a tool to help indicate those areas where the available safety budget can be invested more effectively.

ASTER was funded under the Competitive and Sustainable Growth Programme (V FP). The objective of this two-years-project (2000-2001) was the development of a meth-

odology enabling to set and optimize safety targets for each participant in the air transport system in order to achieve the optimum level of safety for the system as a whole.

In the following subchapters we will describe the main parts of this approach, which basically consist of the identification of target levels of safety in the aviation industry and the development of a simplified model of the air transport system for the quantification of risk. The most contributory factors to direct and indirect costs of accidents/incidents are also taken into account in order to provide the necessary elements for carrying out cost benefit analyses.

Target Safety Level

A Target Level of Safety is the amount of safety that is aimed for. It can be mandatory, i.e., it must be achieved in order to carry out some activity, or strongly desirable, i.e., a target that must be aimed for but not necessarily achieved.

A review was conducted on current Target Levels of Safety (TLS) within the civil aviation community, drawing relevant comparisons with other industry sectors (e.g., railway, nuclear). The work explored the role of stakeholders and the extent of their involvement in TLS derivation processes and considered the need for greater understanding of the relevant stakeholders. Strengths, weaknesses and developments of relevance to our needs were identified accordingly. However, before addressing TLS it is necessary to define safety. The following definition was used (Joyce 2001):

Safety is freedom from unacceptable risk, where risk is a combination of the probability of occurrence of harm and the severity of the harm. Harm is either physical injury or damage to the health of people either directly or indirectly as a result of damage to property or the environment

Consequently, in order to ensure that risks are properly identified and evaluated it is necessary to specify the persons (or groups, or classes of persons) exposed to risk, and establish what type of target is being, or should be, used in a target level of safety. The 'As Low As Reasonably Practicable' or ALARP approach identifies three risk regions (Figure 1), namely:

1. Above the Maximum Tolerable Risk. In this region the level of risk is so great as to be intolerable, unacceptable and therefore, the activity giving rise to the risk is not permitted until the risk is reduced below the Maximum Tolerable Risk.
2. At and below Negligible Risk Level. In this region the level of risk is so small as to be broadly acceptable.
3. Between the Maximum Tolerable Risk and the Negligible Risk Level. In this region the risk is only tolerable (i.e., acceptable) if a benefit is gained from the activity and either risk cannot be reduced or the cost of risk reduction is grossly disproportionate to the improvement gained.

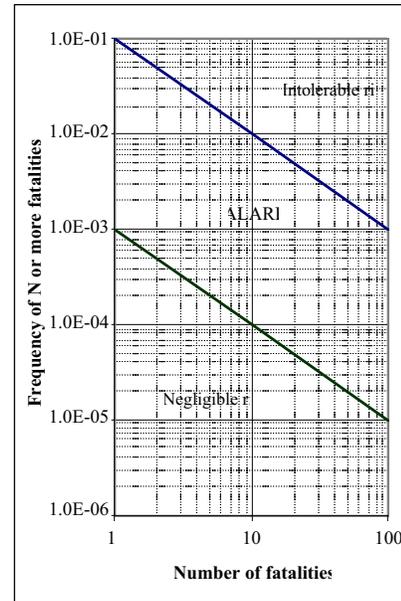


Figure 1: ALARP regions in the maritime domain.

If the Maximum Tolerable Risk is used as a Target Level of Safety then it is clearly a mandatory target since a higher level of risk is not acceptable under any circumstances. Similarly, the Negligible Risk Level is normally considered to be a strongly desirable target since higher levels of risk can be tolerated, albeit with some conditions attached. Target Levels of Safety set in the ALARP region may be either mandatory or strongly desirable. This arises because of the way in which the ALARP principle works to improve safety and drive down risk. Risk is considered tolerable if either the risk cannot be reduced or the cost of risk reduction is disproportionate to the improvement gained. Although a target may be set, failing to meet it would be tolerable if all reasonable options for risk reduction have been investigated and either the reduction cannot be achieved or the associated costs are disproportionate. However, mandatory targets may be set if it is clear that a particular level of risk can be achieved. For instance, a target might be set that is equal to an industry average or the best that is achieved within an industry sector. It is clearly reasonable and practicable to achieve it. Similarly, historical performance can be used to set a mandatory target. Again, it is clearly reasonable and practicable to achieve it because it already has been achieved in the past.

