

RISK ANALYSIS IN COMPLEX SOCIO-TECHNICAL SYSTEMS: GLOBAL AND LOCAL RISK ASSESSMENT

Steve Hall, Ph.D.
Galaxy Scientific Corporation
Egg Harbor Township, New Jersey

Traditional risk analysis methods were designed for use primarily in mechanical-technical systems, but recent efforts have attempted to apply these methods to social-technical systems. Such systems are common in transportation and often encompass multiple technical sub-systems as well as several layers of human operators and several layers of management. The current project is tasked with measuring the safety of Title 14 CFR Part 137 agricultural aviation operations, which combines technical, human, and social systems in the context of inherently dangerous aviation work. While different operators may engage in the same set of flight activities, the risks associated with these operations vary from operator to operator. Thus, there is a need to quantify the risks associated with these operations in general and then adjust these assessments based on the characteristics and protocols of the specific operator. This paper discusses the development of a risk assessment methodology that begins with a risk assessment to evaluate the hazards associated with agriculture flight operations and moderates the risk values based on various operator characteristics and controls.

RISK ASSESSMENT AND THE HUMAN FACTOR

The concept of minimizing the risk associated with unwanted events is nothing new and serves as one of the underlying principles of both systems safety and human factors engineering. Human factors efforts tend to focus on interface and machine design while systems safety emphasizes potential hazards in a series of operations. Through safety conscious design, the likelihood of errors and failures can be minimized, but never eliminated.

Hazards are everywhere in aviation operations and the resulting risks require constant attention to ensure an acceptable level of safety. Much has been done to increase the reliability of aircraft mechanical systems, but the human element has proved much more difficult to contain. The result is that the proportion of accidents attributed to mechanical failures has reduced more in the past 40 years as compared to those blamed on human failures (Weigmann & Shappell, 2001). This is not at all surprising when one considers the fact that humans play key roles in maintenance, aircraft operations, and air traffic control; conversely, technological innovations in electronics, engineering, and manufacturing have made hardware and software components more precise and reliable. Similar trends have been noted in non-aviation domains as mechanical systems have become more refined and reliable and the human operator's role has become more complex (Cacciabue, 2000).

Traditional approaches to risk analysis (through the lens of systems safety engineering) have emphasized the quantification of risk mainly through the assessment of component failure rates. In such a paradigm, the human element in systems is seen as another complex piece of hardware with some quantifiable failure rate. As Redmill (2002) pointed out, hardware components fail in relatively predictable ways for certain reasons (e.g. manufacturing defects, component wear-out, improper maintenance, etc.),

but humans fail for a host of reasons, some of which cannot be predicted or even foreseen (e.g. sabotage, psychological breakdown, whim, etc).

While human reliability analysis (HRA) techniques are available, they are no where near as developed or tested as their hardware component counterparts (Redmill, 2002). Integrating HRA into the probabilistic risk assessment (PRA) model requires a thorough understanding of the requirements placed on the human in the context of the system and a thorough understanding of human performance. This is a move away from the use of human failure rates in PRA to a modeling approach for the prediction of human failures (Mosleh & Chang, 2004; Strutt, Loa, & Allsopp, 1998).

HUMAN PERFORMANCE IN CONTEXT

These new models of human performance hold promise for use in well-scripted and well-defined processes and systems. At one extreme, some systems use human operators as system monitors who are programmed to perform specific actions given certain states of the system. A given system state results in the operator performing a set of prescribed actions. At the other extreme, some systems are more goal-oriented, meaning that it is up to the operator to use a combination of tools and processes to achieve some desired outcome. Modeling human performance is a much more manageable task in the first situation as compared to the later. In other words, the fewer degrees of freedom the human operator has, the easier it is to model human performance.

While such models hold promise for aiding risk assessment, there are practical concerns about the generalizability of such models across local situations. It is often desirable and more efficient to conduct a generic risk assessment for a specific type of system or process, such as power generation or aviation operations, and then apply that information to a specific exemplar of the system (e.g.

aviation operations at a specific airline). The problem with this approach is that the results of the risk analysis will have to be adjusted to account for local factors, such as environmental and management issues, that impact both human and hardware reliability and performance.

