



PII: S0925-7535(97)00052-0

RISK MANAGEMENT IN A DYNAMIC SOCIETY: A MODELLING PROBLEM

Jens Rasmussen

Hurecon, Smorum Bygarde 52, DK 2765 Smorum, Denmark

Abstract—In spite of all efforts to design safer systems, we still witness severe, large-scale accidents. A basic question is: Do we actually have adequate models of accident causation in the present dynamic society? The socio-technical system involved in risk management includes several levels ranging from legislators, over managers and work planners, to system operators. This system is presently stressed by a fast pace of technological change, by an increasingly aggressive, competitive environment, and by changing regulatory practices and public pressure.

Traditionally, each level of this is studied separately by a particular academic discipline, and modelling is done by generalising across systems and their particular hazard sources. It is argued that risk management must be modelled by cross-disciplinary studies, considering risk management to be a control problem and serving to represent the control structure involving all levels of society for each particular hazard category.

Furthermore, it is argued that this requires a system-oriented approach based on functional abstraction rather than structural decomposition. Therefore, task analysis focused on action sequences and occasional deviation in terms of human errors should be replaced by a model of behaviour shaping mechanisms in terms of work system constraints, boundaries of acceptable performance, and subjective criteria guiding adaptation to change. It is found that at present a convergence of research paradigms of human sciences guided by cognitive science concepts supports this approach. A review of this convergence within decision theory and management research is presented in comparison with the evolution of paradigms within safety research.
© 1997 Elsevier Science Ltd. All rights reserved.

1. Introduction

The following discussion of modelling risk management in a dynamic society is based on the experiences gained during several decades from multi-disciplinary research on industrial risk management at Risø National Laboratory and from the interaction within a multi-disciplinary international network that evolved from the Bad Homburg workshop series “New Technology and Work” (Wilpert, 1987-). Looking back on the evolution of our modelling efforts, it appears as if our modelling concepts have moved around a full circle through several research disciplines and paradigms. Apparently, when meeting problems we frequently found the grass to be greener in the garden of yet another discipline.

Our research started as an effort within systems and control engineering to design control and safety systems for hazardous industrial process plants. An evaluation of our results by

comparison with accident records rapidly guided our attention to the human-machine interface problems and we were forced to enter the arena of human error analysis, operator modelling, and display design, involving also psychological competence. From here, we quite naturally drifted into studies of the performance of the people preparing the work conditions of operators and we found it necessary to have a look at the results of research within management and organisational science also to consider the problem of decision errors at the management level. Finally, since a major problem appeared to be management's commitment to safety and the related efforts of society to control management incentives by safety regulation, we also had to involve experts in law and legislation in our studies.

Now, after several decades of modelling efforts we find that the models we create by bringing together results from several disciplines can be very useful for design of work support systems for the individual actors and decision makers, but they are not very useful for analysing the performance of the total risk management system. For this problem, a system model cannot be built by a bottom-up aggregation of models derived from research in the individual disciplines, but a top-down, system oriented approach based on control theoretic concepts is required.

The reason for this is that a system is more than the sum of its elements. Often we found that attempts to improve the safety of a system from models of local features were compensated by people adapting to the change in an unpredicted way.

Hence, we are back to considering our initial conceptual control framework in a more complex setting, that is, as being a control structure embedded in an adaptive socio-technical system.

2. The problem space: risk management in a dynamic society

Injuries, contamination of environment, and loss of investment all depend on loss of control of physical processes capable of injuring people or damaging property. The propagation of an accidental course of events is shaped by the activity of people which either can trigger an accidental flow of events or divert a normal flow. Safety, then, depends on the control of work processes so as to avoid accidental side effects causing harm to people, environment, or investment.

Many levels of politicians, managers, safety officers, and work planners are involved in the control of safety by means of laws, rules, and instructions that are formalised means for the ultimate control of some hazardous, physical process. They seek to motivate workers and operators, to educate them, to guide them, or to constrain their behaviour by rules and equipment design, so as to increase the safety of their performance.

The socio-technical system actually involved in the control of safety is shown in Fig. 1, that represents the problem space we passed through, bottom-up during our research, with several academic disciplines involved at each of the various levels. At the top, society seeks to control safety through the legal system: safety has a high priority, but so has employment and trade balance. Legislation makes explicit the priorities of conflicting goals and sets boundaries of acceptable human conditions. Research at this level is within the focus of political and legal sciences. Next we are at the level of authorities and industrial associations, workers' unions and other interest organisations. Here, the legislation is interpreted and implemented in rules to control activities in certain kinds of work places, for certain kinds of employees. This is the level of management scientists and work sociologists. To be operational, the rules now have to

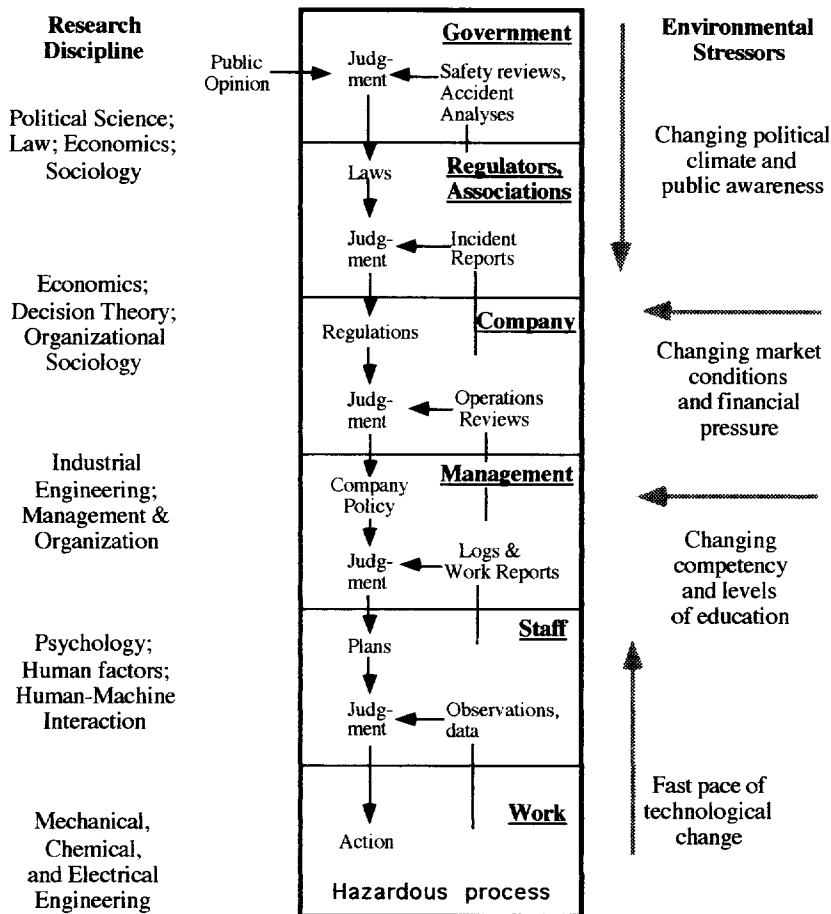


Fig. 1. The socio-technical system involved in risk management.

be interpreted and implemented in the context of a particular company, considering the work processes and equipment applied. Again, many details drawn from the local conditions and processes have to be added to make the rules operational and, again, new disciplines are involved such as work psychologists and researchers in human-machine interaction. Finally, at the bottom level we meet the engineering disciplines involved in the design of the productive and potentially hazardous processes and equipment and in developing standard operating procedures for the relevant operational states, including disturbances.

Control of activities and their safety by the classic prescriptive command-and-control approach deriving rules of conduct top-down may be effective in a stable society where instruction and work tools at all levels can be based on task analysis. In the present dynamic situation, this approach is inadequate and a fundamentally different view on system modelling is required.

3. The present dynamic society

Compared to the stable conditions of the past, the present dynamic society brings with it some dramatic changes in the conditions of industrial risk management:

- A very fast pace of change of technology is found at the operative level of society within many domains, such as transport, shipping, manufacturing and process industry. This pace of change is much faster than the pace of change presently in management structures — Savage and Appleton (1988) talk of “second generation management applied to fifth generation technology” in manufacturing. An even longer lag in response to change is found in legislation and regulation. The different time lags found at the different levels thus present a problem, and the dynamic interaction among levels during a period of change becomes an important modelling issue.

- The scale of industrial installations is steadily increasing with a corresponding potential for large-scale accidents. Very low probabilities of accidents have to be demonstrated for acceptance of operation by society. Consequently, models should not only include normal or average performance, but also very rare conditions.

- The rapid development of information and communication technology leads to a high degree of integration and coupling of systems and the effects of a single decision can have dramatic effects that propagate rapidly and widely through the global society. This has been demonstrated by the effects of less successful computerised trading systems (e.g., the Wall Street turbulence in 1987 (Waldrop, 1987)). It is thus becoming increasingly difficult to model systems in isolation and to make small-scale, local experiments to evaluate models.