No single process for establishing TLS is best in all situations. The process should be appropriate for the commercial, legal and political situations of the industry concerned and should take into account technical issues and time scales. The process should be acceptable to at least the most important stakeholders, should be clearly defined and should result in a TLS that is both achievable and not open to significant dispute. Whilst approaches to TLS are more developed in Air Traffic Management, there is a recognized need for a framework to establish and control how

TLS are (or are not) being achieved, in view of the entire aviation system.

Air Transport System Model

The possibility to determine the level of safety in the air transport system is a requirement if we want to compare the current level of safety with the expected/predicted effects of proposed safety measures on this level of safety.

However, quantifying safety for an industry that already has a good safety record -such as aviation- inevitably implies to express safety in terms of small numbers. It is important that the corresponding limitations are understood and properly handled. Quantification requires the choice of units, in which risk is expressed, and data, by which risk is assessed. Units must be justifiable in terms of their final use and target audience. Data must be collected so that appropriate figures can be employed. For example, 'per flight' means that the number of flights must be counted, 'per km flown' means that distance must be determined. The accuracy of the measured safety will depend upon the accuracy of all data, not just the direct 'safety data' (e.g., deaths, hull losses).

The safety level can be estimated by using historical accident and incident data, or it can be estimated by using mathematical models that represent (parts of) the air transport system. In ATM, mathematical modeling for the assessment of mid-air collision risk is significant as a tool for determining the extent to which TLS are being achieved. Both approaches have inherent advantages and disadvantages. The disadvantage of using historical accident data is that it always provides a picture of the past rather than the present, and the fact that accidents are rare, which limits the statistical reliability of the estimates. Mathematical modeling results are limited by the quality of the model, the validity of which often requires comparison with statistical data with all associated difficulties.

We developed a mixed approach in which a simple mathematical model heavily relies on historical data for quantification. In so doing not only the model remains as transparent as possible, allowing a better understanding of its limitations but also, and most importantly, it facilitates straightforward adaptations to user needs and operational context.

We developed such a model by identifying the elements of the air transport system and the causal factors affecting the safety of each element. Elements and factors must be linked to be able to determine the overall level of safety of the air transport system. This was obtained by a combination of accident scenarios representing typical combinations of causal factors. These accident scenarios have been constructed for the following accident types:

- Loss of control in flight
- Controlled flight into terrain
- Collision with another aircraft
- General disintegration
- Landing gear related

- Runway veer-off/overrun.

The scenarios have initially been developed by analysis of a sample of accidents for which detailed information was available. Individual fault trees were established and combined into generic accident scenarios for specific types of accidents. The trees were further developed by combining them with a Functional Model of the air transport system that systematically describes functional relationships between different actors. Combining the initial accident trees with the Functional Model has the advantage that factors that could play a role but have not materialized in an accident are exposed. In addition, it directly identifies 'ownership' of causal factors.

To describe the scenarios, the model uses a classification of causal factors that is based upon the International Civil Aviation Organization (ICAO) Accident/Incident Data Report (ADREP) standard. The ADREP system describes an accident or an incident chronologically by listing the events that led to this occurrence together with causal factors that triggered these events. Causal factors can be both technical (Descriptive factors) and non-technical (Explanatory factors). We propose the use of a simplified classification, with fewer elements than the ADREP list. To enable a better representation of ATC related factors, elements of the 'Harmonization of European Incident Definition Initiative for ATM' (HEIDI) taxonomy, which has a structure that is similar to ADREP, were added to the classification.

Generic accident scenarios were integrated into a nodal network, with the top event "fatal accident". The nodes of the network consist of the elements of the classification scheme that were used to build the scenarios. States were defined for each of the nodes, indicating the status of that node. The number of possible states for each of the nodes was intentionally minimized to keep the network as simple as possible. To allow probabilistic causal analysis, a Bayesian Network was constructed (Roelen et al 2001)

A Bayesian Network is a combination of causal relations without closed loops in which relations between variables are represented through conditional probabilities. The Bayesian Network forms a representation of a joint probability function of the system, which allows a computation of the probability of any combination of values for the parameters in the model. A Bayesian Network allows easy representation in a graph, which can be helpful to provide more insight during the analysis.

A very pragmatic approach to quantification of the model was followed. Although the size of the model is kept modest (it consists of 49 nodes), full quantification requires more than 500 conditional probabilities to be determined. While a strong effort was made to justify all of the probabilities that are used in the model, limitations in available time and data required that some assumptions had to be made. It is important that users of the model are aware of those assumptions.

