SAFETY SCIENCE Monitor

Special Edition 2006

Article 5

VOL 10

SPECIAL CHALLENGES

GERHARD BECKER
TUV Rheinland, Berlin Brandenburg, Germany

J PARIES
Dedale SA, CNRS, France

CORINNE BIEDER
Dedale SA, France

TOM HEIJER
Safety Science Group, Delft University of Technology, The Netherlands

ANDREW HALE
Safety Science Group, Delft University of Technology, The Netherlands

1. COMPARABILITY OF RISKS (GERHARD BECKER)

The comparison of risks from different sources implies the trial to compare very different forms of harm: death, different forms of remaining handicaps from accidents, diseases – suddenly occurring or slowly progressing. Risk comparison needs to reduce the multidimensional forms of harm to only one dimension, expressed by a figure. This figure is used as a measure for comparison of the risks associated with e.g. different industrial forms of energy production.

The International Commission of Radiological Protection has proposed an “index of harm” already in 1977 as a method for comparing the risk associated with the exposure to radiation to that from other occupational and health hazards.

The paper deals with a study the author and his colleagues Jochen Hennig, Elmar Schaaf and Horst Schnadt have performed some years ago. In this study the risk of various jobs in industry, trade and services has been determined by evaluation of data of the German Berufsgenossenschaften. The construction of an index of harm for the various jobs with these data has led to the question of how to combine the different forms oft harm mentioned above. Depending on the assumptions and weighting factors applied, we obtained considerable variations in the sequence of “risky“ occupations.

1.1 Introduction

The risk debate and various approaches to assess "the risk" accompany the use of nuclear technology since the early years of its commercial application. The probabilistic risk assessment is meanwhile a standard tool used during licensing and operation of nuclear plants. In the course of the growing risk research around the nuclear technology there have also been trials to compare the risk of nuclear energy production with other forms of energy production (e.g. Inhaber, 1979) Such studies cover the comparison of the risk of accidents for the population living in the vicinity of the plant as well as comparisons of the pollution during operation and the health impact which might result thereby.

This paper is not dealing with the trial to compare the risk to the public resulting from the industrial installation. Instead, the occupational risk of working within different industries shall be discussed.

The study, performed already nearly two decades ago (Schaaf et all, 1986) had the aim to compare the risk of being exposed to radiation with the risk of working in other industrial fields. A further aim of the study has been to include a broad spectrum of different occupations in order to judge the adequacy of occupational dose limits. The aim of this paper is to illuminate the problems we had with the combination of the different “dimensions” of risk. What I want to discuss here is if ideas are available to overcome such problems.

1.2 Health risks from various occupations

Death rates have commonly been used as an index of the comparative safety or harm. But in the conventional industry accidental death is only one of the risk to which workers may be exposed. Other risks may be injuries with and without permanent consequences and occupational diseases. Any attempt to review the relative safety of different industries therefore involves two forms of assessment: firstly, the estimation of the size of all significant risks, such as the rates of accidental death, injuries and diseases; and secondly, some evaluation of the relative amount of detriment that is judged to be caused by each kind of harmful effect (ICRP45).

Our assessment of the size of significant occupational risk has been based on data of the Berufsgenossenschaften, which act as the workers´ compensation insurance companies in Germany. The data we collected from various Berufsgenossenschaften cover a large spectrum of jobs/professions with very different risk, from miners to jobs in the office.

In Germany any occupational accident, which results in absence from the work of more than 3 days has to be reported to the Berufsgenossenschaften. A fraction of the reported accidents concerns those where the insurance has so pay compensation e.g. for death, permanent reduced ability to work or retirement from the job. For such cases there is documentation available containing details about the accident and its consequences as well as data about sex, age, year of accident, profession (job area), length of illness, time in the hospital, and the year of death of the victim. The data differentiate between accidents at work and accidents in travel to and from work.

Data concerning occupational diseases could also be obtained from the Berufsgenossenschaften, structured in a similar way as the accident data.

In order to calculate risk data like accident rates the number of employees working in the various areas and differentiated on background variables (sex, age classes) have to be available. In earlier publications on occupational accident rates the authors tried to extract data about the profession/job from census data. It is known that such data from polls are not very reliable. For this reason we looked for other data sources and found one. In 1973 the obligatory registration for the social insurance had been changed. Since then, employers are bound by law to announce the number of employees together with information about their profession/ jobs to the Federal Office of Statistics in Germany. The code used in the database of the Federal Office of Statistics to identify the profession/job of the employees is the same as used by the Berufsgenossenschaften. So we were able to achieve the necessary data about the employees in the different occupations as a special evaluation from this database.

From both data sources mentioned above we were able to calculate rates of occupational accidents causing death, diseases or permanent disability, similar rates for accidents in travel to work as well as rates for occupation caused diseases.

1.3 Travel to work

The travel to and from work is usually excluded from risk comparisons. The data available at the Berufsgenossenschaften include detailed information on such accidents in similar detail to the occupational accidents. The evaluation of these data showed that such accidents vary not only with age and sex but there is also a variation with the occupation, as Figure 1 shows. As a first approach one could say that travel accidents increase with the roughness of the job.

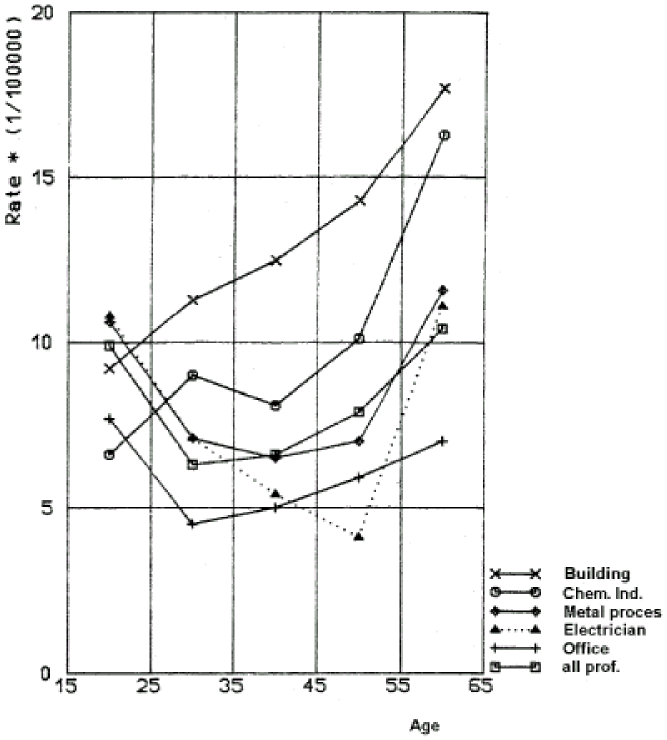

Figure 1 Fatal travel accidents (men)

1.4 Effects exposure to radiation

Occupational radiation exposure with dose rates up to the permissible limits (0,05 Sv per year) may result in "stochastic" effects. (Non-stochastic effects occur only if a substantial threshold dose is exceeded). For stochastic effects only the frequency is related to the dose not the severity of harm. The main effects may be

- induction of leukaemia and other types of cancer
- genetic effects
- effects upon an embryo or foetus

The further considerations are only concerned with the risk of cancer, which is likely to be the significant component, except in the developing embryo, at dose levels received occupationally.

Usually the risk of radiation exposure is measured as the number of additional malignant disease in a population, in addition to the usual incidence of cancer, related to the size of the population and the sum of individual doses. Radiation exposure risk data are always estimations of the hypothetical consequences related to exposition data whereas conventional risk values are based on accidents and diseases, which really occurred.

The physical, chemical and biological processes in the human body during an after radiation exposure are not fully understood up to now. Therefore simplified models have to be used for risk estimation. Any assessments of stochastic radiation damage are extrapolations of observations of the consequences of radiation exposure with very much higher doses. A huge number of competent publications are concerned with the topic and it was not the intention of our study, nor our competence, to present our own new estimations of the risk of radiation exposure. Instead we used the evaluation of the ICRP published in the reports ICRP 26 and 27. From this source we adapted a risk value of all total malignancies of 100 per million persons exposed and 0,01Sv, which is a mean risk value of 10^{-2} Sv^{-1}.

When considering this risk value for fatal malignancies for comparison with conventional risk it might be important to consider additional aspects of such harm:

- the period of illness prior to death
- the period of illness or disability, including any operation, in non-fatal cancers.

Table 1 Risk values for the consequences of fatal or non-fatal cases Lost (RLYOL) resp. Impaired (RIYOL) years of life (expectation) per 1000 employees and year

	Profession or	**Code**	**RLYOL** [1]			**RIYOL** [2]		
No.	**Group of professions**	**of prof.**	**OA**	**TA**	**OD**	**OA**	**TA**	**OD**
1	"Mining"	07...	19,9	2,9	7,5	106,2	7,0	60,7
2	Miners	71	23,1	3,7	10,4	126,7	8,1	78,4
3	quarry workers	81	31,0	2,7	9,8	44,6	8,6	46,7
4	"Building material"	09...	5,8	4,7	2,2	32,1	7,4	5,2
5	"Chemical Industry"	141	3,0	3,0	1,5	13,4	6,2	2,0
6	"Metal processing"	24...	4,2	3,2	0,2	27,1	6,8	2,1
7	metal worker	27	5,4	3,3	0,2	33,0	7,0	2,6
8	Mechanic	28	2,5	4,3	0,1	11,2	9,5	0,2
9	"Electro"	311...	4,8	3,1	0,1	14,9	5,3	0,3
10	Electrician	31	5,0	3,5	0,1	16,7	6,5	0,4
11	"Construction/Building"	44...	10,5	4,2	0,1	50,0	7,2	1,5
12	Roofer	452	23,6	4,0	0,0	87,4	7,1	0,5
13	building labourer	47	12,1	5,1	0,0	53,2	8,0	0,7
14	Painter	511	4,6	3,3	0,0	29,6	6,4	0,7
15	"Energy operators"	541...	5,4	4,2	0,0	19,8	6,5	2,9
16	car/lorry driver	714	13,4	3,0	0,0	33,9	5,1	0,1
17	"Shipping"	721...	56,5	1,0	1,6	65,1	2,3	0,2
18	"Office"	77...	1,4	2,0	0,0	4,4	4,3	0,0
19	"Health service"	853...		1,3	0,0	1,2	2,2	7,5
20	room cleaning	933	3,1	1,2	0,0	11,0	5,8	0,0
21	machine cleaning	937	4,3	3,1	0,0	17,6	3,7	2,5
22	"all occupations"	071...	4,4	2,8	0,3	20,6	5,4	2,0

[1] RLYOL: Risk of the loss of life expectation due to fatal cases
(Life years per 1000 employees and year).

[2] RIYOL: Risk of impaired (loss of healthy) years of life due to non-fatal cases

OA:	Occupational accidents
TA:	Travel accidents to and from work
OD:	Occupational diseases

1.5 Comparability of fatal accidents, injuries and diseases.

Lost life time as an index

The risks to health of conventional occupations are mainly accidental death, injuries with and without permanent disabilities and occupational diseases. An index of harm, which should allow comparing the consequences of different risks, must try to integrate the various dimensions of harm in one figure.

It is obvious that the widely used death rate does not include considerations of the relative importance of different types of disability or disease. But even comparing death from accidents with delayed death from disease raises the questions how this can be taken into account in an index of harm. Furthermore it should be considered at which age the death, which is attributed to a certain risk, might occur. Such considerations have lead the International Commission on Radiological Protection (ICRP) to propose the time lost from a full and normal working of life (in man-years per year per 1000 employed.) as a possible index of harm (ICRP 27).