- Furthermore, companies today live in a very aggressive and competitive environment which will focus the incentives of decision makers on short term financial and survival criteria rather than long term criteria concerning welfare, safety, and environmental impact.

These trends have a dramatic effect on the necessary approach to modelling system behaviour in some very fundamental respects, and they raise the problems of modelling by structural decomposition versus functional abstraction and the problem of cross-disciplinary research versus multi-disciplinary co-operation (Hale et al., 1996; Rasmussen et al., 1994).

4. Modelling by structural decomposition: tasks, acts and errors

The usual approach to modelling socio-technical systems is by decomposition into elements that are modelled separately. This has some peculiar effects. The socio-technical system involved in risk management is, as shown in Fig. 1, normally decomposed according to organisational levels which then are the objects of study within different disciplines.

The effect of this is that risk management at the upper levels is normally studied with a ‘horizontal’ orientation of research across the technological hazard sources. Traditionally, sociological studies at the upper levels are based on analysis of samples of organisations or groups of people with no detailed consideration of the actual processes found at the productive bottom level. Analyses are based on statistics and industry-wide questionnaires. Barley (1988) made a similar observation when studying a particular work domain — radiological work in medicine.

In this way, management theories tend to be independent of the substance matter context of the given organisations. To be a manager is seen from this point of view as to be a profession, independent of whether you are managing a hospital, a manufacturing company, or a bank.

Also the aim of commercial companies presently appears to have changed from being organisations serving a particular substance matter domain toward a narrow focus on financial operations (see e.g., Engwall, 1986). What are the implications of this situation for the societal control of the safety of industrial installations? Following a recent Scandinavian ferry accident, a marine safety official noted on a TV interview that we might see a decrease in naval safety, since ships were increasingly operated by banks and investors rather than shipping professionals. A recent critical review of the effects of this trend on management behaviour in, e.g., public health care, is found in Rees and Rodley (1995).

In this situation we need more studies of the vertical interaction among the levels of socio-technical systems with reference to the nature of the technological hazard they are assumed to control.

While a *system* traditionally is modelled by decomposition into structural elements, the dynamic *behaviour of systems* and *actors* is modelled by decomposition of the behavioural flow into events. Such decomposition is the basis for identification of activity elements in terms of tasks and in task elements in terms of decisions, acts, and errors. The problem is that all work situations leave many degrees of freedom to the actors for choice of means and time for action even when the objectives of work are fulfilled and a task instruction or standard operating procedure in terms of a sequence of acts cannot be used as a reference of judging behaviour. To complete a description of a task as being a sequence of acts, these degrees of freedom must be resolved by assuming additional performance criteria that appear to be 'rational' to instructors. They cannot, however, foresee all local contingencies of the work context and, in particular, a rule or instruction is often designed separately for a particular task in isolation whereas, in the actual situation, several tasks are active in a time sharing mode. This poses additional constraints on the procedure to use, which were not known by the instructor. In consequence, rules, laws, and instructions practically speaking are never followed to the letter. Strikes by civil servants take the shape of 'working-according-to-rules'. Even for highly constrained task situations such as nuclear power operation, modification of instructions is repeatedly found (Fujita, 1991; Vicente et al., 1995) and operators' violations of formal rules appear to be quite rational, given the actual work load and timing constraints. One implication in the present context is that following an accident it will be easy to find someone involved in the dynamic flow of events that has violated a formal rule by following established practice, and who is, therefore, likely to be exposed to punishment. Consequently, accidents are typically judged to be caused by 'human error' on the part of a train driver, a pilot, or a process operator (Rasmussen, 1990a,b, 1993c).

A task instruction thus is an unreliable standard for judging behaviour in actual work. Likewise, modelling human behaviour in terms of a stream of acts will be unreliable for a dynamic environment when behaviour is context dependent. Modelling by 'task analysis' is only useful when behaviour is very tightly controlled by the control requirements of a technical system. In general we have to seek models at a higher conceptual level.

Another example of decomposition of behaviour is the modelling of behavioural control by separate decisions. In traditional decision research 'decisions' have been perceived as discrete processes that can be separated from the context and studied as an isolated phenomenon. In our field studies (Rasmussen, 1992, 1993a), we have found it very difficult to isolate proper decisions. In a familiar work environment actors are immersed in the work context for extended periods; they know by heart the normal flow of activities and the action alternatives available. During familiar situations, therefore, knowledge-based, analytical reasoning and planning is replaced by a simple skill- and rule-based choice among familiar action alterna-

tives, that is, on practice and know-how. When, in such situations, operational decisions are taken, they will not be based on rational situation analysis, only on the information which, in the running context, is necessary to distinguish among the perceived alternatives for action. Separate 'decisions' therefore are difficult to identify and the study of decision making cannot be separated from a simultaneous study of the social context and value system in which it takes place and the dynamic work process it is intended to control. This problem has led to development of the skill-, rule-, knowledge-based behaviour model of cognitive control (Rasmussen, 1983) and the more recent paradigms of 'naturalistic' decision making (For a review, see Klein et al., 1994). In general, the present interest in cognitive science has brought with it a convergence of the economist's concept of 'decision making', the social concept of 'management', and a psychological concept of 'cognitive control' of human activity, a topic we will consider in more detail below.

5. Accident causation

Considering the problem of the frequent deviation from normative work instructions and rules, it is no wonder that it is often concluded in accident reviews that 'human error' is a determining factor in 70–80% of the cases. Furthermore, multiple contributing errors and faults are normally found, because several defenses against accident are usually planned for hazardous processes.

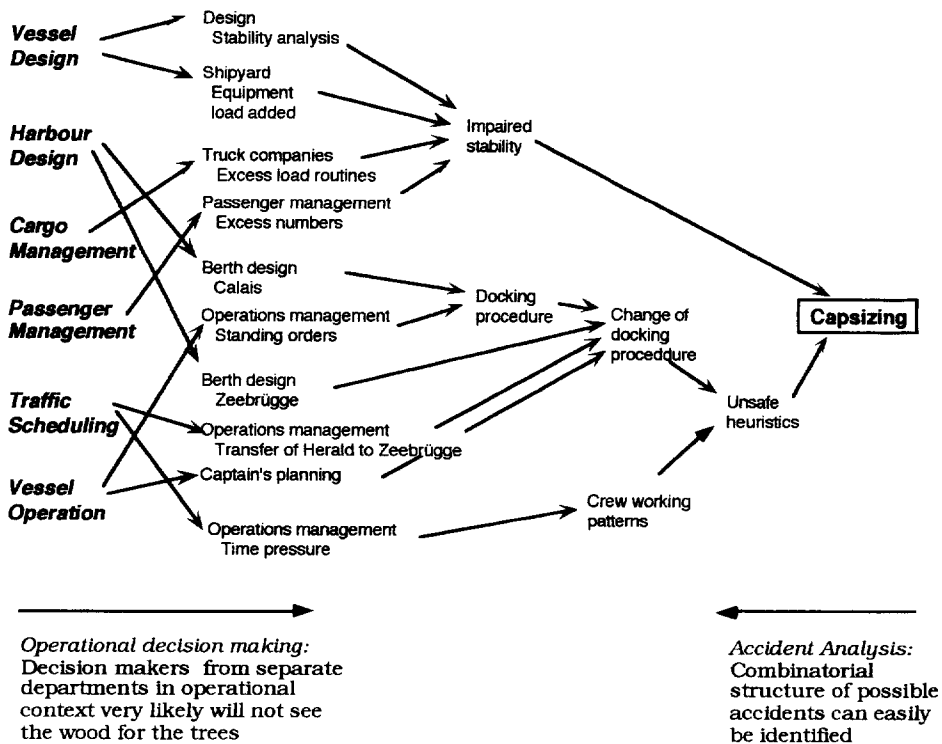


Fig. 2. The complex pattern of the Zeebrugge accident.

However, it should not be forgotten that commercial success in a competitive environment implies exploitation of the benefit from operating at the fringes of the usual, accepted practice. Closing in on and exploring the boundaries of the normal and functionally acceptable boundaries of established practice during critical situations necessarily implies the risk of crossing the limits of safe practices. Correspondingly, court reports from several accidents such as Bhopal, Flixborough, Zeebrugge, and Chernobyl demonstrate that they have not been caused by a coincidence of independent failures and human errors, but by a systematic migration of organisational behaviour toward accident under the influence of pressure toward cost-effectiveness in an aggressive, competitive environment (Rasmussen, 1993b, 1994b).

In this situation, we have to consider the interaction of the effects of decisions made by several actors in their normal work context, all very likely to be subject to the same kind of competitive stress. Fig. 2 shows a causal tree derived from the Zeebrugge accident where several decision makers at different times, in different parts of a shipping company, all are striving locally to optimise cost effectiveness and thus are preparing the stage for an accident. The dynamic flow of events can then be released by a single human act. As seen by the individual decision makers from the front end, it will be difficult to see the total picture during their daily operational decision making. The individual decision makers cannot see the complete picture and judge the state of the multiple defenses conditionally depending on decisions taken by other people in other departments and organisations.