Availability of data remains a problem. The best sources of data are aircraft manufacturers and airlines. In the context of Flight Operational Quality Assurance (FOQA), air-

lines collect and analyze in flight recorded data using systems such as the British Airways Safety Information System (BASIS). After a comparison of the classification scheme and BASIS, it was concluded that both BASIS and FOQA data can be successfully exploited as sources of technical information. The classification could be modified so as to interface better with BASIS, but this would imply a reduction of the classification and consequently a lower precision for similar categories. As far as non-technical (human factors) data are concerned, neither BASIS nor FOQA are capable to provide relevant data at the moment. We could have referred to Human Reliability Analysis (HRA) methods, at least as long as human performance is concerned, in order to assess the probability of human failures in given contexts. However, this would have allowed to only tackle some of the causal factors reported in the classification and would have been out of the project's scope.

Availability of in-flight recorded data is difficult. Data is not generally shared with other operators. Often an agreement has been made with the pilot unions to keep in-flight recorded data within the Flight Safety Department.

Cost Factors

Costs may arise from accidents and incidents in a number of ways, both directly and immediate and indirectly, perhaps over a longer term. An important distinction is made with respect to incidents that can be either Non-Operational or Operational.

Non-Operational incidents are events in which the aircraft was damaged or even destroyed but there had been no intention of flight e.g., a hangar fire, or events during ground handling. These events may in some cases cause delays and cancellations but do not directly relate to any air safety issue. Operational accidents and incidents can be further differentiated as follows:

Events which resulted in death and injuries and/or substantial damage. These events will generate high costs.

Events where there were no injuries and aircraft damage fell below that defined as 'substantial' in ICAO Annex 13. These events might include minor runway/taxiway excursions, hitting animals or birds, minor ground collisions. These are minor damage events, with the cost of repairs not exceeding a few hundred thousand dollars but normally probably much less. Other costs may also arise such as delay but, again, they could normally be expected to be relatively minor.

Events where an accident 'almost' happened but where there was neither damage, nor injury, nor other disruptions. Such events might include loss of separation between aircraft, descent below a safe height, equipment/system failures etc. These 'Serious incidents', as defined in ICAO Annex 13, may be investigated by the State and/or internally by the airline, manufacturer, ATM provider. However, it is suggested that these investigations and their costs should be regarded as a safety action, a cost of safety rather than unsafety.

In our studies, analysis of accident data showed a direct link between aircraft damage and number of occupants killed. Moreover, for many other heads of cost the level of cost is largely dependent upon the accident severity. The accident severity scheme presented in Table 1 was used to model the effect of accident severity on the level of cost.

Level	Damage	Death
Catastrophic	100%	80%
Disaster	100%	30%
Major	80%	0%
Moderate	50%	0%
Minor	20%	0%

Table 1: Accident severity classification scheme

For each accident type considered a typical severity level was established; this provided the following results:

Accident type	Severity level
Loss of control	catastrophic
Controlled flight into terrain	catastrophic
Collision with another aircraft	disaster
General disintegration	major
Landing gear related	moderate
Runway veer-off/overrun	moderate

Table 2: Typical accident severity

The most significant determinants of cost arising from accidents and incidents that came out from our studies are synthesized in the following subchapters.

Aircraft Physical Damage

Under the cost model, aircraft physical damage is one of the key determinants in categorizing accidents. However, the actual costs arising from similar physical damage suffered by different aircraft can vary by perhaps as much as two orders of magnitude – consider an old Boeing 707 with a value of \$1 to \$2 million and a brand new Boeing 747 valued in excess of \$150 million. The destruction of these aircraft will produce very different resulting costs for aircraft physical damage. Costs will also vary with time as inflation causes new prices and costs for repair to gradually rise.

Hence, rather than using actual cost figures for individual aircraft, these figures were normalized so that an 'index' for relative degree of damage was obtained. The index is expressed as a percentage of the aircraft damaged. The average loss of aircraft value with age, as determined by Airclaims, across all Western-built jet airlines types for all years of manufacture and all market conditions was provided to help estimate aircraft damage costs.

Passenger and Crew Deaths and Serious Injuries Human life is precious and beyond price, it is therefore neither possible nor, indeed, desirable to attempt to put a price on it. Nevertheless, attempts are made to indemnify for the purely material losses arising from deaths and serious injuries. Cost benefit models increasingly use a Value of Statistical Life in their calculations where this 'value' generally includes an element of indemnity together with soci-

ety's 'Willingness to Pay' (WTP) to avoid a statistical fatality.