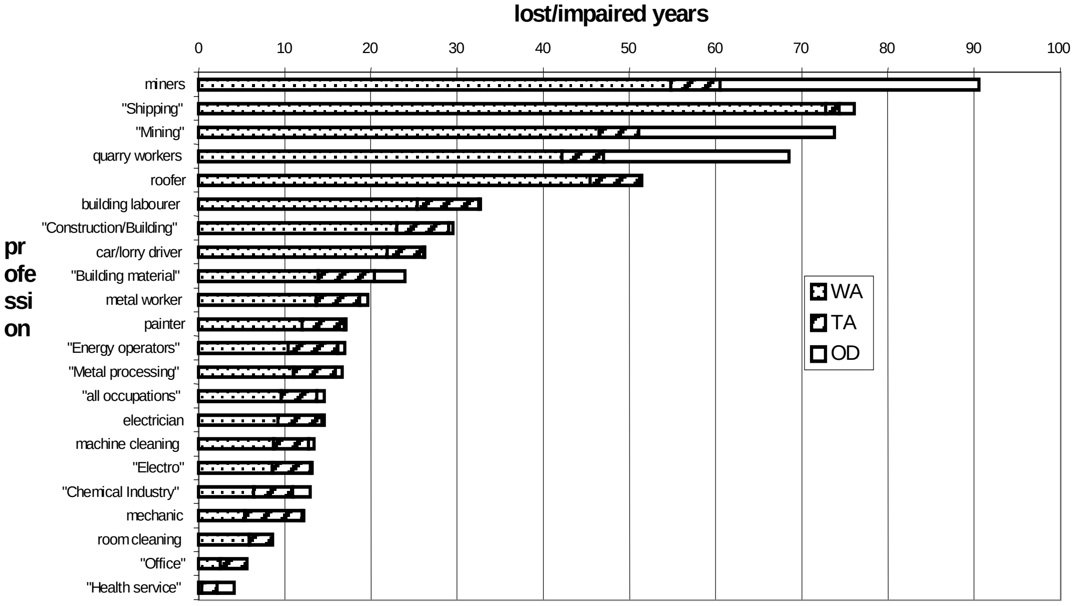

Figure 2 Index of harm (f=0.25)
WA=occupational accidents; TA=travel accidents; OD= occupational diseases

The advantage of this index is that it allows for considering e.g. that the age distribution of occupational fatalities is very different for the various occupations. Furthermore, the index allows comparing the mostly immediate consequences of occupational accidents with the time delayed consequences from diseases like cancer. Cancer is one of the main consequences of radiation, but may also be one of the consequences of other environmental pollution.

How to combine the various components of risk

If one accepts the time lost from a full and normal working life as a promising index of harm there is still a further question to solve. This question concerns the problem of combining e.g. the time lost in case of a fatal accident with the time of permanent injury or handicap as a consequence of a non-fatal accident, which might be handled by introducing weighting factors. This is the topic I want to stress here.

Are the lost years of live to be judged equivalent to years of working time due to an accidental injury? How shall we judge permanent disabilities? Shall a 100% loss of ability to work be judged equivalent to the lost lifetime in case of death or shall it be judged less severe by applying certain weighting factors? (A judgement of the years of an handicap by not affected persons as more severe seems to be expressed in the German saying – "Lieber tot als blind - better dead than blind". Such extreme judgement might also be supported by the suicide of some victims).

Furthermore, there remains also the question if injuries leading to a limited period off work shall be included in the risk index at all. There is also the question if this index shall be used with the risk from radiation, because the dominant risk component in that case is the possible induction of cancer leading to death.

1.6 Occupational diseases

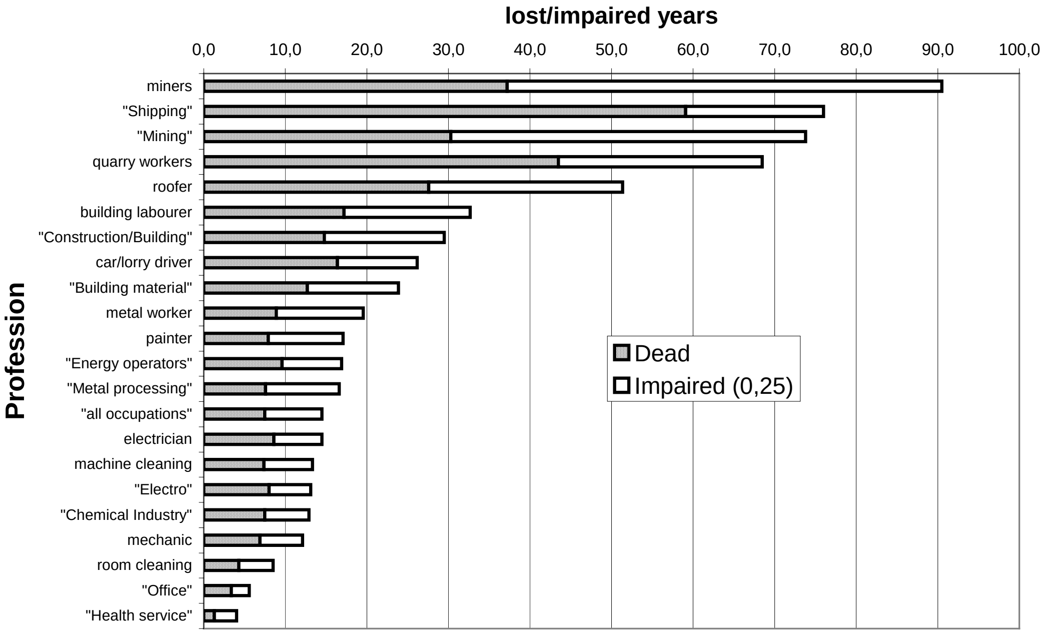

Figure 3 Index of harm (f=0.25)

The proposed index of lost years from a full and normal live allows for the inclusion of occupational diseases, because the lost life time for fatal diseases can be calculated from the data of the Berufsgenossenschaften as well. In principal, diseases leading to permanent disabilities may be treated like injuries from accidents based on the percentage of reduced ability to work, which are contained in the statistical data. But in fact, it is very difficult to calculate valid data concerning occupational diseases. In contrast to accidents, which can be attributed to a certain point of time, the occupational disease is the consequence of long time effects. Therefore, data about earlier times of occupations should be available to calculate correctly. Furthermore, the progress of work safety and health programs over the years has changed the spectrum of health dangers over the years. Thus, there is a qualitative difference in what can be deduced from data about occupational diseases.

In fact, only in a small number of industries (e.g. coal mining) occupational disease lead to major contributions to risk when compared with the risk of accidents. Further more, this is the case for occupations where disease is a remarkable part of overall risk (like in health care occupations).

The questions touched here demonstrate, that an index of harm combining different dimensions of harm in one figure may never be satisfactory for any purpose. From a political economy point of view it may be adequate to judge the loss of production days in the same way, independent of which cause contributed to them. In the traffic area the equal judgement of lost time (as a cost of loss of resources) is applied, in order to define priorities for the application of resources for traffic safety. In the same way, this may be adequate for the definition of priorities in work safety matters.

But an index of harm comparing the harmful effects of radiation exposure with the risk in conventional occupations should try to consider the different perceptions of risk, which different persons attribute to different harmful effects. A judgement of the time lost from a full normal (working) life in case of death, time of severe disease or the disability to work with a permanent physical handicap will be quite different depending on the perspective of different persons like the affected worker, his family, the occupational insurance or from the point of view of the political economy.

1.7 Considerations to combine the different dimensions of harm

Earlier publications (ICRP 27; Solomon & Abraham 1980) estimated a stable relation between the lost time in case of non-fatal accidents and the frequency of fatal accidents. This assumption, caused by the lack of adequate data, could not be proved with our data. Furthermore, we could not confirm the hypothesis that the fraction of accidental death increases more rapidly than that of all accidents with increasing hazard. Instead, we found from our data that the fraction of fatal accidents varies over the various professions to a very large extent. If follows that the calculations of a hierarchy of risk professions/jobs depends on the weighting factors used for the impaired life time in combination with the lost life time.

The data concerning permanent disabilities from accidents and from disease available to us measure the severity of accidents and diseases by the percentage of disability. This seemed to us to offer a better basis for the calculation of impaired time compared to other estimations used earlier.

The considerations above define important input data to calculate an index of harm in a much better way than earlier available but the central question of adequate weighting factors for combining the risk of injury with the risk of death remained.

1.8 Constructions of an index of harm

The key question for the construction of the index of harm combining the different dimensions of harm is – as pointed out above – the definition of adequate weighting factors. Previous studies mentioned above used a weighting factor of 0.1 to combine the impaired time of live with the lost time for fatal accidents (ICRP 27), other used the same factor but in some examples also a 1:1 relation (Solomon & Abraham 1980).

One possibility to achieve such weighting factors might be, to ask the persons which were affected from occupational accidents and suffer from the consequences. We were not able to perform such a study, which should include representative cross sections of people with different forms of handicaps, age, sex, professions etc. But we propose such a study to solve the judgement problem for further discussion.

In our study we found no further arguments to propose a privileged weighting factor. Instead, we presented combined results for the index of harm using different weighting factors. The factors applied for the impaired lifetime were 0.1; 0.25; 0.5 and 0.75.

To sum up, the index of harm should represent the frequency as well as the severity of the harm. As measures of severity we derived till now:

- for fatal consequences:
    - the time lost from a full and normal life
- for non-fatal consequences:
    - the period spent off work
    - the period of hospital treatment
    - the percentage of permanent disability to work
    - the impaired years of life, due to full or partial permanent disability

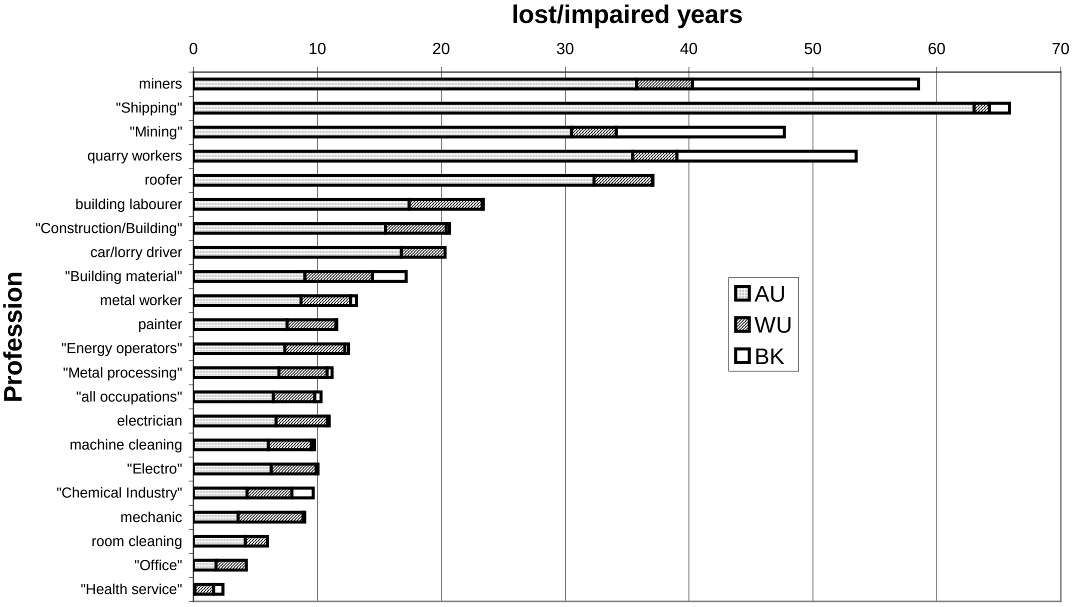

Figure 4 Index of harm (f=0.1),
AU=occupational accidents; WU=travel accidents;BK=occupational diseases

Reliable detailed data are not available for the period spent off work for all accidents with out permanent disabilities and for the period hospital treatment for the same class of accident. But the overall figures show that the time off work and the period of hospital treatment per 1000 employees are small (265 days) by comparison with the impaired years of life with full permanent disability per 1000 employees (about 14,2 years). So it is justified to neglect these upper two dimensions of the non-fatal consequences.