In this situation, modelling activity in terms of task sequences and errors is not very effective for understanding behaviour, we have to dig deeper to understand the basic behaviour shaping mechanisms.

6. Modelling by functional abstraction: migration toward the boundary

This discussion points to the existence of a natural migration of activities toward the boundary of acceptable performance and we have to find a way to represent the mechanism underlying this migration.

Human behaviour in any work system is shaped by objectives and constraints which must be respected by the actors for work performance to be successful. Aiming at such productive targets, however, many degrees of freedom are left open which will have to be closed by the individual actor by an adaptive search guided by process criteria such as work load, cost effectiveness, risk of failure, joy of exploration, etc. The work space within which the human actors can navigate freely during this search is bounded by administrative, functional, and safety related constraints. The normal changes found in local work conditions lead to frequent modifications of strategies and activity will show great variability. Such local, situation-induced variations within the work space call to mind the 'Brownian movements' of the molecules of a gas. During the adaptive search the actors have ample opportunity to identify 'an effort gradient' and management will normally supply an effective 'cost gradient'. The result will very likely be a systematic migration toward the boundary of functionally acceptable performance and, if crossing the boundary is irreversible, an error or an accident may occur.

In any well designed work system, numerous precautions are taken to protect the actors against occupational risk and the system against major accidents, using a 'defence-in-depth' design strategy. One basic problem is that in such a system having functionally redundant protective defenses, a local violation of one of the defenses has no immediate, visible effect

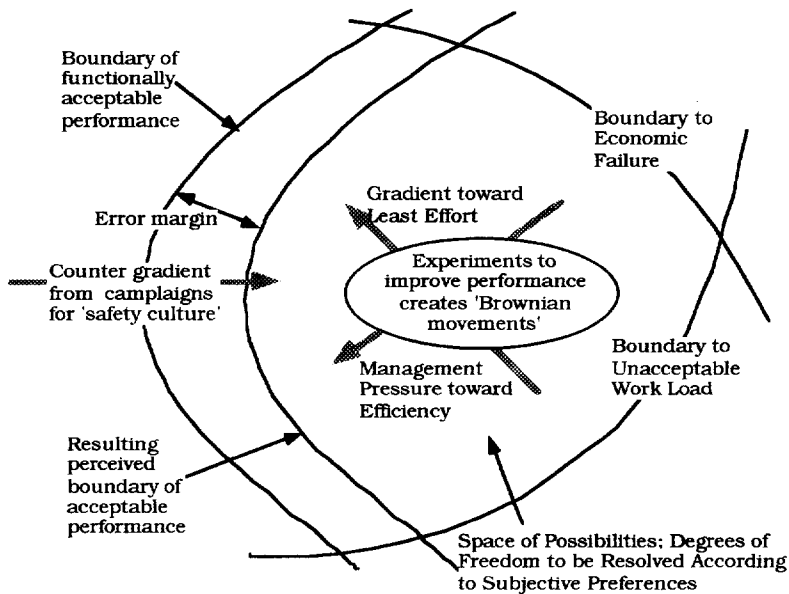


Fig. 3. Under the presence of strong gradients behaviour will very likely migrate toward the boundary of acceptable performance.

and then may not be observed in action. In this situation, the boundary of safe behaviour of one particular actor depends on the possible violation of defenses by other actors, see Fig. 3. Therefore, in systems designed according to the defence-in-depth strategy, the defenses are likely to degenerate systematically through time, when pressure toward cost-effectiveness is dominating. Correspondingly, it is often concluded by accident investigations that the particular accident was actually waiting for its release (Rasmussen, 1993b).

The important issue is that the stage for an accidental course of events very likely is prepared through time by the normal efforts of many actors in their respective daily work context, responding to the standing request to be cost-effective. Ultimately, a quite normal variation in somebody's behaviour can then release an accident. Had this particular 'root cause' been avoided by some additional safety measure, the accident would very likely be released by another cause at another point in time. In other words, an explanation of the accident in terms of events, acts, and errors is not very useful for design of improved systems.

When the decision makers managing institutions and companies involved in risk management adapt individually to their normal commercial stresses, the resulting interaction will very likely not match the overall safety control requirements, when judged after the fact. Fig. 4 shows the interaction conflicts between institutions and companies at the various levels as identified from super tanker and ro-ro ferry accidents (see Shell, 1992; Estonia, 1995; Stenstrom, 1995). Similar conflicts are found in the aviation domain (see Schiavo, 1997).

It is evident that a new approach to representation of system behaviour is necessary, not focused on human errors and violations, but on the mechanisms *generating* behaviour in the actual, dynamic work context. The analogy was mentioned with thermo-dynamic models in terms of the boundary conditions and gradients of a field which take us further than the representation of classical physics in terms of particles, paths of travelling, and the events of

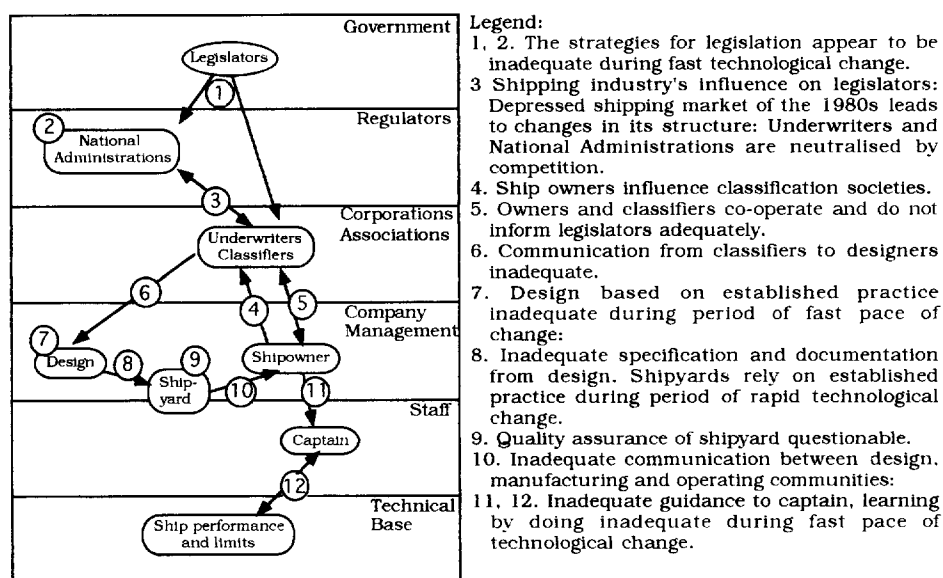


Fig. 4. Map conflicts among actors in shipping.

collision. Correspondingly, for human behaviour, we need a representation at the higher level of functional abstraction than the level used for task analysis. Such a representation involves identification of the boundary conditions of the work space and the gradients in terms of process criteria guiding the drift across this space. This approach involves a detailed study of the means-ends relations of the work system and, in turn, to the need of studies focused on particular types of systems characterised by their peculiar work processes and hazard sources. For this aim we need a taxonomy of hazard sources and their control characteristics.

In addition, we need a framework for identification of the objectives, value structures, and subjective preferences governing the behaviour within the degrees of freedom faced by the individual decision maker and actor (for details, see Rasmussen, 1994c and Rasmussen et al., 1994). This modelling approach is related to the Gibsonian concepts of invariants, affordances (Gibson, 1966, 1979), and the related representation of the 'space of safe driving' (Gibson and Crooks, 1938). For a review of the state of the art within this ecological approach, see Flach et al. (1994).

7. Control of system performance

This change in approach to modelling accident causation also invites a new approach to the control of system performance. Rather than striving to control behaviour by *fighting deviations* from a particular pre-planned path, the focus should be on the control of behaviour by *making the boundaries explicit and known* and by giving opportunities to develop *coping skills at boundaries*.

Efforts to improve the safety of skilled activities take a number of different shapes with respect to this control:

One way is to *increase the margin* from normal operation to the loss-of-control boundary. In this case, the natural adaptation to the boundary will very likely compensate (see the discussion of the effect of radar; anti-blocking car brakes, and 'risk homeostasis' in a section below). The resulting level of safety consequently depends on the recovery characteristics of the system. If the width of the recovery zone in a re-designed system has been decreased or the transition to loss of control has become more abrupt, the actual level of safety may deteriorate as a consequence of an 'improvement'.

Another way is to *increase the awareness* of the boundary by means of instruction and motivation campaigns. Safety campaigns may create a 'counter gradient' to the cost-effectiveness gradient and thereby serve to maintain the margin, see Fig. 3. However, an empirical struggle for a good 'safety culture' will never end because it only works as long as it acts as a *continuous pressure* compensating the *functional pressure* of the work environment. With the present high level of industrial safety attained by this empirical approach, significant further improvement will probably require an approach directed selectively toward the behaviour shaping factors of particular process environments.