WTP values are based on individual preferences, which include perception and attitude to risk. It is not necessary that these preferences, perceptions and attitudes are the same for all types of risks. The WTP for a given reduction in number of deaths can vary by a factor of more than three for different contexts. Even accidents in different transport modes seem to have different values for the individuals; reduction in underground accidents has been found to be valued one and a half times the value placed on road accidents. From the WTP a Value of Statistical Life (VOSL) can be derived.

Loss of Reputation

A major accident may change the way the general public and, directly and indirectly, how business views the airline or the aircraft manufacturer which may be associated with the crash. This can have both short term and longer term implications for the company. A very important factor is the view of the public on 'who is to blame' in the accident. This public perception is largely determined by the media. It should be noted that although the airline involved may suffer a substantial loss of turnover, other airlines might profit from the crash, merely leading to a redistribution of revenue. From a (societal) economic point of view this offsets the impact of a crash suffered by the airlines involved.

Insurance

Probably because of the potentially catastrophic nature of air accidents and the large amounts of money at risk, insurance plays a large part in commercial aviation, removing most of the possible direct loss resulting from such accidents. It may be assumed that, in most cases, it is the airline's insurance which responds first to a loss with other parties becoming involved through subrogation or contribution. Separately, personal insurance e.g. life insurance, may also be expected to respond.

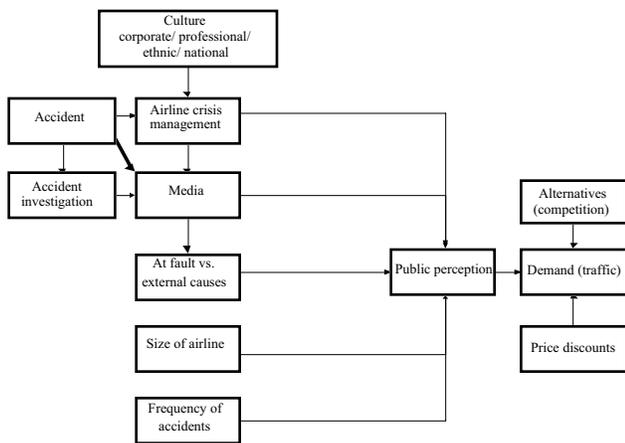


Figure 2: Interrelation of accidents and reputation.

Insurance, hull, passenger liability etc, removes most of the direct costs of an accident from the airline and other parties which may be involved (the insureds) so that the monetary

impact of an accident falls more immediately on insurers and re-insurers.

However, in undertaking the work for both DESIRE and ASTER, when discussing 'the cost of accidents' with anyone outside of insurance, a common response to the suggestion that most of the cost of an accident is taken away by insurance is, 'well of course if you have an accident, although insurers will pay the claim, you end up paying it all back in a couple of years in increased premiums.'

In most cases this is not true. A major loss, all else being equal, may result in an increased premium for two or three years but the insured would not be expected to 'pay back' the loss in higher premiums. However, if it is assumed that the cost of all claims paid by the insurers are eventually recovered from the insureds as a whole through the general level of premiums paid plus any additional loading placed on the premiums of those insureds which have suffered losses, then the cost of the loss will end up back with the airline community as a whole. Therefore, it would be double counting if the cost of insurance premiums are included in the cost of accidents, even if it is only the expected additional premium cost incurred following a loss.

Costs of Safety Measures and Regulation

Costs can arise at all stages of a lifecycle: design, manufacturing, operation, disposal. The lifecycle can be applicable to products, but also regulation and procedures. In the case of regulation, life cycle stages are research & development, implementation, compliance and termination. Within each phase, particular types of activities take place and were identified. The relative size of costs and the duration associated to each of the phases depends on the nature of the (safety) measure. The (safety) benefits will in general only appear during the operational stage.

Cost Benefit Analysis

If we want to optimize safety targets across the air transport system, this optimization has to be based on a number of criteria, one of which is the balance between safety benefits and costs. The method used by economists to appraise (government) policies and projects and to optimize costs and benefits across all parties is cost-benefit analysis (CBA). In CBA, all effects are expressed in monetary terms. Not all effects are priced on an economic market however. Non-priced effects must be valued by using surveys, behavioral data or shadow prices.