The index of harm may then be calculated summing up the following lost/impaired years of lifetime components in the following way:

- fatal occupational accidents
- non-fatal accidents x percentage of disability x Weighting Factor. (LNFA*F)
- fatal occupational disease
- non-fatal occupational diseases x percentage of disability x Weighting/Factor)

Table 8 presents the basic data of lost/impaired years of life, which we have calculated from the rates of occupational accidents (OA) and travel accidents to and from work (TA) and occupational diseases (OD) taking into account the age distributions of the cases.

The figures show the hierarchy of risky jobs/professions for different weighting factor applied to remaining permanent disabilities. In the figure 2 we applied a weighting factor of 0,25 and the professions are sorted according to the risk level achieved. The figure indicates the contribution of occupational accidents, travel accidents and occupational diseases.

Figure 3 presents the same result for the factor of 0,25 but it differentiates the fatal /non-fatal consequences of the accidents and diseases.

In figure 4 we applied a weighting factor 0.1 but the sequence of professions is the same as in figure 2.

We conclude that the calculated risk figure depends to a very large extent on the weight, which is attributed to the non-fatal consequences of the harm.

1.9 Discussion

In our study we tried do develop an index of harm for well-known risks using data about cases, which really happened. But even in this situation we found it very difficult to recommend a way to reduce the number of consequences (or dimensions) which any risk may imply to only one figure. But just this reduction of dimensions is necessary to enable us to compare different risks.

The task of risk comparison may be even more difficult if e.g. environmental risks, which cause probable consequences in the far future, shall be included in a risk assessment. Similar difficulties may result from the fact that some risk can only be estimated in a very rough manner and the risk dimensions/consequences may be not visible at the moment.

2. LEARNING FROM EXPERIENCE. (J. PARIES AND CORINNE BIEDER)

How do we recognise early signals? How do we detect what is going to hit us next? Risk is often specified as an average of many instances e.g. fires in houses, but in reality a house is on fire from a specific cause. If only one could know in which house out of the population of all houses the chance of a fire went to 1!

In all modern industries coping with high intrinsic risk (e.g. aviation, nuclear electricity, chemistry, rail transportation) contemporary accidents are rare, but the residual ones are often catastrophic. It is not socially acceptable to merely wait to analyse accidents in order to learn from them how to improve safety. Among complementary prevention sources, the use of feedback from operational experience concerning incidents, failures, errors, deviations, and other anomalies of daily operations, has been established – at least formally- as a standard for many decades, particularly in aviation.

Using incidents to prevent accidents seems a reasonable and easily attainable goal. From an incident to an accident, the path seems to be shorter, more visible than from a normal situation to an accident. Hence, incidents are considered to be forewarnings of accidents. Incidents often look like accident embryos. We expect to see accident precursors revealed by incident analysis. And the recurrence of similar features in incident scenarios, or similar causes behind them, should reveal unknown weaknesses in the system's safety, and allow for some proactive correction.

In the eighties, the emergence of cheap computers facilitated the management of large amount of data and boosted the development of incident databases. During the last twenty years, mandatory incident reporting systems, and penalty free confidential reporting systems have flourished within industries like aviation, nuclear electricity, chemical industry, and the like. A continuous flow of reports has been feeding an impressive data storage. However, transforming all these data into useful safety information turned out to be a different story. The illusion that detecting flaws in the system will be spontaneously brought by the power of a computerised database is not uncommon. But as a matter of fact, many safety analysts collapse under a mass of data they are just unable to process properly. And the computers remain silent about where the problems are. Unfortunately, while it is rather easy to take the "free lesson" of a quasi accident unfolding its scenario up to the edge of disaster, it turns out to be much more difficult to detect and exploit the potential signals embedded in the background noise of everyday variability.

Examples of patent failures of the operational feedback process are numerous. The investigation into the Mont Sainte-Odile air accident (January 1992) disclosed the existence of at least four precursor incidents, involving the same scenario of autopilot vertical mode error. The problem was not really to guess what the potential next step in the scenario development was like. All qualified crews, all flight safety experts could have described the deadly potential of that scenario. However, the air transportation system as a whole –airlines, manufacturer, and authorities – failed to react, apart from some awareness action in training. More recently, the Concorde accident near Charles de Gaulle Airport had also been preceded by several incidents, of which at least two were very serious. Such a failure to react to what seems to be, with the benefit of hindsight, an obvious warning is not the privilege of the aviation domain. The nuclear accident of Three Mile Island in 1979 was also a follow-up of two serious incidents involving the same thermo-hydraulic phenomenon. A report describing the main lines of such an accident had even reached the power plant some time before the disaster. Here again, the "system" did not properly react. It is worth asking why.

A common and simple answer is to blame the economic cynicism of an industry in which everyone's eyes are turned on profit charts, and which refuses to react to yet crystal-clear signals. The contention of this paper is that while economy - or other interests like the public image - obviously plays a central role in such decisions, the problem is less "cynicism" – i.e. insensitivity to safety concerns generated by productivity obsession – than acuteness. It is not so much that the industry is deaf; it is mainly that safety signals are muffled, in spite of all the obviousness that will inevitably appear after the accident. In other words, the contention is that there is a real difficulty to see the wood for the trees, to elicit the relevant safety lesson from low-level events. The proper lesson most often remains a "weak signal".

Indeed the question is how to extract a safety lesson from a series of events. Not all events are valuable. In most airlines equipped with an effective operational incident reporting system, the proportion of reports that actually bring something usable in terms of flight safety is about 5%. Like in gold mining, someone has to sift through the dirt, and wash it away to find some gold dust. And nuggets are the exception. Then the critical issue is not so much the size of the sifting pan, but rather the capacity to discriminate a piece of gold from a bit of fool's gold or gravel. Most efforts invested in incident reporting have increased the size of the sifting pans and the number of miners. Considerably less attention has been given to the 'serendipity' associated with the process, in other words to this insight, this ability to "see the signals", and to be really surprised by one's discovery.

Additionally, there is some innocence in the way the challenge is perceived. Most current approaches to "seeing the signals" proceed from the same assumption, which could be called the "concatenation paradigm", and has been endlessly illustrated by the pyramid metaphor. The pyramid metaphor is based on a graphical illustration of a statistical "truth": the number of incidents is a decreasing function of their severity. However, the concatenation paradigm goes beyond that simple statistical acknowledgement: it strongly suggests that the space of risk is a continuum, that minor incidents and severe accidents have basically *the same causes*, although may be in different numbers. The belief that the causes of major and minor safety events are the same is very strong among safety analysts, although it has been criticised by several authors, and even labelled an "urban myth" by Andrew Hale (2000). The implicit understanding of the frequency/severity relationship is that when more causes combine their effects, more damage will result, although less frequently, as the probability is then a combined probability. In other words, risk results from the action of elementary "risk quanta", and accidents result from the improbable combination of numerous quanta. Consequently, reducing the frequency of minor incidents will reduce the probability of major accident. Actually, one can reasonably argue for some degree of overlap in the causal fields of minor and severe occurrences. However, it will be this paper´s contention that there is also a *difference in the nature* of these respective fields of causality, that makes extrapolations from minor to severe occurrences for prediction purposes particularly "touchy".

2.1 Weak signals: species to be defined

In brief, we need to better understand why it is a real difficulty to elicit the relevant safety lessons from low-level events, or why these proper lessons most often remain "weak signals". One therefore first need to elaborate on the notion of weak signal. The following discussion will be restricted to "weak signals" associated with incident analysis. The expression "weak signal" will be used to mean any event, or phenomenon, bearing risk related information, although in such a subtle, remote, indirect, implicit, or hidden manner, that the clue would escape normal perception by individuals and organisations. To understand how an event or a phenomenon can become a weak signal, it is worth reminding that an event is intrinsically neither weak nor a signal. Its potential status of being a signal, strong or weak, results from its connection to a representation of risk factors, itself an outcome of a more or less elaborated and comprehensive model of risk and safety in the "world" of interest. Consequently, there are several categories of reasons why a signal can be "weak":

- A first category stems from the non-linearity of accidental processes in the world itself. One strives to predict or anticipate a phenomenon by monitoring the birth and the growth of its "causes". But most of the time there is no symmetry between causes and consequences, for several reasons.

- It may be first that we are dealing with divergent phenomena. The propagation of a crack in a strained piece of metal is a good metaphor of that. For a long while, the crack stays at a microscopic scale and remains invisible to normal observation. Then it develops to a size that makes it visible if searched for, while it does not bear any consequence. Finally, its growth speed increases considerably, and eventually a rupture – an accident – occurs very soon. We have a "weak signal" if the crack remains invisible, or very difficult to detect, while the crack development already reached the accelerated phase. We also have a weak signal when the crack is visible, but its propagation law is unknown, and expected to be linear. It is always extremely difficult to anticipate non-linear processes. This category of weak signals corresponds fairly well with ageing processes. They include a slow accretion or percolation of micro-deviations that suddenly reach a threshold at which new properties suddenly emerge – or at which existing properties are suddenly lost. This generates an important discontinuity in a system's behaviour (here, its safety). Signals are "weak" as long as one cannot monitor the margin to the threshold.

- Second, it may be that we are dealing with combinatory phenomena. Most of the accidental conditions in a complex system involves the improbable combination of common failures, individually harmless for the system (Rasmussen 1997, Reason 1997). Then the signal is "weak" because what we see is a random distribution of insignificant events, with a high individual occurrence frequency, while what makes sense in the safety space is their combination, and only their combination: none of them taken individually has any real harmfulness.

- But it is very difficult to see combinations and coincidences building up. The prediction of these combinations implies associative, imaginative and inductive capabilities that are the privilege of human intelligence and expertise only. These capabilities are considerably weakened when the volume of data exceeds the capacity of human memory, and when a computer-based data processing is used. Indeed, apart from the storage of a narrative section, the very principle of coding an event report into a database, through context and causal fields, breaks the meaningful concatenation and dismantles the “accidentogenic” construction initiated during the incident. It is then particularly optimistic to expect that processing the data originating from a population of events will reveal these constructions.

- More generally, a signal can be “weak” because we are facing the limitations inherent to the notion of cause itself (Hollnagel 1998). This notion implies that the same conditions invariably produce the same effects. Unfortunately, in a complex system, extremely close conditions can easily produce extremely different effects. As soon as human actions are involved, the definition of a causality link between an event and these human actions is a complex sociological and psychological judgement process. The notion of error itself is representative of that difficulty. If it is defined with reference to safety, then the link to risk is a tautology (something like “actions (or inaction) that can potentially or actually result in unsafe situations can cause an accident”), but the link to psychology is disputable. If it is defined with reference to psychology (intentions), then the link to psychological mechanism is real, but the link to safety is disputable. In the *post hoc* analysis of an incident or an accident, the knowledge of the final outcome – the benefit of hindsight - overshadows the complexity and the partially arbitrary features of the of the action control process (Woods et al 1994; Decker 2001A; 2001B). Furthermore, the causation models that are used are most of the time linear (one event or fact produces an effect, and so on). As noted by Rasmussen (1997), “one major difficulty in the use of linear causal reasoning is that it is unreliable for analysing the behaviour of systems including closed-loop, feed-back functions. In that case, linear causal reasoning becomes circular”.