The most promising general approach to improved risk management appears to be an explicit identification of the boundaries of safe operation together with efforts to make these boundaries visible to the actors and to give them an opportunity to learn to cope with the boundaries. In addition to improved safety, making boundaries visible may also increase system effectiveness in that operation close to known boundaries may be safer than requiring excessive margins which are likely to deteriorate in unpredictable ways under pressure.

8. Risk management: a control task

It follows from this discussion, that risk management is to be considered a control function focused on maintaining a particular hazardous, productive process within the boundaries of safe operation and that a systems approach based on control theoretic concepts should be applied to describe the overall system functions. As a basis for the discussion of the implications of this point of view on the design of risk management systems, a brief review of the relationship between control strategies and design approaches will be useful.

Management and work planning in any organisation apply different control strategies, depending on time horizon, stability of systems, and predictability of disturbances. One classic strategy that has been effective during a period with production for a stable market is the centralised planning based on prognosis from last years result with management by formal materials and resource planning (MRP) systems. In this mode, management is based on monitoring of performance with reference to plans, budgets and schedules, and control is aimed at removing deviations, 'errors'. The planning model is then revised from observation of the accumulated result over the planning period. This strategy can only cope with disturbances if they have been predicted and included in the planning model.

This strategy has presently been widely replaced by the more dynamic customer-controlled, just-in-time production strategy. It has not always been realised, that the two strategies are good examples of an open-loop and a closed loop strategy and that the two strategies call for very different design approaches (Rasmussen, 1994b).

Modelling risk management in a dynamic society in which all actors continuously strive to adapt to changes and the pressure of dynamic markets, we clearly have to apply such an active, closed loop feedback point of view. An abstract representation must be used of the

information network of the entire system involved in the control of the hazard presented by the technical core at the bottom. This control function and its stability cannot be studied at a general level across systems, but must be analysed for a particular system, or type of system. Therefore, studies of risk management must be based on a categorisation of hazard sources according to their control requirements.

For a particular hazard source, the control structure must be identified, all relevant controllers (actors) identified, their objectives and performance criteria determined, their capability of control evaluated, and the information available to them about the actual state of the system with respect to production objectives and safety boundaries must be analysed from a feed-back control point of view, as reviewed in the following paragraphs.

8.1. Identification of controllers

The first step is the identification of the decision makers (*controllers*) who may be involved in the preparation of the landscape through which an accident may propagate. One approach to this identification has been developed from accident analysis (Svedung and Rasmussen, 1996) by mapping the relationships among the decision makers who contributed to accident causation by their location in the system of Fig. 1, see the 'AcciMap' in Figs. 5 and 6.

From a control point of view, the interaction among the decision makers potentially involved in accident causation has some very special features. All these decision makers are busy managing their particular work domains and their attention will be focused on the control of the means and ends of their normal productive tasks while they strive to meet their

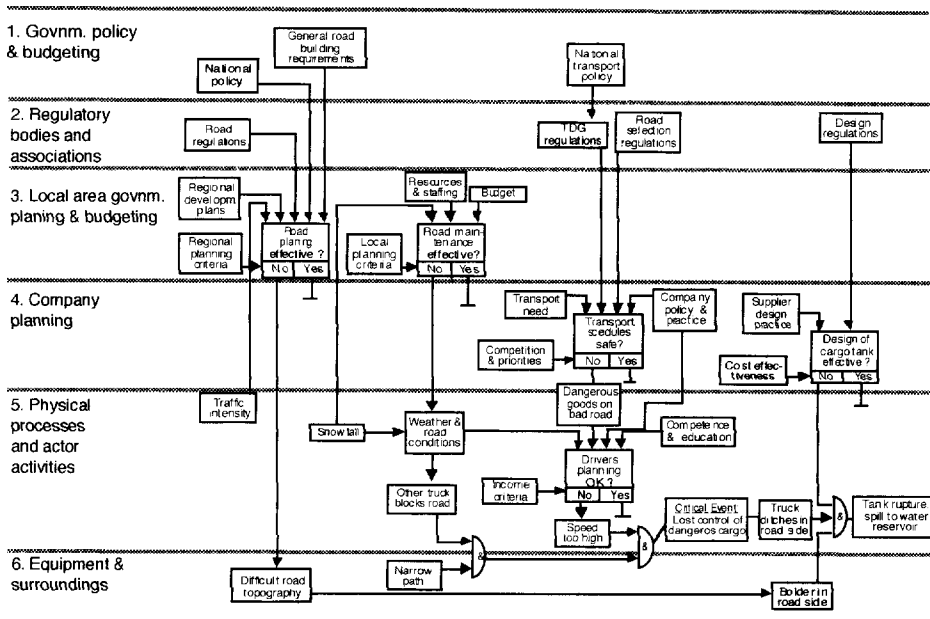


Fig. 5. A map showing the results of the analysis of a traffic accident involving oil-spill to a drinking water supply.

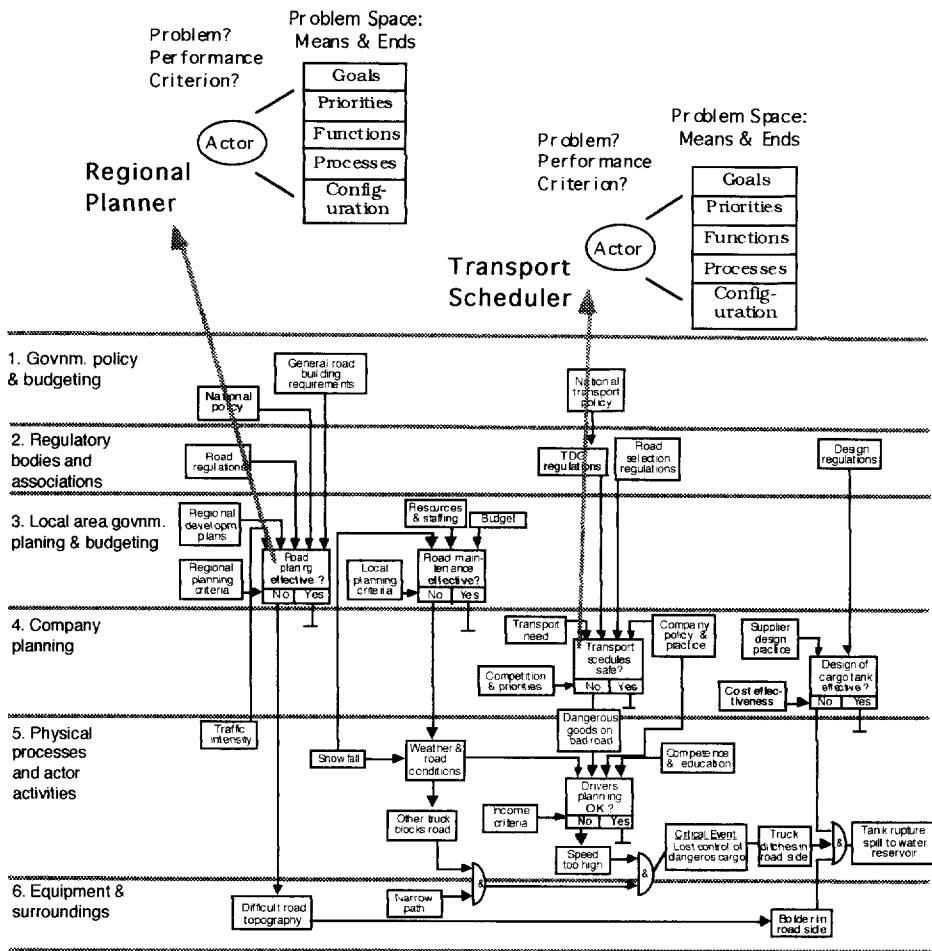


Fig. 6. The course of an accidental event is created by side effects of decisions made by decision makers busy coping with their local work requirements.

production targets, often under considerable stress to optimise process criteria such as time spent and cost-effectiveness. This must be done while respecting the constraint defined for their local context, including the boundaries defining safe overall operation. A critical issue is that the boundaries relevant to a particular decision maker depend on the activities of several other decision makers found within the total system and that *accidents are created by the interaction of potential side effects of the performance of several decision makers during their normal work.*

8.2. Work objectives

It is critical for any control function that the controllers have information about proper *action targets* (productive and safety related objectives) in a form and at a level corresponding to their action opportunities.

According to the previous discussion, this is a very complex issue. For the normal work

activities, basic business objectives propagate downward through an organisation and will be formulated by different concepts at the different levels of the company involved. Objectives for action have many shades. They can be expressed in terms of product specifications and targets for production volume; process optimisation criteria such as cost effectiveness; or constraints on processes, such as safety, work conditions, etc. A critical issue one meets when analysing performance criteria in actual work is that they very often are implicit in company or local work practice and difficult to make explicit.

At the higher levels, objectives and values are formulated in rather general terms. When objectives propagate downward through an organisation, degrees of freedom for action multiply since there are several ways of implementing objectives. Accordingly, objectives are to be interpreted according to the particular local context.