HUMAN FACTORS WITHIN ORGANIZATIONS

The traditional human factors perspective on human performance emphasizes the roles of interface and equipment design, task design, human physiology, and human cognition in determining human performance. A factor that is less often considered, especially in the context of risk analysis, is the role that organizational structure and management play in human performance. From a safety perspective, the human factors literature has addressed the role of management under the rubric of the safety culture and work in the area of crew resource management has evaluated the influence of team dynamics on human performance. Wiegmann and Shappell's (2001) Human Factors Analysis and Classification System (HFACS) is an excellent example of a human performance model that explicitly incorporates organizational structure and processes into the estimation of human performance.

The disposition of management with regard to safety practices and performance monitoring can certainly impact human and system reliability and performance, but the fact that management *is* a part of the overall system is seldom accommodated in the risk analysis process. In the case where a risk analysis is performed in the local context, it can be argued that the influence of management on safety is accounted for by default, but when the goal is to generalize a risk analysis to a variety of local settings, it is clear that the role of management must be accommodated in the process.

THE CURRENT PROJECT

The Federal Aviation Administration (FAA), through the Systems Approach to Safety Oversight (SASO) project, has sponsored research to examine methods of system safety monitoring in both commercial (Title 14 CFR Part 121) and agricultural application general aviation (Title 14 CFR Part 137) operations. One of the goals of the agricultural aviation research team is to develop a safety measurement system that will assess system safety at the operator level as a proof of concept. While the operators involved in this segment of aviation all engage in basically the same type of activity, it is unrealistic to expect that they operate with the same level of safety. This is due to the fact that the various operators work in different geographic areas, operate different models of aircraft, and enforce a variety of different organizational policies and procedures. Operator safety is also influenced by the pilots and mechanics that work for the operator, with some pilots and mechanics obviously being better and safer performers than others.

The development of a safety measurement system must balance the need for precision on one hand and the need for simplicity on the other. Precision comes with complexity and at the expense of generalizability; conversely, a simplistic measure with high generalizability may not provide all of the information desired. The applied nature of the current project (i.e. develop a safety measurement tool that will be used in the field by inspectors and operators) emphasizes the need for an easy to use safety metric that is based on data that are available in current databases or could be collected prior to or during a site visit.

The safety measurement framework

The premise of the safety measurement system under development is that safety at the local level can be quantified by assessing the risk of Title 14 CFR Part 137 operations in general, identifying safety measures that an operator could use to reduce risk, and assessing risk at the local level based on the specific safety measures that the operator has in place. In essence, the crux of the safety measurement process is the identification of management-level policies and procedures that are thought to enhance and promote operational safety, namely pilot safety. In other words, the safety metric is based on the efforts that management is making toward improving safety.

The novelty of the current approach

Most risk assessment efforts are conducted within the context of a specific organization or operation. That is, most risk assessments are local. While this approach produces a very precise assessment of risk, it is not a reasonable approach to use when assessments must be done across multiple organizations or multiple assessments must be conducted over time. This is primarily due to the fact that risk assessments are very time consuming. The current approach seeks to assess risk using a generic task framework then identifying operator characteristics and controls that are associated with risk reduction. The result is that a given hazard may pose a greater degree of risk to one operator relative to another. The basis of the approach is to compute a separate risk value for each combination of operator characteristic and risk control that has been identified as having an impact on the likelihood of a given hazard resulting in an accident. Information from a specific operator can then be used to adjust the generic or baseline risk values to identify hazards and risks at the operator level.

The major benefit of this approach to risk assessment is that a single assessment can be conducted and local risk assessment values can be computed by collecting information specific to a single operator. This eliminates the need for a time-consuming local risk assessment. Since the local risk values are based on basic information provided by the operator, the local risk estimates can be

rapidly updated to reflect the implementation of new operator policies and procedures.