- Any attempt to predict the future from a set of present conditions implies a perspective reversal that reveals all the complexity of real time control processes. Future is open, and about everything may happen. In the case of dynamic situation control (e.g. flying an aircraft; controlling a nuclear plant), a fundamental process interposes itself between the situation parameters and the associated (future) risk: the dynamic management of risk by front line operators. In their control activity, these front line operators do not strive to detect and correct all errors, anomalies and deviations, not even a combination of them. They do not try to conform to a unique way of doing things properly. They act and react according to their mental representation of the situation, mainly according to their perception of current and future risk (Amalberti & Deblon 1992). They use the emergence of cognitive signals perceived at the approach of a loss of control – among which their own errors - to adjust their margins and recover risky deviations (Rasmussen 1997; Wioland & Amalberti 1996). Consequently, while what is stored in the database as a signal of an accident course mainly consists of deviations from a “right” course of action, the relevant signal would be that of a beginning of a loss of control, e.g. a mismatch between perceived risk and actual risk.

- Signals can also be “weak” because they are swamped by background noise. A metaphor may help to better grasp the problem. Using incidents to predict and then prevent accidents is like gold mining. Not all events are valuable. Someone has to sift through the dirt, and wash it away. Tons of mud must be filtered to find some gold dust. Nuggets are the exception. Questioned in 1999 about the percentage of incident reports that eventually turn out to bear usable safety information, the safety managers of ten major airlines around the world gave answers averaging around 5% (Pariès et al 1999). Then the critical issue is not only the size of the sifting pan, but rather the capacity to discriminate a piece of gold from a bit of fool's gold or gravel. Most efforts invested in incident reporting have increased the size of the sifting pans and the number of miners; considerably less attention has been given to the discriminatory ability. Arguing that we don't really know what we're looking for until we see it, there is now a gold rush mentality - incident reporting systems have become an unstructured 'report everything' panorama of possibilities. The uncontrolled growth of incident reporting practices can then lead to the paradoxical result of burying useful information even further in a mass of irrelevant detail.

- Finally, a signal can be “weak” because it is inconsistent with the framework or the parameters of the organisational decision it is expected to trigger. The difficulty that safety issues recurrently meet to force their way through managerial decision making processes has been described on many occasions. In general terms, such decisions imply at some point a choice between the reasonably probable benefit of a productive investment and the remote probability of an accident loss. The prospect theory (Kahneman & Tversky 1979) anticipates in this case a strong decision bias in favour of production. The central objective of all productive organisation is to produce (build aircraft, transport passengers), to make profits and to stay in business. Safety does not clearly generate profits, at least not on a short term, local – e.g. one airline - perspective. Safety is a constraint. This does not necessarily imply that safety and productivity must be antagonist. Many of the features that make an efficient organisation also make it a reliable and safe one. However, and particularly on a short term perspective – and the managers' assessment time span is short term, typically one to three years - safety has a cost. Someone must pay for simulators, training time, life jackets and the like. Managers therefore need to make decisions that imply to find a compromise between production and safety. By their very nature, these decisions are influenced by three main factors, inducing specific bias tendencies:

- The random component of accidents: Only a part of the safety equation is under the control of managers. This means that even the best safety policy does not protect against "bad luck" strikes. Conversely, considering the very low frequency of accidents, even risk taking organisations can evade accidents for a long time. As Reason (2000) puts it "luck works both ways: it can afflict the deserving and protect the unworthy". This is not very motivating.

- The obscurity of outcomes: In safety issues, only failures are visible. There is no accepted indicator of the safety level of an organisation. Intuition can be very misleading. For example, more incidents do not necessarily mean lesser safety. They may, as well, reflect a better efficiency of incident reporting. Or even more challenging: they can feed and maintain a better efficiency of the "immune system" defences. In the aftermath of Valujet Airlines crash in the Everglades swamp near Miami, Florida, (May 2000), it was revealed that an FAA report finished just days before the disaster showed that Valujet had an unusually high rate of nonfatal accidents and incidents. The FAA was blamed for inaction, with the good sense assumption that higher rate of such mishaps means a greater risk of disaster. Arnold Barnett, a risk management professor at MIT, was asked to assess such an apparently obvious, and broadly shared assumption. Pr. Barnett found (www.gwu.edu/~csm/aviation/track_i/barnett.htm) that the coefficient of correlation between passenger death risk and nonfatal mishap rates among major US jet carriers during the period [1/1/90-3/3/96] was –0,10 for incidents and –0,34 for serious (non fatal) accidents. In purely statistical terms, this negative correlation means that passengers would have reduced their flight risk by preferring incident prone airlines, and even better: nonfatal accident-prone ones. The discussion of this interpretation will not be conducted here. It is sufficient to say that using frequent event data to make judgements about passenger risk in a particular airline is lacking empirical support.

- The absence of direct causality: The relationship between preventive measures and the expected benefit is unclear. Who can pretend to quantify the safety benefits brought by an additional training? In a system as complex as a large industrial plant, causality is circular rather than linear. Implementing "more protections" can induce more confidence and even excessive confidence and lead to a final balance on the opposite side from the expected one.

In summary, production is the main internal objective of productive organisations. Results of investments can be modelled and anticipated, causality is clear, outcomes are expected on a short term, within the managers' time span, and they are rather fairly distributed - they award the deserving. Success is visible, and reasonably reliable instruments are available to assess it. Unlike production, safety is an external constraint. Results of safety investments are rather uncertain; causality is obscure, with a strong random component. Benefits can only be conceived on a long term and global basis, and only the failures are really visible. Clear and reliable assessment tools are lacking.

All this makes the task of managers in charge of safety pretty difficult when it comes to talking to their bosses. To be able to exist, to pass from a state of expert-owned information to a state of a problem on the senior decision makers' agenda, safety issues, as any problem in the organisation, need a sponsor, a promoter (Bourrier & Laroche 2001). From this perspective, potential safety problems promoters are particularly badly served by the features of safety issues. There is by essence a gap between their discourse and decision makers' natural language syntax and semantics. They talk about something that has no positive perception, but demands permanent efforts. As Weick (1987) put it, "safety is a dynamic non event". To be really incorporated into the organisation´s decision-making process, it should become a dynamic event. How can the "signals" generated by traditional incident reporting and analysis allow this goal to be met?

2.2 Signals and safety model

At this stage, it is worth reminding how feedback from operational experience actually works. The traditional way to try and extract safety lessons from operational experience through reported events analysis is a three-fold approach (after Koorneef 2000):

- The first strategy – that we will call the *clinical* approach **-** is like a small-scale accident investigation. It seeks risk management strategies through the reconstruction of the “causal” path to the incident, from direct causes (unsafe acts committed mainly within the time span and space of the incidental event) back to organisational influences. For example, a runway incursion was caused by an unclear clearance, itself caused by a poor phraseology, itself caused by a poor training, itself caused by a poor manpower management, and aggravated by a workload peak, itself caused by poor workload anticipation by the team, and so on.

  The clinical approach limitations are well acknowledged (Amalberti & Barriquault 1999). The causal attribution is an outcome of the *judgement* made by the investigators on the premises of their understanding of what can cause accidents, as well as an outcome of the organisation’s safety model embedded into the reporting/analysis format. Then the definition and attribution of “causes” reflect the thinking of the analysts and his/her organisation about safety. Additionally, inter-rater reliability is generally very low, in spite of all harmonisation training efforts. The extensive use of key words does not really fix the problem, while it brings additional inherent limitations: key words are by nature inflexible and restrictive. Furthermore, the causation models that are used are most of the time linear (one event or fact produces an effect, and so on), while the use of linear causal reasoning is unreliable for analysing the behaviour of systems including closed-loop, feedback functions.

- The second strategy is called the *epidemiological* approach. It builds on a compilation of the available clinical analyses. Indeed, as far as incidents are concerned, organisations generally need more than one event to undertake costly actions to reduce causal recurrence. Some accumulation of evidence is needed to support costly decisions. Consequently, individual incident analysis outcomes, including the “causes”, as attributed by the analysts, are generally stored in a database, then trend analysis is performed. One seeks for strong or repetitive patterns of “causes”. For example such trend analysis will “reveal” that 70% of incidents (of some kind) are the result of a deviation from Standard Operational Procedures, and that this percentage is growing.

This kind of data can assist in the identification of priorities, and the definition of prevention strategies. However, one must be very cautious in the interpretation of such “data”. These percentages, graphs, trends and the like are all but the result of a real statistical vision. The causal attribution in individual incident analysis is an outcome of the *judgement* made by the investigators on the premises of their understanding of what can cause accidents and with reference to the organisation’s analysis format. Hence the causality distribution shown by such data is a mirror of the analyst's and the organisation’s thinking, as well as a reflection of “reality” (Pariès et al 1999). Consequently, the prediction of risk based on such methodology tends to be a self-fulfilling prophecy: one sees what one is expecting to see according to one’s safety model in mind.

- The third strategy is the *statistical approach*: one seeks for correlation between descriptive factors (time, airport, phase of flight) and incidents. Theoretically, no causal model is needed to “see” a pattern of relationships building up between some of the parameters (e.g. un-stabilised approaches occur at that airport at a frequency above average when: i) runway 25R is active; ii) weather is fine iii) runway 25L is closed iv) type of aircraft is AAA v) departure airport is BBB; time is between 0500 and 0800 AM).

The limitations of the statistical approach in operational domains involving human operators are also well acknowledged. Conditions for statistical validity of an emerging correlation are difficult to meet. As heroes of their own story, reporters produce biased reports. We never know what can be kept invisible, so there is a lack of baseline data, and most of the time the representativeness of the samples referred to is highly disputable. Additionally, such an approach is not totally neutral. An implicit accident causation model determines at least the descriptive factors selection (if the pilots’ colour eyes is not felt to be a potential causal factor, it will never be mentioned) and the relationship pattern interpretation (only potentially meaningful patterns in terms of the safety model will be retained). Furthermore, organisations need a “causal explanation” before making costly decisions and acting. It means that once a correlation has been discovered, there is a need to make sense of it in terms of a safety model.

In all three strategies, the processing of information is of a ´bottom-up‘ type. It seeks to spot configurations that may be a symptom of high-level failure risk from either frequency relationships or causality relationships in the events contexts. In both cases a model of what can cause accidents, or symmetrically, a model of what makes the system safe, is involved. In the clinical and epidemiological approaches, it defines the influence of various factors on safety. In the statistical approach, it defines what parameters are to be monitored, and leads the interpretation of a detected correlation. The problem here is not that one refers to an accident causation model; it is that one does it *implicitly*.

The fact that the safety model is not used explicitly has some perverse consequences:

- The safety model itself is protected against feedback from experience. The safety model is taken as a truth. As it is not expressed, it cannot be falsified. For example, asserting that front-line operators' errors are or can be "causes" of incidents may look patently obvious, but it is actually an assumption, based on a specific model of system's safety, in which errors and deviations are fundamentally unsafe. Alternative models of safety can be suggested, in which the key issue is risk management and situation control. According to these models, errors are a normal component of performance accomplishment. Human performance normally varies (Hollnagel 2001), because it is globally more efficient to do things approximately and then adapt the action according to feedback from the real outcome. Particularly, human operators do not seek safety through the absence of errors. The front-line operators have a permanent real time risk management activity, not only directed towards their own erring ways, but also towards all the potential threats, coping with all the contingencies. That risk management process includes the monitoring of one's action (including error monitoring), but also planning, anticipating, selecting a strategy, assessing one's own capabilities to do the anticipated task, and the like. From such a perspective, unsafe situations might occur, not when errors are committed, but when the risk management process fails.