So far the normal work objectives. For the overall risk management, similar considerations apply for the propagation of objectives downward through the socio-technical system of Fig. 1. For this propagation, special care should be taken to analyse the influence of the different time-lags in the response to change at the various levels.

The present general trend in legislation and regulation away from prescriptive rules to performance-centred objectives should be considered carefully when analysing the propagation of objectives and values downward through the socio-technical system in Fig. 1. A trend is found away from prescriptive and toward performance-based legislation for several types of risks in the US. Since the regulatory process typically requires 6–10 years to develop adequate prescriptions, the fast technological pace of change has led to the introduction of the ‘general duty clause’ that has substantially enhanced the regulator’s ability to protect workers during the latest decades (Baram, 1996). This clause states that each employer “shall furnish to each of his employees a place of employment which is free from recognised hazards that are causing or are likely to cause death or serious harm to his employees”.

In this way, it is required that certain generic functions are carried out to avoid accidents, leaving the details as to how the functions should be carried out to the companies (or other regulated organisations). Such trends are clearly an implementation of the closed-loop, feedback design concept. In this way, detailed rule-making takes place at a level where the context is known, and this change clearly also changes the role of decision makers within the social control hierarchy. Analysis of the dynamic interpretation of safety criteria in the organisation becomes an important research issue. The change of legislation strategy clearly brings with it a very substantial change of the structure of the distributed decision making system of Fig. 1. It changes the need for information about the detailed work context at the upper levels and makes the traditional, separate safety organisation obsolete. It also clearly will change the need for interaction between regulatory decision makers and substance matter experts — see the discussion of competence below.

When safety is controlled by stating performance objectives as is the case with generic regulation, safety becomes just another criterion of a multi-criteria decision making and becomes an integrated part of normal operational decision making. In this way, the safety organisation is merged with the line organisation.

Such modern attempts to delegate decisions and to manage by objectives call for an explicit formulation of value criteria and effective means for communication of values down through society and organisations. Interesting developments have been presented for this kind of distributed organisation and formal strategies have been proposed for ‘ethical accounting’ to ensure that the impact of decisions on the objectives and values of all relevant stakeholders are adequately and formally considered (Bøgetoft and Pruzan, 1991).

8.3. Information on actual state of affairs

An important issue in a closed-loop, feedback function is the observation or measurement of the actual state of affairs and the response to control actions. No control system will perform better than its measuring channel. Important questions therefore are whether information about the *actual state* of the functions within their control domain is available to decision makers and whether this information is compatible with (comparable to) the objectives and constraints as interpreted by the agent.

The general trend toward generic legislation has raised public concerns that it is too loose and not easily enforceable or effective because it leaves so much up to the companies. It has therefore typically been reinforced by government efforts to monitor the presence of proper feed-back of performance information. Company documentation is required to demonstrate how it is implementing the rule with disclosure of the documentation to the agency which enacted the rule and to local officials and persons at risk (e.g., workers, community).

8.4. Capability and competence

The question of the content and form of the competence of the various controllers (decision makers) is extremely important. Several questions become critical, when interpretation of generic regulation is delegated to the local decision makers. Are they thoroughly familiar with the control requirements of all relevant hazard sources within their work system? Do they know the relevant parameters sensitive to control actions, and the response of the system to various control actions?

Capability or 'competence' here is not only a question of *formal knowledge*, but also includes the *heuristic know-how* and *practical skills* acquired during work and underlying the ability of an expert to act quickly and effectively in the work context.

Knowledge of the normal competence of co-operating agents is necessary to judge what information to communicate, up, down, and horizontally in the system. In particular during a period with a fast pace of change, it is important to analyse how effectively information of changes of technology, processes, and policies are communicated. This analysis identifies the control structure, the information flow *content* and serves to determine whether the decision makers *can* make the appropriate risk management decisions and whether their mutual interaction *is capable* of a coherent safety control function.

8.5. Commitment

These aspects discussed so far determine whether a decision maker *can* control safety adequately. An additional question is whether decision makers *will* act properly. This question has two aspects:

Are *priorities* right? Will decision makers be committed to safety? Is management, for instance, prepared to allocate adequate resources to the maintenance of defenses? Do regulatory efforts serve to control management priorities properly? This question also points to the influence of different time horizons of market dynamics, personal career planning, financial forecast, and protection against major accidents.

Another question is: Are decision makers *aware* of safety constraints? Considering the nature of 'naturalistic decision making', are decision makers *prompted* to consider risk in the dynamic flow of work? Are they made aware of the safety implications of their business and every day work planning decisions.

8.6. Conclusion

This analysis of the work situation of the various decision makers cannot be based on a traditional *task analysis* due to the dynamic condition of a modern work place. Instead, an analysis of the requirements and constraints of the problem space should be used, formulated at several levels of a means–ends hierarchy (Rasmussen et al., 1994; for a short introduction, see Rasmussen, 1994c). A representation of the problem space in terms of a means–ends hierarchy, as it is indicated in Fig. 6, serves to make explicit the constraints and the options for choice of the individual decision makers and, thus, to serve as a context within which the boundaries of safe choice can be made explicit and (hopefully) visible. In addition, the analysis should include a careful identification of the problem formulation and performance criteria of the individual decision makers.

This kind of analysis calls for a tightly co-ordinated analysis across all levels of Fig. 1 with a deep understanding of the subject matter of the work and related competence at each level. This is clearly a cross-disciplinary issue posing some basic problems, as discussed below.

9. Identification of constraints and safe boundaries

The success of this approach depends on whether it is possible to identify explicitly the constraints of the work system and the boundaries of acceptable operation in a dynamic society. For well structured and tightly coupled systems such as industrial power and process plants, protected by multiple, technical defenses, predictive risk analysis has been developed, capable of identifying explicitly the preconditions of safe operation, even if such analyses have mostly been shaped for acceptance of new installations, not for operational risk

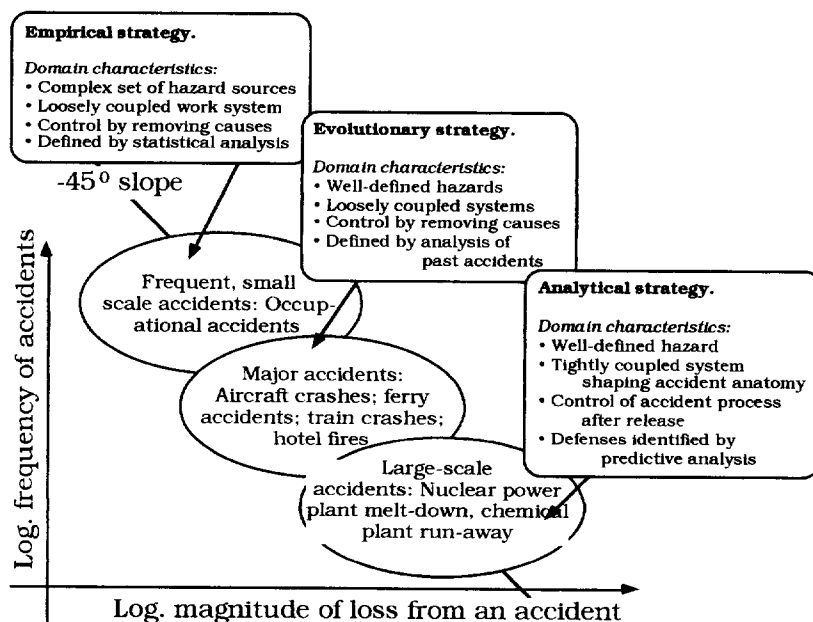


Fig. 7. Hazard source characteristics and risk management strategies.

management. (For recent reviews of the state of the art, see Leveson, 1995 and Taylor, 1994). For less structured work systems development of risk analyses for making explicit the boundaries of safe operation at all levels of the socio-technical system is still a research issue.

Depending on the nature of the hazard source, quite different risk management strategies have evolved across the different hazard domains (Rasmussen, 1993c, 1994b), see Fig. 7. A detailed study of the control requirements of the dominant hazard sources of a work system is therefore mandatory for risk management.

The three categories in Fig. 7 are characterised by their frequency of accidents and by the magnitude of loss connected to the individual accident. In a mature society, there appears to be an inverse relationship between the accepted frequency and magnitude, as shown in the figure (Johnson, 1973).

1. *Occupational safety* focused on frequent, but small-scale accidents: The hazard is related to a very large number of work processes and the level of safety over time can be directly measured by the number of LTIs (lost-time-injuries) and casualties. Consequently, the average level of safety is typically controlled empirically from epidemiological studies of past accidents.

2. Protection against *medium size, infrequent accidents*: In this category, safer systems evolve from *design improvements in response to analysis of the individual, latest major accident*. Examples are accidents such as hotel fires, aircraft accidents, train collisions, etc. Safety control is focused on the control of particular, reasonably well defined hazard sources and accident processes. Consequently, several lines of defenses against accidents have been established by an evolutionary, incremental effort toward improved safety. In this category, risk management is focused on monitoring the defenses against particular accident processes.