Another potential use of the risk assessment data involves the selection of safety measures at the operator level. For example, the risk assessment process will identify specific safety measures that are linked with risk reduction; therefore, the operator can use this information to assess the anticipated impact of specific safety measures given the specific characteristics of the operator.

The proposed risk assessment approach

The safety measure will be based on the results of the risk assessment process. The risk assessment process entails a series of steps designed to identify the types of accidents that occur in agricultural aviation operations, determine why those accidents occur, and identify risk controls that are currently in use to keep accidents from occurring. Inherent in risk assessment is the notion that some accidents are more likely or have more serious consequences than others. Accidents that are likely to occur and produce severe outcomes are considered to pose more risk than accidents that seldom occur or have minimal consequences.

For the current project, the risk assessment phase has been broken down into several steps.

1. Identify the types of accidents that occur.
2. Rank-order these accidents according to risk-frequency and severity (i.e. risk).
3. Identify the sequence of hazards that can lead to unwanted events.
4. Identify proactive steps that operators can take to interrupt these sequences to avoid an accident.
5. Determine which sequences are most likely.

Thus far, the first two steps of the process have been completed and the research team is preparing for upcoming meetings with Subject Matter Experts (SMEs) to complete steps three through five.

Accident event identification For the current project, data from the FAA Accident/Incident Database System (AIDS) were used to categorize the types of accidents that have occurred over the past 20 years in agricultural aviation operations. About 3,800 usable records were identified from over 5,500 recorded accidents. Over 41 accident types were identified in the database, with nine different accident types accounting for over 80% of the accidents and 13 different accident types accounting for over 90% of the accidents.

Rank-order accident types The accidents were categorized according to phase of flight. The frequency of each accident type within each phase of flight was divided by the total number of accidents within each phase of flight to obtain the relative frequency of each accident type. A six-point severity scale was created using AIDS information about aircraft damage and fatalities. A value of zero was assigned to accidents that involved no aircraft damage and no loss of life and a value of five was assigned whenever a fatality was involved. The remaining points

related to the amount of aircraft damage (1 = minor; 2 = substantial; 3 = destroyed). The frequency and severity values were multiplied together to compute the risk value of each accident type within each phase of flight. The accident types were then rank-ordered from high to low and the accident events that accounted for 80% of the accidents within that phase of flight were moved on the hazard chain construction phase.

Hazard chain construction The process of determining how accidents can occur is accomplished in the hazard chain construction phase. Accidents are seldom the result of a single hazard; instead, they are the result of a sequence of hazards and events. The lack of empirical data regarding the sequence of events prior to accidents makes the use of SME (e.g. FAA inspectors and agricultural application pilots/operators) input necessary. The hazard chains will consist of proximate, intermediate, and root causes, though early construction trials indicate that the chains will seldom be this "neat". The SMEs will be asked to first identify plausible reasons why a specific unwanted event might occur (i.e. identify the proximate causes). From there, the remainder of the hazard chain is constructed for each proximate cause. The SMEs will also be asked to identify contributing factors such as weather conditions, fatigue, etc.

Identify real-world risk control measures For each proximate cause chain, SMEs will be asked to identify policies, processes, and procedures designed to keep each proximate cause from occurring. The emphasis will be on identifying such controls that they have actually seen in practice. SMEs will also be asked to identify specific operator characteristics that might be linked with the occurrence of specific proximate causes. For example, some operators may be more susceptible to certain proximate causes given the geographic region where they operate.

Assess likelihood of proximate causes The final step in the risk assessment process is to estimate the likelihood that each proximate cause will occur *and* result in the unwanted event. While traditional risk assessment approaches produce a single risk value for a given proximate cause, the current approach will evaluate the risk for each proximate cause given every possible combination of operator characteristic and risk control measure.