- On the contrary, the safety model tends to be self-confirmed. If errors are believed to be an obvious risk factor, many errors will be found in the database as cause of incidents. Indeed, because errors are committed in every flight, it will not be difficult to find errors during incidental flights, and to establish their causal status. As only incidental events are reported, there is a total lack of baseline data. Consequently, a naive idea of real operations develops, in which everything is consistent with specifications. This reinforces the tendency to understand incidents as a result of errors and deviations, seen as exceptions rather than common phenomema.

- The capacity to "read" experience is reduced. The safety model is the relevant filter to look at experience. As in any perception process, the selection and the structuring of stimuli available in the environment (bottom-up process) is driven by a scheme, a predefined order, a cognitive model of the world (top-down process), until an acceptable coherence between the perceivable and the conceivable can emerge. One can much more easily perceive what one knows. One can better see when one knows what is to be seen.

2.3 Developing an alternative approach

From the above discussion, it is a short step to the idea of reinforcing the "top-down" component of the experience feedback data processing. This implies to make the safety model explicit. Top-down approaches are commonly used in dependability methodologies: they start with the identification of unwanted events, then strive to identify all the combinations of failures and circumstances (e.g. causal trees) that can potentially lead to an unwanted event. This sets the scene for developing defence strategies.

The originality of the approach described hereafter is to apply such a top-down approach to operational experience feedback. The goal is to analyse incidents through pre-identified risk management strategies. The approach is currently under development in three main domains: aircraft design (Airbus); Air Traffic Management (Eurocontrol, HERA 2 Project), and Aircraft Maintenance (ADAMS 2 Project). It has led to the development of a software-based assistance tool for safety managers (modestly) called SMART (for Safety Management Assistance & Repository Tool). The software development is at the mock-up phase (2002).

The core idea is to confront incidents directly with the reasons why they were not meant to happen (in other words to the corresponding safety assumptions), in order to assess and map the strengths and weaknesses of these safety assumptions, as shown by the following diagram:

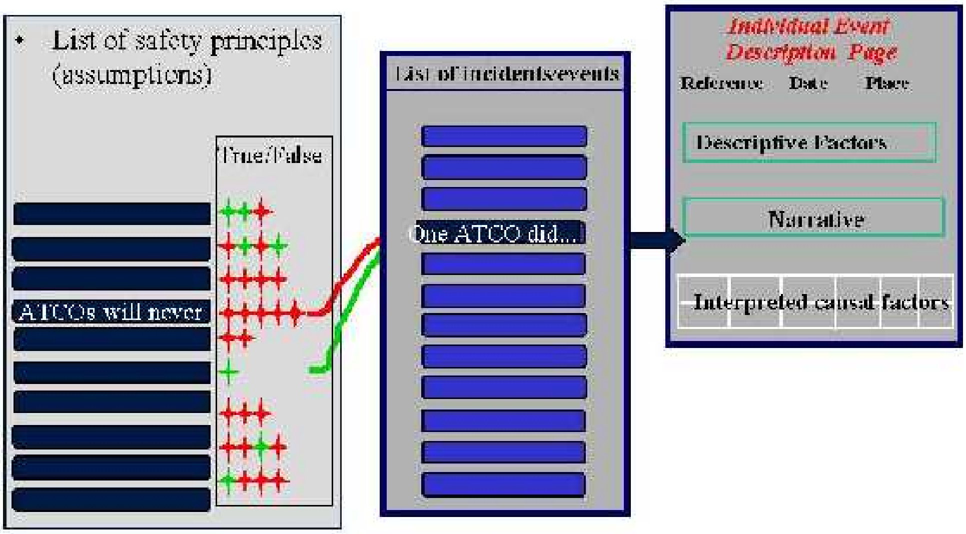

Figure 5 Basic SMART structure

One main goal of SMART is consequently to explicate – from an aircraft manufacturer perspective, or an ATM perspective, or any other perspective- the safety assumptions (called Safety Principles - SPs) supporting the safety of the system, in order to assess/challenge them through operational experience feedback.

However, the number of SPs meant to ensure the safety of a system like aviation is so large that it would be impossible to list them from scratch. It would also be impossible for an analyst to check all of them when analysing an individual incident. There is a need for a screening function to identify the relevant subset of SPs associated with a specific incident. Additionally, SPs do not work independently. They are linked by logical relationships: for example, a specific incident may be prevented by the existence of a procedure *and* the fact it is relevant *and* the ability of the crews to implement it, *or* the intervention of an automated system. This synergy must be represented.

For that purpose, the notion of “Generic Initiator“ (GI) is introduced. A Generic Initiator is defined as *any event (or non event) from which an accident would develop, should no specific recovery action be positively taken*. The underlying vision of safety is a dynamic vision. Following Weick’s words, safety is seen as a “dynamic non event”. It means that safety is not seen as an absence of unsafe events (e.g. errors, violations), but as the result of the system being under control, in a dynamically stable, intrinsically safe state. The different states of the system can then be represented metaphorically as shown by the following picture:

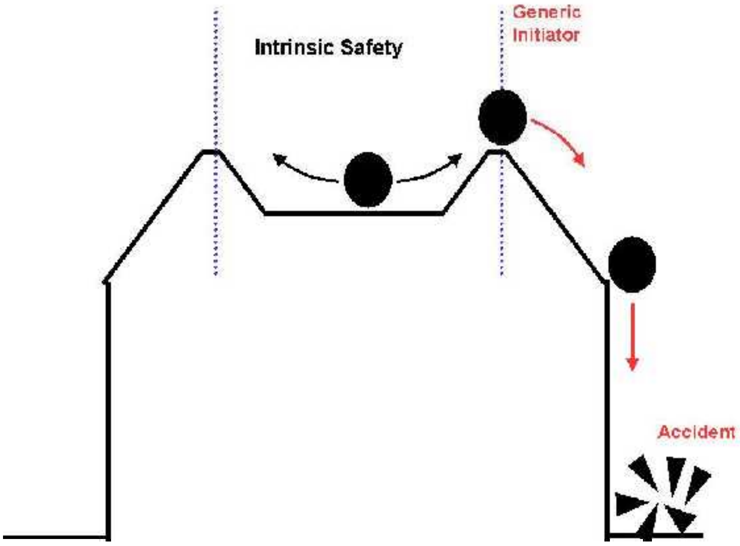

Figure 6 A basic model of safety

The idea conveyed by the concept of Generic Initiator is twofold:

- It is an **initiator,** which means that at one point, the system switches from a stable, controlled, intrinsically safe state to an unstable, uncontrolled, intrinsically unsafe state.

- It is **generic**, which means that it is independent from a particular instance or circumstance, a specific layout or piece of equipment or organisation, as long as these elements do not have a major influence on the safety management strategy. Indeed, a lesson can only be learned from an event if some generalisation is achieved. A significant part of the 'causality' of an event is context related, and this context will not repeat itself. The challenge to extract a safety lesson is then to find cross-contextual elements without losing the "sense" of the incidental story. As Rasmussen (1997) puts it, "completeness removes regularity". The solution suggested here is to seek the needed regularity through generic, prototypal paths to accidents. A prototypal path to an accident is a failure in the organised set of protections that are expected to prevent a specific type of accident.

In the next step, for each of the GIs, the Safety Architecture (logical combination of SPs) that is relied on will be developed. SPs can be first categorised according to their "distance to the accident":

- SPs intended for preventing the Generic Initiator from happening will be called *Prevention* SPs.

- SPs intended for preventing the Generic Initiator from developing into an accident will be called *Recovery* SPs.

- SPs intended for preventing the Accident from developing into its worse consequences will be called *Accident Consequences Mitigation* SPs

The development method starts with the "high level" safety principles. Then each SP is decomposed into a logical combination of lower level SPs, and so on. Consequently, the method goes from the most abstract (strategy) to the most concrete (expected behaviour of the system, its components and interactions). It may be helpful to see the successive levels of SPs as levels in a means-ends abstraction hierarchy[1].

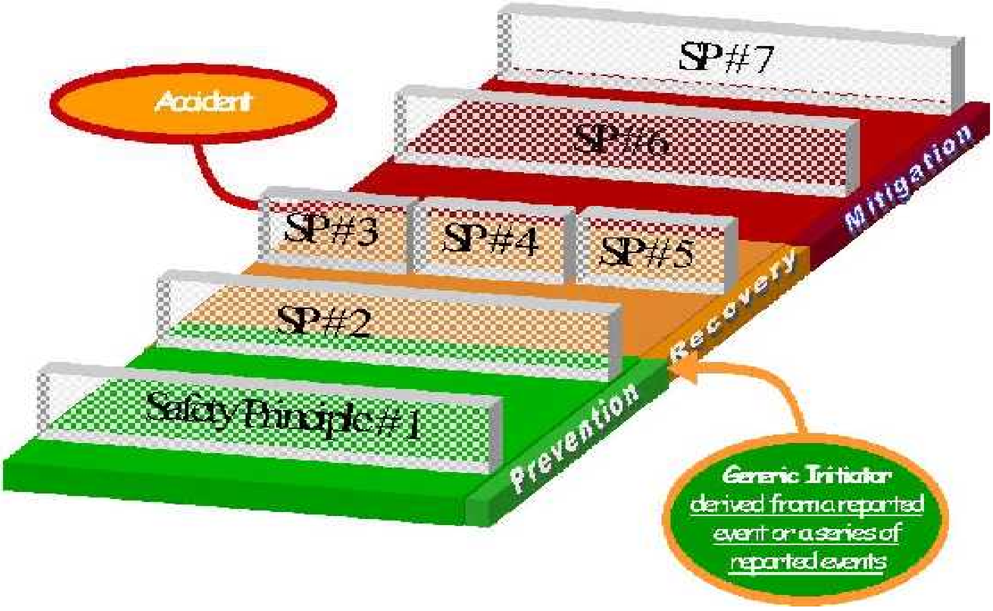

Figure 7 Safety Architecture (i.e. combinations of Safety Principles) associated with a Generic Initiator

The figure above represents a Safety Architecture at the highest level. Safety Principles on the same line are linked by an "and" while Safety Principles on different lines are linked by an "or".

[1] In practice it may turn out to be difficult follow such a means-ends abstraction hierarchy to closely. It should only by used as a guidance.