3. Protection against *very rare and unacceptable accidents*. The potential damage from accidents in some large-scale systems (e.g., nuclear power) is very large and the acceptable mean-time between accidents so long that design cannot be guided by empirical evidence from past accidents. Consequently, design and operation must be *based on reliable predictive models of accident processes and probability of occurrence*. For this purpose, probabilistic risk analysis (PRA) has been developed and system design is then based on the application of several, functionally independent protective systems. A full-scale accident then involves simultaneous violations of all the designed defenses. The assumption is that the probability of failure of the defenses individually can and will be verified empirically during operation even if the probability of a stochastic coincidence has to be extremely low. In this way, the reference for monitoring the performance of the staff during work is *derived from the system design assumptions, not from empirical evidence from past accidents*.

This brief review illustrates different categories of hazard sources and the related different risk management strategies that have evolved in the past. It is, however, necessary to reconsider this empirical categorisation in the present fast pace of change of the conditions of risk management.

Development of a classification system for hazard sources and their different control requirements is useful for several reasons. The empirical approach to risk management has been very effective in the past for prevention of small and medium scale accidents. Recently, however, the quest for a proactive, 'no-accident-is-tolerable' strategy is being voiced (see e.g., Visser, 1991). This quest calls for a wider application of more focused, analytical risk management strategies and a classification of hazard sources, their control requirements, and related effective risk management strategies will be necessary to select a proper risk management policy and information system for a particular hazardous installation.

A basic issue is that every work place has many different, potentially hazardous activities and, consequently, risk management adopted by a company will require a set of strategies directed toward several hazard sources. Therefore, consensus among decision makers at all levels of the socio-technical levels of Fig. 1 with respect to the characteristics of the hazard sources within a company is necessary to classify activities with reference to management strategies and the required formats for communication among decision makers for direct hazard control, for management planning, and for regulatory monitoring.

Some preliminary dimensions of a taxonomy for classification will be:

- The nature of the significant *hazard source* to control: Is it well defined as is the case with hazards from loss of control of large accumulations of energy (chemical process systems, power plants); from ignition of accumulations of inflammable material (hotel fires); or from loss of containment of hazardous material (chemical plants, transport of hazardous goods)? Or is it poorly defined and bounded as is the case for occupational safety at a construction site?
- The *anatomy of accidents* following release of the hazard: When the anatomy is well bounded by the functional structure of a stable system (such as a chemical process plant designed according to the defence-in-depth philosophy with multiple barriers), then the protection against *major accidents* can be based on *termination of the flow of events* after release of the hazard.

In contrast, when the anatomy of the accidental flow depends on the particular circumstances, then protection against accidents must be based on *elimination of the causes of release* of the hazard (capsizing of ro-ro ferries; leaks and explosions in petrochemical plants).

It should be mentioned here that the two strategies may very well be applied for the same system. For instance, loss of production due to shut-down of a process plant is expensive, and great care is therefore taken to avoid release of safety action by operator errors or insignificant technical faults. That is, protection of continuity of production depends on fighting the potential causes of a release of the hazard and the defenses.

- The degree to which *defenses can be based on predictive analysis*. The design of barriers against propagation of hazardous material or released energy in a chemical or nuclear plant today is only accepted on the basis of a predictive risk analysis demonstrating an acceptable overall risk to society. When the predicted risk has been accepted, the process model, the preconditions, and assumptions of the prediction then become specifications of the parameters of risk management. For the PRA to be operational in this way, preconditions and assumptions must be explicitly stated. It is no longer acceptable that predictive safety analysis is considered an art (Amendola, 1989). A focused further development of PRA toward an operational management tool is required. Fortunately, it is not necessary for this purpose to predict performance of operators and management. When a plant is put in operation, data on human performance in operation, maintenance, and management can be collected during operation and used for a 'live' risk analysis. Thus, predictive risk analysis for operational management will be much simpler than the analysis for a prior acceptance of the design.

The different features of hazard sources, system configurations, and related risk management strategies clearly call for carefully co-ordinated studies involving several technical and human sciences, and the models required to plan effective risk management strategies cannot be developed by the integration of the results of horizontally oriented research within the various academic disciplines across these different hazard domains, as it is sometimes assumed when universities establish risk research centres. Instead, vertical studies of the control structure are required for well bounded categories of hazard sources, characterised by

uniform control requirements. This research calls for a cross-disciplinary competency including deep familiarity with the various work domains and their operational practice. As discussed elsewhere, these requirements match poorly the academic teaching environment where the principal criteria for decisions on tenure and promotion are the number of publications in refereed journals, the number that show single rather than multiple authorship, and the frequency of citation in acknowledged, scientific journals representing the central paradigms of the particular discipline. (For a detailed discussion, see Rasmussen, 1994a, Kahn and Prager, 1994 and Albert, 1985).

10. Present trends in the paradigms of human sciences

The approach described above is based on the assumption that the behaviour of a dynamic socio-technical system cannot be represented in terms of task sequences and errors referring to a 'correct' or 'rational' performance. Fortunately, some developments within the various academic disciplines of relevance for risk research can be identified which greatly facilitate a cross disciplinary approach. The aim of the following discussion is not to present a systematic review of this wide field of research, but rather to illustrate a common trend as it has been perceived from the point of view of my problem focused research.

From a review of research paradigms it appears that the concept of 'human error', of 'decision bias' etc., is typically found in a certain phase of the evolution of descriptions of human behaviour within several professional fields. That is, description of human behaviour typically starts by identification of rational behaviour by normative models, then actual behaviour is described in terms of some kind of deviation, that is, error, with reference to the normative behaviour. Presently, the widespread acceptance of concepts related to models of cognition within several human sciences brings with it a further trend toward modelling the actual behaviour directly in terms of behaviour shaping constraints of the environment and the adaptive mechanisms of human actors in the environment.

Fig. 8 illustrates the parallel evolution of the paradigms within decision research and management research and the concurrent change of paradigms within branches of safety research.

The figure also indicates how the efforts to control the behaviour of individuals, organisations, and society shift along with the concepts underlying behavioural research. In the following sections, the disciplines included in the figure are briefly discussed.

10.1. Decision research

Normative, prescriptive theories and models	Economic decision theory (von Neumann, Morgenstern) Expected utility theory (Keeny, Raiffa)
Descriptive models in terms of deviations from norms	Psychological decision theory: Prospect theory, Judgement biases (Tversky, Kahnemann).
Descriptive models of actual behaviour	Social judgement theory, cue-utilisation (Hammond, Brehmer). Dynamic and natural decision making (Brehmer, Klein)

In decision making research, the shift in research paradigm from normative, prescriptive models over descriptive models in terms of deviation from rational performance towards

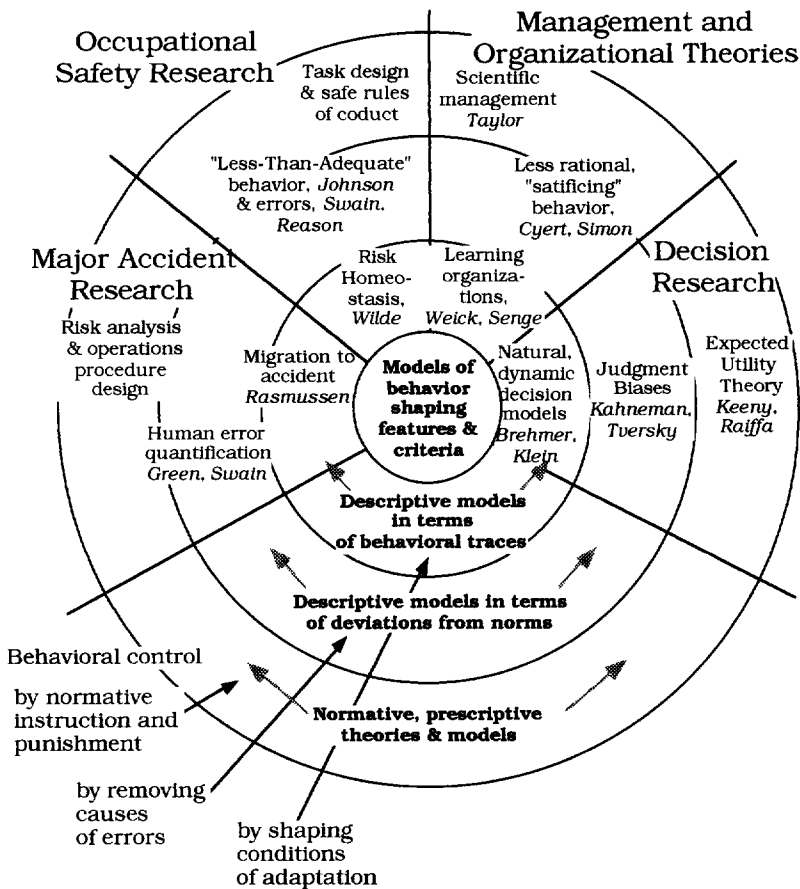


Fig. 8. The convergence of human science paradigms toward models in terms of behaviour-shaping work features and subjective performance criteria.

modelling actual behaviour is very visible in the different schools of decision research, (see, e.g., the review in Hammond et al. (1980).