For example, suppose that three separate real-world risk control strategies and one demographic factor (with 2 levels) have been identified for a single proximate cause. For each level of demographic factor, there are eight unique combinations of the risk control methods (see Table 1) for each type of certificate holder. Thus, a total of 16 likelihood values are possible for this one proximate cause. Recall that the severity values are based on the unwanted event that results from this proximate cause and stays constant regardless of the likelihood values. Risk values for the unwanted event can be computed based on the product of the proximate causes' likelihood and the single severity value associated with the unwanted event.

Table 1. Example Risk Evaluation Matrix For A Single Proximate Cause: Cells Hold SME Likelihood Ratings

Certificate Holder Demographic	Risk Control Present							
	None	RC1	RC2	RC3	RC1 & RC2	RC1 & RC3	RC2 & RC3	RC1, RC2, & RC3
Type A								
Type B								

Data analysis

There are several analyses that will be applied to the risk data. First, inter-rater reliability and agreement should be analyzed to determine whether or not the ratings are reasonably reliable. Intra-class correlation coefficients (ICC) will be used to assess inter-rater reliability, where SMEs are seen as a random factor, each unique combination of control and characteristic is seen as a fixed factor, and the absolute agreement definition is used. Rater reliability and agreement are both important given that individual points on the likelihood scale have different qualitative as well as quantitative meanings.

The format of the risk data collection allows for the estimation of effect for each identified risk control mechanism on the various proximate cause likelihoods. Such analyses could be used to identify which controls are viewed as most effective by the SMEs.

The combination of risk data and control effectiveness information make it possible to develop a safety audit program that applies different weights to risk controls given demographic information about the operator. The implication here is that some controls will be more effective for some operators than for others. Theoretically, the operator could combine this information with cost of implementation information to obtain information about the relative utility of various risk controls. This concept is in alignment with the general risk management principle of balancing the cost and benefit of risk controls.

The pilot factor

Since a substantial portion of aviation accidents are a result of pilot error, it seems reasonable to conclude that information pertinent to a specific pilot within a specific operation would be an effective component of any safety metric. The reason that pilot factors were not included in the current approach was a matter of practicality and logistics. The available accident data could be used to link factors such as pilot experience and pilot age with accident events, but these factors would likely explain a relatively small amount of variability in assessed risk relative to the management-level factors being considered.

DISCUSSION

All aviation operations involve risk stemming from multiple sources including equipment failure, pilot error,

and management practices. Traditionally, systems safety engineering has focused on the issue of equipment failure while human factors specialists have attempted to reduce the likelihood of pilot error. What has been missing to some extent is the inclusion of management level factors in the effort to reduce operational risk.

The current project aims to identify and assess risks across an entire segment of aviation operations, namely agricultural aviation operations. This is a daunting task because a generic risk assessment encompassing typical agricultural aviation tasks will not provide any means to discriminate between individual operators in terms of safety. Additionally, approaching safety assessment via operator level risk assessments is not practical.

The proposed solution involves performing a generic hazard identification and risk assessment with an emphasis on collecting information about operator level characteristics and risk control measures that might impact operational safety. The premise of this approach is that management-level policies and protocols set the context for safe pilot behavior, or lack thereof. This perspective emphasizes the notion that humans do not perform tasks in a social vacuum; instead, the social and organizational context within which humans perform tasks can have a marked influence on safety in the cockpit.

REFERENCES

Cacciabue, P. P. (2000). Human factors impact on risk analysis of complex systems. *Journal of Hazardous Materials*, 71, 101-116.

Mosleh, A. & Chang, Y.H. (2004). Model-based human reliability analysis: prospects and requirements. *Reliability Engineering and System Safety*, 83,241-253.

Redmill, F. (2002). Human factors in risk analysis. *Engineering Management Journal*, August, 171-176.

Strutt, J.E., Loa, W. & Allsopp, K. (1998). Progress towards the development of a model for predicting human reliability. *Quality and Reliability Engineering International*, 14, 3-14.

Wiegmann, D.A. & Shappell, S.A. (2001). Human error analysis of commercial aviation accidents: Application of the Human Factors Analysis and Classification System (HFACS). *Aviation, Space, and Environmental Medicine*, 72(11), 1006-1016.