In the next step, each Safety Architecture will be confronted with the lesson of operational events. The first thing an analyst must do when analysing a new incident is to match it with a Generic Initiator. Once the incident is matched with a Generic Initiator, the Safety Principles to be considered (i.e. potentially called for during the event) are simply those involved in the safety architecture associated with the Generic Initiator identified. Assessing these Safety Principles then consists in going through the list of SPs, to point out and record:

- which of them failed,
- which of them were actually successful,
- which of them were called upon, but where behaviour is not available in the report,

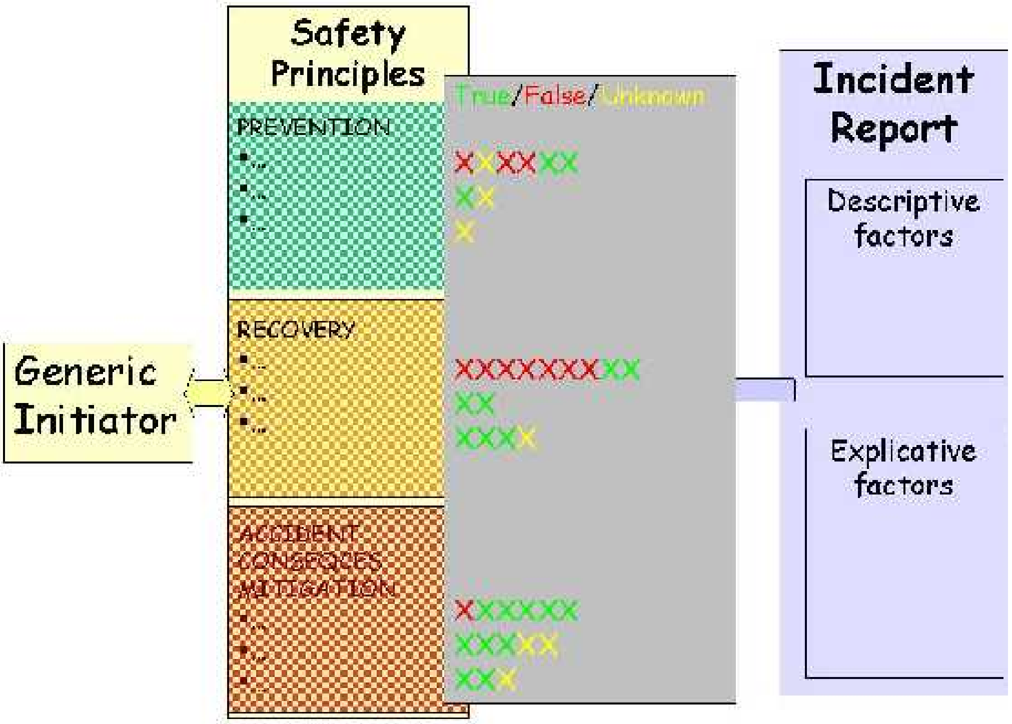

Figure 8

- and also, which of them were clearly not prompted. This recording will generate data about the prompting frequency of a SP.

The behaviour qualification (success, failure, unknown outcome**,** or not prompted) of the Safety Principles involved in one event translates into colour codes attributed to the link established between the corresponding SP and the event.

While the process of assessing SPs behaviour is repeated, event report after event report, the SMART system starts to accumulate the coloured links declared for each SP (be it success, failure, or called upon with unknown outcome). At this stage, Safety Architectures are no longer discriminated. If several Safety Architectures share a SP, and if links to event reports have been declared for this SP through several Generic Initiators (hence several Safety Architectures), then all these links will be gathered. The series of links associated to one SP will then form what could be called the “Health Map” of that Safety Principle, as illustrated below:

The next step is to interpret the "health map" of each Safety Principle to assess its empirical robustness, in other words, the level of trust that can be vested in it, on the grounds of the knowledge gained through an interpretation of the available feedback from experience.

Finally the consequences of the attribution of a certain level of trust to a specific Safety Principle can be assessed. Even though the “veracity” of a SP during a reported event is assessed through a Generic Initiator, the impact of its possible invalidation is assessed not only for this Generic Initiator in the specific situation of the event, but also in slightly different situations. In the long run, when all Generic Initiators and associated Safety Architectures are made explicit, the impact of a Safety Principle possible invalidation is even assessed across all the Generic Initiators involving that Safety Principle. This proactive exploration allows determining respectively the "local, "extended" and "generalized" criticality of the Safety Principle.

SMART therefore is a support to safety-related decision-making. It provides decision criteria derived from the criticality of unreliable Safety Principles, this criticality being determined on the basis of its role in the whole safety system.

In addition, because it “shows” the Safety Architectures associated with Generic Initiators, SMART may also help suggesting alternative solutions when a modification is decided upon, and help a simulated testing of the efficiency of these options as far as Safety is concerned.

For that approach to be fully efficient, all the possible Generic Initiators should ultimately be identified, and for each of them, the Safety Architecture explicated. Such a task could be achieved through an exhaustive top-down approach based on a functional safety analysis –then it would take a long time before the system can be used. The development could also be based on reported events: one could determine, on a case-by-case basis, the Generic Initiators associated with one reported event (or a series of similar reported events), and then develop the corresponding Safety Architecture.

2.4 Conclusion

Traditional incident analysis systems have shown their limitations. They are intrinsically dependent of the causal approach, and all the associated restrictions. The approach suggested in SMART fundamentally differs from traditional approaches in that it tends to describe why the system is supposed to be safe, instead of trying to understand why it fails. It seeks to learn safety lessons through a permanent comparison of expected and actual behaviour of the safety system. Beyond the "negative" experience, it also affords to build up information on what functioned in the way it was expected to. Additionally, it allows the building up of experience on all protection layers. Through the notion of Generic Initiator, the SMART approach is an attempt to go beyond the specific circumstances of an incident, while memorising the reasons with they could breach the system’s protections. All this information is enriched as reported events are analysed, through a growing and living database. Thus, the decision-making assistance proposed in the SMART approach benefits from the complete aggregated experience, rather than from a collection of single event experiences.

3. SYSTEM BOUNDARIES AND RISK TRADE-OFFS (TOM HEIJER, ANDREW HALE)

In complex technologies, particularly those concerned with transport, the risks of the technology affect different groups of people, the passengers, and the company employees operating the transport system, workers in ancillary functions such as maintenance and the public living around the transport infrastructure. This means that there are different interests represented in any risk assessment and in the decisions about risk control measures to be taken. The ideal is that the measures taken benefit all groups and are therefore not controversial. Many are, such as measures to increase the intrinsic safety of planes or trains, but some are not. This paper concerns itself with those cases where there is conflict of priority or of choice of preventive measures between two or more of the groups concerned. This is because they present specific issues of comparison of risks, questions of comparative standards of protection for the different groups, or problems of comparing different types of risk.

We present here two short case studies of such choices and examine some of the dilemmas of dealing with them. Our main purpose is to raise questions for debate, for which we hope to get new ideas, if not solutions, during the workshop.

3.1 Case study 1: maintenance on the track

The Safety Science Group of the Delft University of Technology was contracted to produce an opinion on an ongoing dispute between the Dutch Railways and the Labour Inspectorate. The dispute concerns the safety of track workers during large-scale renovation of the rail infrastructure. Because track renovation proceeds relatively slowly, the work, even on a line of limited length, takes a number of weeks. Closing the line to all traffic for such a period of time is unacceptable to the railway organisation, especially on those parts of the network that are heavily used for commuter traffic. Therefore, Dutch Railways proposed to carry out the works on one track while allowing train traffic on the other and using an enhanced safety regime. This meant that the work would have had to stop each time a train went past, in order to clear the line in use. Since timely detection of an oncoming train is vital to this scenario, a regime of “guaranteed warning” was proposed. This entails the use of hardware train detection devices and both visual (a colonnade of flashing lights) and auditory (110 dBA) alarms that provide a 40 second advance warning period.

The Labour Inspectorate however rejected this scheme on the grounds that the frequent evacuation of a part of the personnel (10 trains per hour) poses too great a risk for those workers. They demanded a complete stop of train traffic for the duration of the works despite the fact that the commuters on that line would suffer considerable delays, or, due to the reduced capacity, would have to transfer to other means of transport.

The Safety Science Group was asked to produce an opinion on the integral safety of several scenarios, ranging from work performed during a complete stop of all train traffic to continuation of both work and traffic without any speed limitation (=140 km/h) on the trains. The scenarios between these extremes involved continuation of train traffic with reduced speeds (80 and 40 km/h) in the vicinity of the works.

Approach

At first sight the position taken by the Labour Inspectorate seems almost unassailable, seen purely from its own perspective, which is to protect the workforce, i.e. the track workers (and the train staff). The passage of a train (with any speed) during the execution of the works will always pose a significant additional risk to the track workers. Therefore they considered cessation of all train traffic during the works the most reasonable choice despite the admitted "discomfort" to travellers. It is not the duty of the Labour Inspectorate to worry about the safety of other non-workers in the system. However, in a risk based approach to such decisions one must be careful to consider *all* effects on the risk of *the whole* population that is directly affected by such a decision. In this case, an important aspect was left out of the considerations: the effects on the risk of the commuters. They not only suffer delays as a result of the cancellations but at least some of the commuters will choose a different mode of transportation for the duration of the works. The alternative, in this and most other cases, will be to drive to work by road. Since the probability of individual accidents for travel on a motorway is about one order of magnitude greater than the risk involved in train travel over the same distance ($2e^{-10}$ per train kilometre vs. $2{,}1e^{-9}$ per kilometre per car), we should include this shift in the risk in our considerations, because this increased risk will be run by a considerable population. The approach taken to this problem was the following:

- Develop a simple fault-tree based model to represent the risks to the track workers *and* train passengers as a function of different exposure scenarios

- Estimate the probabilities of the basic events in the model in such a way that, if there is doubt about the figures, we are always certain to *over-estimate* the resulting risk to the track workers. This was because the risk to the track workers was the starting point of the exercise and we were being asked to provide an analysis for the Labour Inspectorate, whose 'clients' they are.

- Because continuation of train traffic during the works slightly raises the risk of derailment of a passing train, this increased risk has also been considered. Derailment threatens not only track workers but also passengers and therefore this risk of passengers being killed in a derailment has been included.

- Since most accidents between trains and track workers are fatal accidents, we then apply statistics of road fatalities to estimate the number of commuters that need to travel an equivalent daily distance on the motorway for the duration of the work to produce an equal number of expected fatalities.

Thus we obtain a criterion to decide for or against each of the scenarios: starting with the scenario preferred by the Labour Inspectorate (cessation of train traffic) we can calculate whether the number of expected road fatalities significantly exceeds that expected of the workers and train passengers. If so, we can move to the next best scenario until we obtain a scenario with a more balanced outcome. Because of the over-estimation of the risk for track workers (and passengers) we may be fairly certain that this procedure will not select the more dangerous scenarios for them too quickly.

Results

The fault tree model, depicted in figure 9, consists of 2 main branches: a branch representing scenarios where the timely evacuation of workers fails and a branch representing other collision scenarios e.g. control errors that put a train on the wrong track and scenarios in which the passing train hits objects in its path that subsequently hit workers. Without addressing all details of the model, there are some aspects worthy of note. One is that the probability of a person working in or near the track of the passing train has been set at 1: there is always someone working in that zone due to the location of controls on the machines used in this maintenance work. Therefore the probability of a man being hit by a train is only dependent on the success or failure of the evacuation procedures. This suggests a possible safety measure to be the relocation of the controls on the machine to the safe side! Secondly there is the problem of combining probabilities in the branches. Some of the risks to the workers depend on conditional probabilities: the probability that an evacuation fails when a train approaches (there is no evacuation if there is no train).

Other risks are more continuous risks e.g. the probability that a train is misdirected into the work zone by a controller. Furthermore some risks are contingencies: e.g. the probability that an object was moved into the track just prior to a train passage. To work around these problems all probabilities were converted to *probabilities per working hour*. On this basis the probability of the top event "any fatal accident" in figure 1 was estimated at 1.18e^{-5} accidents per working hour.