10.1.1. Normative models

Classic *decision theory* is a normative theory based on the expected utility theory developed by economists and mathematicians (von Neuman and Morgenstern, 1944). The approach is focused on decision making from a prescriptive point of view and makes no claim to represent the actual information processing of human decision makers. The emphasis is not on what they do, but what they should do. Later theorists (Keeney and Raiffa, 1976) emphasise the mathematical modelling of subjective probability and utility and promote the use of the theory to aid decision makers to achieve logical consistency.

To account for the behaviour of practical decision making, this theory was supplemented by *behavioural decision theories* initially developed by Edwards and Tversky (1976) and

based on the Bayesian probability theory and on experiments to describe how closely human information processing approximates to the Bayesian process, and how information is used to revise subjective probabilities.

10.1.2. Descriptive models in terms of deviation from normative behaviour

Psychological decision theory is based on the work of Tversky and Kahneman (1974) with their concepts of representativeness and availability bias, and the use of heuristics. Their approach rejects the use of optimal decisions as frame of reference for description. It turns instead to a search for the psychological mechanisms that people use to evaluate frequencies and likelihood. Psychological decision theorists generally consider subjective utility theory to be empirically unsuccessful and to be replaced by a theory that *explains* human behaviour in terms of psychological biases with reference to rational behaviour.

10.1.3. Descriptive models of actual behaviour

Social judgement theory can be seen as a further development towards modelling actual behaviour within decision research (For a review of this approach see Hammond et al., 1980. This theoretical development as well as the basic concepts of 'ecological validity of cues' and 'utilisation of cues' originate in Brunswik's theory of perception (Brunswik, 1952). The primary intention of the approach is not to explain, but to *describe* human judgement processes, and to provide guides for the development of decision aids. The basic framework of the social judgement theory is Brunswik's 'lens model' which has also been the basis of research concerning diagnostic judgements in several professional activities such as, for instance, stockbrokers, clinical psychologists and physicians (Brehmer, 1987).

Much of decision research so far has been based on well-structured and carefully planned laboratory experiments. Furthermore, 'decisions' have been perceived as discrete event: an actor realises that a problem has turned up, then a situation analysis is performed to diagnose the situation and define the problem, this is followed by an evaluation of the present goals, and an action is planned. Recently, however, human interaction with the environment is increasingly being considered to be a continuous control task. Separate 'decisions' therefore are difficult to identify. The state of the environment is perceived in terms of action possibilities, actions are chosen, and the resulting response from the environment is perceived as a background for the next action. Errors then are difficult to localise in the stream of behaviour; the effects of less successful acts are a natural part of the search for optimal performance. In this way, further developments of decision theories can be seen toward direct description of the actual behaviour based on analysis of behaviour/performance in complex work environments with little emphasis on the identification of errors or biases with reference to normative models. One line of development is research on *naturalistic decision making* (for a review see Klein et al., 1994; another is research on *dynamic decision making* (Brehmer, 1992).

It is interesting to note that models at higher levels in terms of characteristics of the environment are found when decision theories evolve into models of cognitive control of action. A 'field theory' of human performance was early proposed by Lewin in 1951 but unfortunately, his concepts do not appear to solve our problems. As Brunswik noted, Lewin's concepts are post-perceptive and pre-operational (see Lewin, 1951). The theories of invariants and affordances of Gibson (1966, 1979) seem to point in a more fruitful direction in our context.

10.2. Organisational theory.

Normative, prescriptive theories and models	Scientific management theories (Taylor, Gulick, Urwick, Weber).
Descriptive models in terms of deviations from norms	Satisfying choice, bounded rationality (Simon, Cyert, March)
Descriptive models of actual behaviour	Learning organisations (La Porte, Rochlin, Senge, Weick)

10.2.1. Normative theories

Normative, rational models take different shapes. *Scientific Management* (Taylor, 1911) is focused primarily on manufacturing and similar production activities. It employs economic efficiency as ultimate criterion and seeks to maximise efficiency by rational planning procedures. Goals are known; tasks are repetitive. *Administrative Management* (Gulick and Urwick, 1937) assumes that ultimately a master plan is known, against which specialisation, departmentalisation, and control are determined. *Bureaucratic Models* (Weber, 1947), follow similar patterns and the complexity of human components is controlled by divorcing private life and office roles by means of rules, salary and career. In all these models, the behaviour of the individual is controlled by normative rules, pre-planned by system design and reinforced by training, instruction and punishment of the staff and planning of the work place to induce proper work performance (time-and-motion studies, assembly lines).

This perception of management underlay until recently the command and control concepts that have governed industrial organisation and the development of centralised planning tools (such as MRP — Materials and Resource Planning — systems), a situation characterised by Savage and Appleton (1988) as second generation management of fifth generation technology.

10.2.2. Models in terms of deviation from normative behaviour

Following after the scientific management paradigm, focus was directed toward study of organisations, based on the *informal organisations view* that sentiments, cliques, and social control by informal norms are patterned, adaptive responses in problematic situations (Roethlisberger and Dickson, 1939), and the *interaction view* in which organisations are not autonomous entities but interact intimately with other organisations and the public: Even the best laid plans of managers have unintended consequences and are conditioned or upset by other social units (Barnard, 1938; Selznick, 1949; Clark, 1956). Such considerations switch the focus of study onto the study of particular organisations in terms of biases and deviation from rational behaviour (Simon, 1957a; March and Simon, 1958; Cyert and March, 1963). Organisations must, to cope with complexity, develop processes for searching and learning, as well as for deciding; decisions are *satisficing* rather than maximising, and are based on 'bounded rationality' (Simon, 1957b); decision makers are not rational, but 'muddle through' the work requirements.

10.2.3. Models in terms of features shaping organisational behaviour

The line of research introduced by Simon et al. is continued by Thompson (1967). The basis of Thompson's approach is Parson's suggestion that organisations exhibit three distinct levels of responsibility and control: technical, managerial, and institutional (Parsons, 1960). Parsons and Thompson introduce a control point of view that later is explicitly formulated in 'distributed decision making' models of organisational behaviour (see Rasmussen et al., 1991),

Parsons stresses the qualitative break in the simple continuity of 'line' authority because the functions at each level are qualitatively different.

"Those on the lower levels are not simply lower-order spellings out of the higher-level functions. The articulation of levels and of functions rests on a two-way interaction with each side withholding its important contribution, in a position to interfere with the function of the other and of the larger organisation."

Self-organising features of organisations similar in nature to Thompson's synthetic organisations have been studied by the high reliability organisations program of the Berkeley group (See e.g., Rochlin et al., 1987). The group has studied various successful organisations by very careful field studies with the team participating in lengthy missions. They conclude that the high reliability is explained by a high degree of learning and self-organisation which is facilitated by the rapid turn-over of staff in the navy, in contrast to prior expectation.

In a study of aircraft carriers, they have modelled the shift in organisational structure forced by changes in the actual operational situation. They were able to identify three distinctly different organisational modes: a formal rank organisation, a high tempo operational organisation and an emergency organisation. The high tempo organisation was characterised by having a very high degree of functional redundancy, protecting against errors and resource violations.

In all, learning organisations, as stressed by Weick (1977) and Senge (1990a,b) becomes an important research topic. In all of this there is also the desirability for ensuring organisational learning at all levels. Indeed in order that the entire organisation or system shall be adaptive, it is absolutely necessary to ensure interaction between the levels. It is not enough for a change in plans or actions to occur: the entire organisation must be made aware of it, and aware of its implications for the various levels and subgroups within the organisation, in each case in a language appropriate to the kind of decisions which need to be made by that part of the organisation.

10.3. Occupational safety research

Normative, prescriptive theories and models	Task design and safe rules of conduct
Descriptive models in terms of deviations from norms	Studies of errors (Altman, Chapanis, Gibson, Rasmussen, Rigby, Rook, Swain). 'Less-Than-Adequate' management behaviour (Johnson). Management errors, resident pathogens (Reason)
Descriptive models of actual behavior in terms of behavioural traces	Risk homeostasis (Wilde) Migration toward accidents (Rasmussen)

10.3.1. Models in terms of normative practice.

During the period when the 'scientific management theory' and behavioristic theories of workers' performance were predominant, task behavior was typically controlled by prescriptive instruction and by motivation campaigns. Response to errors was usually to introduce improved training and selection schemes to eliminate 'error-prone' individuals and have the rest 'trying harder' through 'zero-defects' programs.

10.3.2. Models of behavior in terms of normative performance and errors

During the 60s and 70s this picture is changing. Researchers and human factors specialists are looking for a description of the work process and the origin of accidents and, therefore, an interest in the concept of human error emerges. Accidents were related to human error by several Human Factors specialists, e.g., Altman (1970), Chapanis (1970) and Christensen

(1972) during the 60s, and the use of the human error concept in accident analysis was developed through a series of studies of home accidents by the American Research Institute.