A first step in CBA is to define the scope of the assessment. An aspirational or mandatory target needs to be set. The actual level of safety must be determined by analyzing historical data, using mathematical models, or a combination of both. Comparison of the actual level of safety with the target may lead to the conclusion that improvement is necessary. Improvement alternatives may be identified by using (a combination of) expert judgement, common sense, literature review, research results, accident and incident investigation and operational experience.

The expected costs and benefits of alternatives are considered in relation to a base case. The base case is the situation that would exist if the alternative were not undertaken.

In identifying impacts it is necessary to carry out a stakeholder analysis. If costs and benefits of alternatives have been assessed, the actual CBA can be carried out. A frequently applied measure in CBA is the Net Present Value (NPV). The NPV is defined as the cash equivalent now of a sum receivable or payable at a future date. In order to calculate the NPV, it is necessary to discount future benefits and costs. This discounting reflects the time value of money. Benefits and costs are worth more if they are experienced sooner. The higher the discount rate, the lower is the present value of future cash flows. Hence the value of the discount factor reflects the preference of society for today's income versus income later in time. The discount rate represents also the cost associated with diverting investment resources from alternative investments or from consumption.

Sensitivity Analysis

It must be kept in mind that in conducting any cost benefit analysis a number of assumptions have to be made, for instance with respect to the efficiency of certain measures, expected traffic growth, future technological developments, displacement effects, hardware costs, etc. The amount of uncertainty in such assumptions can be substantial. For this reason it is recommended to perform a sensitivity analysis, where the effects of systematic variations of key parameters on the economic measures of the project viability are determined.

Strengths and Weaknesses

The proposed method is generic in the sense that it can be applied to the aviation system as a whole or to parts of it, such as geographic regions or specific actors. The method can be a valuable tool for regulatory bodies as well as airlines, ATC organizations, aircraft manufacturers, etc. However, because it is generic, users of the method are expected to have their own sources of data that can serve as input to the described method.

Application

The approach described in this paper has been applied to two cases: Cost Benefit Analysis of an Electronic Pilot Activity and Alertness Monitor (Speyer et al, 2001) and Cost Benefit Analysis of wake vortex measures [Ref. 4]. Both cases demonstrated the practical applicability of the method, and highlighted the importance of being able to compare and express safety benefits in monetary terms.

Conclusion

An approach has been developed for assessing benefits of safety measures in aviation. The method can be used to quantify the effect of safety measures, both in terms of cost savings due to safety benefits and in terms of implementation costs.

An important element of the approach is the air transport system model, a fully quantified Bayesian Network developed in a HUGIN software environment. This model was used for the two cases to assess the effect of a change in the ATS on the level of safety. While the model can be used directly for calculations, in many cases expert knowledge from the user is required to tailor it to the specific circumstances that are under consideration. Similarly, although considerable care was taken in selecting the factors, the interconnections, the probabilities and the financial equations and data, it is expected that users of the method will have their own sources of data that they would want to use.

Acknowledgments

The authors thank all persons that provided helpful insight and comments on writing this document, in particular Carlo Cacciabue (JRC), Paul Hayes (Airclaims), Tony Joyce (Eurocontrol) and Roelof Jan Molemaker (Ecorys).

References

1. T. Joyce, G. Graham, S. Kinnersly, M.J.A. van Eenige, A.L.C. Roelen, *A study into Target Levels of Safety (TLS) within the aviation industry, including comparative analyses with the rail and nuclear power sectors*, ASTER WP 1, NLR-CR-2001-145, NLR Amsterdam, April 2001.
2. A.L.C. Roelen, Y.S. Cheung, M.H.C. Everdij, E.A. Bloem, M. Pedrali, P. Hayes, *Development of a scenario-based method to assess safety levels within the aviation system and to determine the effect of changes on the level of safety*, Part 1: ASTER WP 2/5 main report, NLR-CR-2001-458-PT-1, NLR Amsterdam, October 2001.
3. J.J. Speyer, A.L.C. Roelen, *Evaluation of the DESIRE methodology: Cost benefit analysis of an Electronic Pilot Activity and Alertness Monitor*, DESIRE-WP5-TR5C, NLR Amsterdam, April 2001.
4. A.L.C. Roelen, R. Piers, R.J. Molemaker, P. Hayes, A. Joyce, G. Graham, *A cost benefit analysis of wake vortex measures*, NLR-CR-2001-610, NLR Amsterdam, 2001.