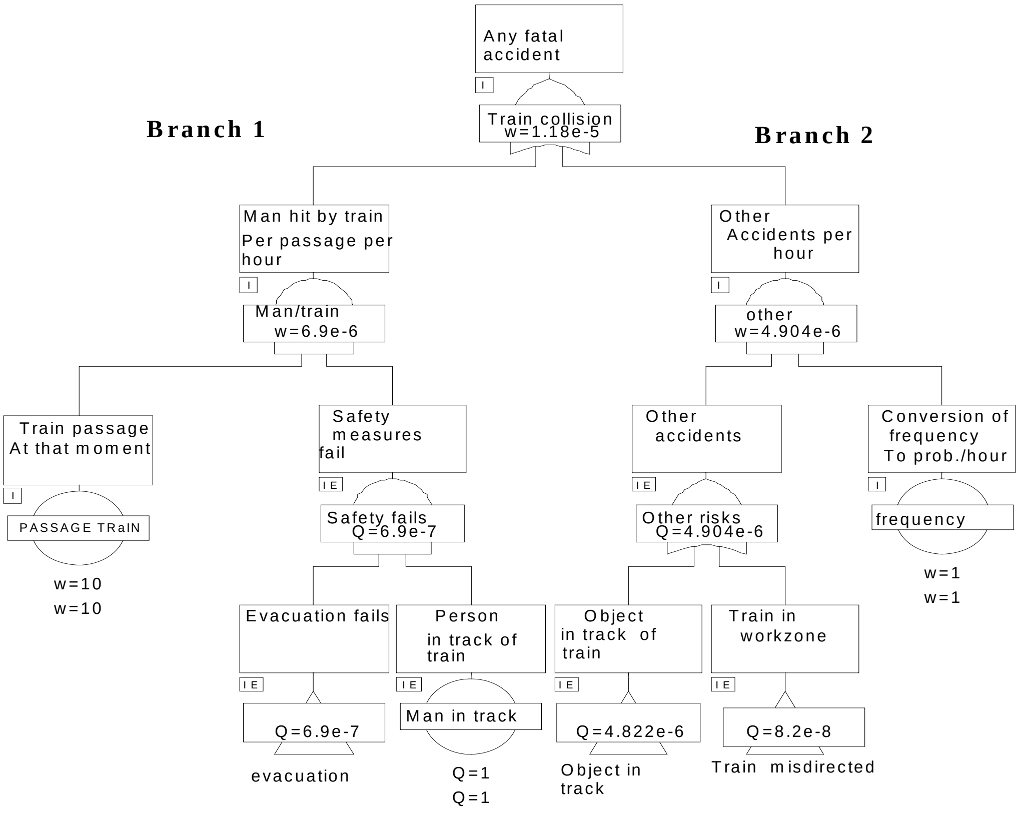

Figure 9: Fault-tree for fatal accidents in railway operations

Most train accidents in the Netherlands (e.g. collisions with track workers) do not involve passenger fatalities; only one high-speed derailment in the past 20 years produced 5 victims among the ca. 400 passengers. The number of fatalities is of course also dependent on the occupation of the train, which varies strongly over a day. Since no further data were available, we have simply postulated that any derailment of a passenger train will produce 5 passenger fatalities. Here, the probability of derailment has been estimated at 8,2e^{-8} per working hour, which results in a probability of 4,1e^{-7} fatalities per hour. The same basic probability of derailment has also been used to approximate separate risks for the environment that are caused by freight trains with toxic or flammable loads. Again to exaggerate the risk to workers, we assumed that dangerous substances would be released in 20% of all derailments of freight trains, killing all track workers (15 workers). This risk, estimated at 3,28e^{-5} fatalities per hour, is incorporated in the estimate of the top event.

On the other side of the risk balance we find the probability of fatalities amongst passengers displaced to road traffic. Again, in order not to favour dangerous scenarios for track workers, the risk figure for those passengers has been equated to the risk on motorways, which, at 1.9e^{-9} fatalities per kilometre driven (car kilometres), is 6-10 times lower than the risk on urban and rural roads.

Of course there is another alternative for road transportation of passengers: a bus. In the accident statistics in Holland, the risk for bus passengers does not differ significantly from that for train passengers. However this alternative was excluded from our considerations because the number of spare buses required to transport such a large number of daily passengers over an extended period of time are simply unavailable in the Netherlands.

To be able to do the mathematics of the risk balance for the actual situation, we need some more data about length and duration of the works. We have assumed, based on information provided by the railways, a length of line worked on of 48 km, and a duration of 40 working days with an average effective working time of 19 hours/day.

For the most dangerous scenario, one in which trains pass the work zone with undiminished speed, the calculations work out as follows:

Number of fatalities expected for the duration of the work:

- Track workers (branch 1) + track workers and passengers (branch 2) = 0,00995 + 3,28e^{-5}+ 4,1e^{-7} =0,00998.
- Number of person-kilometres to be driven on the main road to equal this risk: 0,00998 / 1.9e^{-9} = 525263
- Assuming that most travellers travel back and forth over the 48 km each day, this equates to 525263 / 96 = 54715 trips, and over 40 days this implies 54715 / 40 =1368 trips per day.

We concluded that the most dangerous scenario for the track workers yields fewer fatalities than if 1368 train passengers per day take their individual cars and opt for road traffic. Since the project in our case concerns one of the main branches of the Dutch rail network where over 10,000 passengers per day are transported, the number of 1368 will certainly be greatly exceeded. So, even for the most dangerous scenario for the track workers and in a calculation method in which their risks were over-estimated, probably with 1 to 2 orders of magnitude, the advantage of keeping the trains running is overwhelmingly clear. All other options that reduce the risk of the track workers without increasing the risk of diverting the passengers to the road are, however, welcome. We, therefore, stressed the fact that all possible measures should be taken to reduce the risk for track workers. Our considerations resulted in the following recommendations:

- Continue the train operations only with *passenger* trains with a speed that is reduced as much as is practicable: higher speed produces more severe accidents, but lower speeds significantly increase interruption time to the work per train and so extend the working (= exposure) period, beyond the 40 days.
- Since there are numerous large-scale renovation projects coming: develop a standard risk-based procedure to determine the safety regime on the basis of daily number of passengers transported
- Develop an even more robust safety system that allows last-minute intervention in train access if anything goes wrong during the evacuation.
- Modify the design of the maintenance machines to bring the controls out of the danger area of the adjacent track.

The works have indeed been carried out while continuing the train service at a reduced speed. The regulators agreed to this, but were not entirely happy with this rather innovative trade-off.

Comment

The initial failure to consider risk trade-offs in this case came from the Labour Inspectorate limiting their concern to the track workers. The Railway Regulator and the infrastructure provider on the other hand saw the problems this created for railway capacity. The issue of passenger safety was mentioned as an aside by one of the directors of the infrastructure provider but it was not incorporated in their considerations in our research contract.

The problem was initially a too limited system boundary, whereby the decision of the Labour Inspectorate effectively exported the risk to those who are not its clients. It took the hiring in of a consultant to expand the system boundary and find a way of optimising the risk for the greater system. Apparently this was not something that could be resolved between the two regulators in interdepartmental discussion. We can raise the question whether an umbrella authority over the two regulators could have provided this forum for risk optimisation.

Once the system boundary had been widened, it was possible to find a common denominator to compare the various options, namely the total risk of death for all exposed persons. We did not have to use any notions of acceptable risk for the different groups and could value passengers' and track workers' lives the same.

3.2 Case study 2: crosswinds at Schiphol.

On Christmas Eve in 1997 at 22.47 a Transavia Boeing 757 landed in good visibility, but with strong crosswind conditions with gusting winds at Schiphol. The aircraft landed hard, with a force above the design limits of the nose wheel gear and doghouse, which broke. The aircraft skidded 3000m down the runway and then veered off onto the grass before stopping. There were only minor injuries to 4 passengers. The accident was investigated by the Dutch Transport Safety Council, whose report placed part of the blame on the runway selection procedure and recommended a review of the limits in use at Schiphol in the light of international standards and the state of the art. This selection procedure is strongly influenced by the noise contours of the airport, which try to limit the number of aircraft flying over residential areas in order to reduce the total noise dose they are subject to. The runway preferences are guided by a combination of factors. These include, apart from the noise considerations:

- the crosswind and tailwind velocities which influence the manoeuvring of the aircraft in the last crucial moments before landing (or after take-off) as well as on the runway;
- the friction coefficient of the runway (influenced by standing water, snow or ice), which affects the ability of the aircraft to stop within the runway limits under the prevailing wind conditions;
- the visibility conditions for landing, which also affect the pilots ability to line up for the runway and compensate for any cross and tailwinds and
- the presence of automatic landing aids on the runway, which influence whether the aircraft can be put down using its automatic systems. The limits of the safe performance of the automatic pilot under cross and tailwind conditions are also an issue of some importance.

As such it is a protocol which has to balance safety, in this case largely that of the passengers and crew of the landing aircraft, with the noise nuisance to the residents of the surrounding area. We ignore for the further discussion the safety risk to the local residents from crashing aircraft, since this is much smaller than that to the passengers and correlates reasonably with their noise nuisance.

The limits current at the time of the accident were:

Crosswind 15 knots, tailwind 5 knots (including gusts) for daytime in dry conditions (15/5[2]).

Crosswind 10 knots and no tailwind (including gusts) for wet conditions in daytime (10/0).

At night the criteria for good visibility conditions could be relaxed to 25/5, and for poor visibility to 15/5 in order to favour the use of runways that avoided using the approaches over residential areas.

The actual crosswind at the time of the accident was estimated to have been at least 35 knots with gusts to 43 knots.

Independent committee analysis

The parties concerned at Schiphol (the airport authorities, the ATC, the Dutch airlines), after discussion with the ministry concerned, decided to set up an independent expert committee to look at the balance being struck and to make recommendations about possible changes to the crosswind and tailwind criteria in the light of the accident and of developments in technology and international regulation. The report of the committee (CCTC 2000) contains detailed considerations of all of these issues and comes to the conclusion that there should be one criterion for both day and night-time crosswind and tailwind limits for landing at Schiphol. For daytime they recommended that the limit could safely be safely relaxed (from 15/5 to 20/7), under a number of provisos, whilst the situation for night-time is a tightening (from 25/5 to 20/7). The committee recommends that this limit can be relaxed further (at least to 25/7), if improvements are made to wind information systems, simulator training of pilots and some procedural issues relating to ATC – pilot communication. The research carried out for the committee (v.d. Geest et al 2000) provides evidence that this relaxation could be to a criterion of 25/10, whereby the noise contours can be maintained in almost all cases, whilst the airport can achieve an increased capacity for flights of an estimated 16% in 2010 (25/10 compared with 15/5).

Trade-offs

It is not the purpose of this paper to examine all the arguments of this case in detail, or to arrive at any conclusion about the correctness of the committee’s recommendations. What we want to do is to examine how the trade-off between the two risks, safety of passengers and crew vs. noise nuisance to residents, was handled.

[2] This notation is used to indicate the cross- and tailwind criteria in the rest of the paper.

What strikes the reader of the report in the first instance is that at no point is there a direct trade-off made in terms of so much more risk for so much less nuisance. The reasoning proceeds via intermediates, notably the already established norms and standards for each of the risks, and arguments about specific factors which mean that the risk should be considered to be less than it had been considered to be up to then. We will give some examples of this in a moment.

The trade-off used is in fact not just between noise nuisance and safety risk, but also between those and the airport capacity. When trading off between three variables one method is to hold one variable constant and to look at the trade-off between the other two. Normally one expects that the most important variable will be held constant (we must achieve at least this criterion on this dimension) and the room for manoeuvre will be taken up by comparing the results for the other two. It is striking that the committee bases its recommendations on the calculations in which the fulfilment of the noise contours as a minimum requirement is held constant, and the trade off is of risk against capacity. The research underlying the committee's recommendations estimates that relaxation of the cross- and tailwind criteria under these constant (legal) noise constraints will result in an approximate doubling of the Schiphol airport risk from crosswind accidents in moving from the 15/5 to a 25/10 criterion. This is an increase from one crosswind-related accident in an average of 32 years[3] to one per 17 years[4], a figure which also takes account of the 16% expected capacity increase allowed by the shift. The risk per individual landing under the conditions of 15/5 and 25/10 are given as 2.6e^{-7} and 11.4e^{-7}, i.e. somewhat more than a fourfold increase. The increase of capacity comes from being able to use the noise-preferred runway more often with higher permitted crosswind. Since Schiphol has a legally defined 'noise quota' and must limit its flight movements to stay within that, the capacity of the airport is ultimately noise limited. By using this legal quota as its departure point the research and the committee recommendations arrive at the trade-off allowing the 16% capacity increase if the 25/10 criterion is ultimately allowed. It settles for an interim increase in capacity of 6-10% by going in the first instance for a 20/7[5] criterion, all of which it justifies in the following words:

> ".. risk calculations show that in the scenario with the highest criteria (cross/tailwind of 25/10 in 2010), the expected mean time between accidents is more or less halved compared with the 15/5 condition. At the criterion of 20/7 the increase in mean time between accidents is *significantly improved*. (our italics) ….
>
> The majority of crosswind related incidents or accidents (97%) incurred no fatal on-board casualties. Given the mean time between accidents of 25 years (scenario 20/7 in 2010), fatal accidents would occur less than once per eight centuries[6]."

In further recommendations the committee argues that a further relaxation of the criterion to 25/10 is not justified and suggestions one of 25/7 as a further compromise. For this variant the research made no calculations, which leaves the nature of the trade-off vague. Although the committee does not say so in so many words, its recommendations mean that it regards this increase in risk as justified to achieve the increase in capacity at the constant noise nuisance levels set down in the law.

Its arguments mean that a 22% increase in risk (from once in 32 years to once in 25 years) is worth 6-10% increase in traffic (27,000 – 47,000 extra flight movements). This is expressing risk in relative terms, which could be regarded as exaggerating it, since the absolute risk level is 'low'. Indeed the arguments in the report are designed to indicate that the absolute increase in risk is minimal. The additional arguments about the measures which can be taken to reduce this risk further (better information and training to pilots, better measurement of wind, better autopilot technology) seem intended to imply that this absolute small increase can be removed by the additional measures. The fact that the calculated risk is a doubling for the passengers and crew actually flying in the plane that lands at 20/7 instead of 15/5 is further not recalled in the summary and recommendations. Nor is the question ever raised whether this risk of death is indeed low. If we translate the probability of an accident with a fatality (once in eight centuries according to the studies – see above) into a probability of death per year, and assume (conservatively) that the fatal accident will only result in 4 fatalities, we arrive at a risk of 6e^{-3}/year. Is this small? What are reasonable comparisons? The maximum acceptable individual risk for a local resident at Schiphol is more than two orders of magnitude lower (5e^{-5}, with a proposal to reduce it to 1e^{-5}), but this is the risk of someone standing unprotected on the contour for 24 hours a day for the whole year, whilst the crosswind

[3] 95% confidence limits 17-240 years

[4] 95% confidence limits 9-104 years

[5] Risk per individual landing 5.4 x 10^{-7}, risk of crosswind-related accident at Schiphol every 25 years

[6] NB these are accidents due to crosswind. The mean time between accidents from all causes at Schiphol is approximately 10 years.

fatality risk is for the whole 'installation' and is spread over all passengers using the airport. So, what else can we use?

There is also a table in the report that shows the decrease in number of houses falling inside the 35 cost unit[7] contour with each relaxation of the cross/tailwind criteria. This can be combined with the table indicating the increase in risk for each of these steps, though the report never does this and the text does not suggest it. If we make this combination and translate the mean time between accidents, as we have done above, into probability of death per year, given four fatalities per fatal accident, we produce the following table:

Shift in wind criteria	Absolute decrease in houses within 35cu contour	Percentage decrease in houses exposed	Change in p. of death/year (e^{-4})	Percentage increase in risk
15/5 - 20/7 in 2003	9011-8223=788	9	11.1 - 9.1 = 2.0	22
15/5 - 25/5 in 2003	9011-7435=1576	17.5	14.3 – 9.1 = 5.2	57
15/5 - 20/7 in 2010	12016-11254=762	6	12.0 - 9.3 = 2.7	29
15/5 - 25/5 in 2010	12016-10491=1525	13	16.7 – 9.3 = 7.4	80
15/5 - 25/10 in 2010	12016-10058=1958	16	17.6 – 9.3 = 8.3	89

The percentage increases in death risk in each case are 2.5 to 5.5 times the decrease in households suffering this definition of noise nuisance. If we calculate what the additional risk of death is per household removed from the 35cu contour we arrive at figures ranging from 2.5e^{-7} to 4.8e^{-7}, or expressed the other way round, between 2 and 4 million households removed from that contour per death.

Comment

We note here that it is possible, despite all of the caveats that must always accompany risk measurements, to compare quantitatively the increases in the safety risk to passengers and crew, which are entailed in reducing (or keeping constant) the noise nuisance to local residents. What is striking is that the report assiduously refrains from making that comparison directly. Is this because it wishes to avoid raising the sort of issues that this workshop will address?

The methods of reasoning which the report uses to avoid the direct confrontation of the risk trade-off between different parties is to use one legal exposure limit as an absolute reference and to make the comparison by the proxy of the capacity increases. Risk is traded with airport capacity in the report and not with noise nuisance directly.

3.3 Discussion

We can raise a number of issues for discussion. We do so here rather briefly, since we hope the workshop will expand on many of them further.

1. The risk comparison issue is always one of drawing system boundaries. Which risks will we consider and to which people? Often the tendency is to draw the boundaries of comparison along the lines of the responsibilities of certain parties – work related, linked to a particular installation, etc. The track worker case shows how this can lead to problems with too small a definition. However, risks do not respect system boundaries. No matter where we draw the boundaries there are always internal and external risks and internal and external people. So, what is then a sensible boundary? We can provide another example related to Schiphol, which expands on this issue. We have mentioned in the crosswind case study that there is a proposed new individual risk contour of 1e^{-5} around Schiphol. Within this all houses must be demolished and between this and 1e^{-6} no new houses may be built. Different rules apply to industrial development (less strict) and to 'sensitive' constructions such as schools and hospitals (more stringent). No restrictions at all are put on road developments, despite the fact that the motorways and access roads around the airport are regularly full of stationary traffic jams, often coinciding with the peaks of arriving and departing flights, making their occupancy density quite as high as permanent buildings. The people working within the Schiphol perimeter do not count for external safety calculations and so are not

[7] Cost units (cu) are units for expressing the noise nuisance, based on a night (23.00-06.00) weighted exposure calculation to aircraft noise over the 24 hours. It bounds the area which is called 'relatively high noise exposure' in the legal limits for Schiphol and is the contour for which the number of houses exposed must be kept under 10,000. As comparison, houses exposed to more than 40cu qualify for general noise insulation measures and bedrooms exposed to more than 26cu qualify for noise insulation to that level.

protected by any restrictions. This implies different values for the life (and for protection against noise and other nuisances) for different people. We have known for many years the importance of the distinction between voluntary and involuntary risk when it comes to risk perception. These values appear to be incorporated in some way in the distinctions made here, but the exact ways they are related are not at all clear. The state of the current risk comparisons and risk criteria at Schiphol seem to be more a matter of which Ministries have been active in protecting the interests of their 'clients'. Hence it is the clients of the Ministry of Housing, Planning and Environment who are counted for the risk contours and not the clients of the Social Affairs Ministry (workers) or the Transport Ministry (drivers & passengers)

2. The issue is most important when actions to control one risk (or to control risk to one group of people) increase that for another. If we can find an action which will improve the risk of all concerned, or reduce all types of risk, it is not hard to decide whether to take it. Otherwise we need to ask whether we value the different groups equally – and if not, why not? The question for comparing types of risks is how we can find a dimension on which we can ask the question whether we value the risks equally? Is probability of death adequate, or are there preferred ways of dying? How do we value suffering and annoyance? In the track workers case we could compare risk of death directly and did not feel the need to value passenger deaths higher than track worker deaths. If we had adopted the notion that those who have control over risk should be protected less than those who have no control, we might have had a hard time giving weightings. The train passengers clearly have no control over the risk of derailment, but do have a major influence on their risk if they take their cars to cover the same trip. However, if they carpool and are passengers, they are again relatively powerless. The track workers have some control over risks, except for those involving the incorrect routing of a train onto their track. They also gain their living from carrying out this work – another factor which we might take account of in valuing risks to them. The passengers (even if they become car drivers) are probably only using the trip as an incidental necessity to get to their work and so get less out of it to set against the risk. Are these useful issues to take account of? In the crosswind case we raise the issue of comparability of risk by making the false comparison between external safety criteria and the risk of additional deaths from less stringent crosswind criteria. These two risks can both be reduced to probability of death per year, but each has a totally different reference population. We need to be very careful that our units are not only the same for all options we consider, but are appropriate to the comparison in question.

3. In the crosswind case we have to compare across types of risk. We want to trade risk for noise nuisance. In the case study we calculate what the trade-off is which the proposed decisions imply. But the next question is whether this is a reasonable trade-off? The measure of noise nuisance is not even a 'total volume of noise nuisance', which would be in some way comparable to the 'total volume of death' represented by the risk measure. It is simply a count of houses subject to a previously defined 'relatively high exposure'. It will surely matter in considering how justified the trade-off is, whether we see the three million households being saved from just mild annoyance, or from a miserable existence, by the sacrifice of the one life per year which is being proposed.

4. The crosswind case also illustrates the use of legally defined limits as anchor points from which to argue about trade-offs. The (arbitrarily) defined limit of not more than 10.000 houses in the 35cu contour has been set by law, as will the tightened individual risk contour of 1e^{-5} be. A simple way of making risk comparisons is then in terms of how far we are from the legally defined limit. Anything above that limit is regarded as unacceptable and therefore not to be traded-off. The fact that there was no legally defined risk limit for passengers meant that the crosswind case could trade risk and capacity within the bounds of the noise limit. If all risks have legally defined criteria we can give priority to actions in terms of how much they reduce the amount by which a given risk exceeds its legal limit. Once we get risks below their legal limits, we often meet another way out of phrasing the same priority question. We often shift from requiring a certain result (compliance) in terms of protection for a given group, to requiring a certain effort to protect. This is the ALARA principle. Here we give priority to the action that uses the least resources to produce the most reduction below the legal limit.

5. We seem at times to use backwards reasoning to justify risk priorities or differences in criteria. We decide (in a variety of ways) that cannot afford to spend more than a certain amount, or to cause more than a certain disruption. That is the maximum achievable in the current state of society. Reasoning backwards, we then arrive at the risk criterion or the value we can afford to put on people's lives. Schiphol can again be used to illustrate this point. When it was first decided to impose an external individual risk requirement, the only one existing was that for chemical plant under the Dutch regulations implementing the Seveso Directive. This was 1e^{-5} for existing plant. If this had been imposed at Schiphol, based on the calculations of the risk made for 1990, too many houses would have had to be demolished. So a criterion

was set of 5e^{-5}. When recalculation of the risks in 2000 led to the insight that risks were lower than expected and hence the individual risk contours were closer to Schiphol than expected, the policy makers suddenly agreed to tighten the criterion to the same as chemical works. We should not interpret this change as a realisation that residents have increased in value five times since 1990! This backwards reasoning seems to us to be the main reason why proposals for standards for risk exposure of workers lie one or two orders of magnitude higher risk than for residents. It is not that workers are intrinsically less valuable than residents; simply that this is approximately what the risk difference now is.

6. All of this illustrates the difficulty of making risk coherent. This has led in the past to the cry that it becomes much too complicated if we try to optimise risk across too complex a system, so we should leave it as a piecemeal tapestry of clashing colours. Perhaps this workshop will indicate if there indeed are ways to order these colours and patterns into useful and coherent comparisons of risk.