In the search a classification of accident processes, the concepts of energy and barriers as a basis for accident analysis were introduced by Gibson (1961) who looked for a behavioural approach to safety analysis and by Haddon (1966) who sought a basis for preventive medicine.

The general reliance on an energy concept to control occupational safety clearly has weaknesses, when the concept of energy is generalised to include various categories of chemical, electrical and kinetic energy forms and the process of release of energy is generalised to include falling, nipping, shearing, cutting, etc. In that case, the energy model cannot lead to one well defined strategy for control of hazards. However, the energy concept proved effective as an attention focusing and search guiding strategy.

When models of the accident process, such as propagation of energy releases, are formulated in terms of accidental courses of events, then focus will be on the behavioural sequences of the actors involved. This model of accident causation and release processes introduces systems thinking into occupational safety and invites a transfer of the concepts and methods developed for high hazard systems into the general occupational safety work. This route was taken by Johnson (1973) from the US National Safety Council through a contract with the US Atomic Energy Commission to “formulate an ideal, comprehensive systems concept of accidents and their control, and to test the usefulness of the concept in two ways: 1) Design of an improved accident report and analysis technique; 2) Development of improved measurements of safety performance”.

Johnson adopted the results of the systematic analysis of human error and attempts to establish data bases that were systematically pursued by Rigby (1970), Rook (1962), and Swain (1963) for study of human reliability in nuclear weapons manufacturing during the early 60s. Production of nuclear weapons is a tightly controlled and pre-planned process and description of the actual human behavior in terms of errors is an effective work planning and control parameter.

Rigby (1970) discusses in detail the reference point for judgement of errors and points to the fact that such references are very context and situation dependent:

“Thus, in the most general and most practical sense, a human error is any member of a set of human actions that exceeds some limit of acceptability. An error is only an out-of-tolerance action, and the most important ingredients in any discussion of error are definitions of both the set of actions and the tolerance limits that define errors for those actions.”

Thus, it is necessary to consider that an error can only be found if a standard of judgement exists and whether or not an act is judged an error depends on the perspective and reference for judgement chosen.

The subjective and context depending aspects of error analysis have later been further developed (Rasmussen, 1990a, 1993c) with a focus on the ambiguity found in any causal analysis. In particular, the stop rule applied to identify ‘root causes’ depends on the aim of the analyst (to understand behavior, to punish, or to improve system safety).

The combination of the two basic views that (1) accidents should be understood in terms of an energy related process and (2) hazard management therefore should be directed toward planning of the release route led Johnson (1973) to focus on the management as being responsible for the planning of the context within which accidents unfold, that is, he stressed the role of ‘less than adequate’ management decisions and developed MORT — the ‘Management Oversight and Risk Tree’ tool for accident analysis.

Later Reason (1990) has focused analysis on management errors and organisational factors, such as 'resident pathogens' making organisations vulnerable to accidents.

10.3.3. Models of actual behavior in terms of behavior shaping features

Efforts to improve safety by counteracting the human error sources identified by causal analysis of accidents tend to be ineffective. Causal trees are not models of the system functionality, only records of particular cases and, they do not include the effect of the compensating adaptation of the people involved. Consequently, a trend can be seen toward study of occupational safety in terms of models of actual behaviour without reference to 'errors'.

The need for such a research direction is clearly demonstrated by the observation that human adaptation frequently compensates for attempts to improve system safety. Such compensation has been found in traffic safety for instance in response to anti-blocking car brakes (Aschenbrenner et al., 1986; Status, 1994), and introduction of separate bicycle paths (Ekner, 1989). In psychological traffic research, this tendency has been referred to as 'risk homeostasis', that is, adaptation seeking to maintain a stable level of perceived risk (Wilde, 1976, 1985). This finding can be an artefact caused by a too narrow focus of modelling behavior from accident and error analysis. According to the more general concept of adaptation illustrated in Fig. 3, performance is likely to be maintained close to the boundary to loss of control, in a kind of 'homeostasis', being controlled by perception of dynamic control characteristics of the interaction not by an abstract variable such as 'risk'.

A similar point of view has been stressed by Taylor who applies a hermeneutic analysis of traffic accidents. He observed that drivers tend to try to keep their arousal at a desired, constant level and, consequently, to go faster if conditions become too undemanding, in order to generate more arousing incidents (Taylor, 1981). This point of view matches the arguments presented above very well. The consequence drawn by Taylor (1987) is that traffic safety is hard to improve beyond a certain limit. He criticises the present 'mechanistic' approach and argues that accidents cannot be studied in terms of causes, but should be analysed in terms of reasons.

The conclusion from this discussion is that safety in the direct interaction with the work environment must be based on an identification of the boundary of safe performance by analysis of the work system, and the criteria that drive the continuous adaptive modification of behavior. Efforts for improvement must then be directed toward the control of performance in interaction with the boundary, not on control of errors. In other words, as argued in this paper, we need models of the behaviour of individuals and organisations at a higher level than tasks and acts. This points to basic research in cognitive control of behaviour, such as paradigms mentioned in the previous sections.

10.4. Major accident research

Normative, prescriptive theories and models	Risk analysis and design of standard operating procedures
Descriptive models in terms of deviations from norms	Human error analysis and quantification (Green, OECD group, Swain) Management errors, resident pathogens (Reason)
Descriptive models of actual behavior	Migration toward accidents (Rasmussen) Studies of process operators (Moray, Rasmussen, Rouse, Woods). Studies of organisational behavior (LaPorte, Rochlin)

10.4.1. Normative task models and work instructions

The defence-in-depth protection based on multiple barriers and redundant protective devices was developed systematically for nuclear systems from the very beginning, because small scale experiments are impossible (the minimum critical mass problem). The predictive risk analysis, based on stochastic models was initially developed for evaluation of redundant safety systems, including the influence of test and repair policies on their availability (Siddall, 1954; Jacobs, 1957). In Europe, studies of the reliability of safety systems, based on stochastic models (Rasmussen and Timmermann, 1962) showed that safety systems based on redundancy could easily be designed to have extreme reliability given the failure rates of the available components, as long as certain constraints on repair policy and quality are respected.

Work instructions for maintenance personnel and operating procedures for operators were developed from the functional design and the related technical reliability analysis. Performance was controlled by formal standard operating procedures and effective training (use of simulators).

10.4.2. Models in terms of deviations from normative performance

However, critical human factors problems were identified that initiated studies of the role of man-machine interaction and human reliability. Very influential were the initiatives of Eric Green of the United Kingdom Atomic Energy Authority (UKAEA) who started work on quantification of operator errors, inspired by Alan Swain's work in nuclear weapon production (Green et al., 1968; Green and Bourne, 1972).

The developments by Swain have been very influential for predictive risk analysis of nuclear power installations (Swain, 1963; Swain and Guttman, 1983). The reliability estimation method — THERP (Technique for Human Error Rate Prediction) — is focused on prediction of the probability that a given task procedure is terminated successfully based on data on successful performance of the individual task elements. In general, prediction was limited to prediction of the probability of success in a task, and therefore can be based on gross error data, i.e. whether an act was successfully performed or not. This is the primary aim of analysis of performance related to industrial production and military missions, and the approach will be adequate as long as only routine tasks are considered for which a stable level of training can be reached and, therefore, a 'normal' task sequence identified as reference for judgement of errors. This is actually the case for maintenance of the defences of industrial process plants and the methods for human reliability prediction have been very important in this context.

Also our analyses of incidents and accidents in nuclear power (Rasmussen, 1969) pointed to the problem that incidents were, to a large extent, related to human error, not to technical failures in safety systems. However, studies of operators' roles in industrial accidents (Rasmussen, 1969) and the cases found in the 'licensee event reports' to the US nuclear authorities (Rasmussen, 1980) showed that difficulties in performing a reliable diagnosis during disturbances also contribute to accidents. Not only is human reliability during pre-planned tasks important, we also needed to include cognitive tasks during rare disturbance in a human reliability analysis. This situation opened two new avenues of research. One is the study of operator behavior (human-machine interaction), diagnostic strategies, and proper interface design to control human error mechanisms. Another issue was that it became clear (Rasmussen, 1979) that it was not sufficient to predict the probability of success in a task, the probability of particular, risky erroneous actions should also be estimated and a more elaborate taxonomy was needed.

The consequence of this was a need for a structured taxonomy to distinguish between different human 'error modes' to be used in planning of practical data collection. Therefore, initiatives were taken by Eric Green (UKAEA) to organise a series of expert working groups on 'Rare Event Analysis' (1976) and on 'Human Error Analysis and Quantification' (1977–81) under the Committee of the Safety of Nuclear Installations (CSNI) of the International Organisation for Economic Cooperation and Development (OECD). The aim of these workshops were to suggest a taxonomy of human error based on detailed analysis of event reports and actual operator performance in power plants. In spite of such efforts toward useful databases of human error rates, reliable error data remained scarce, partly because of difficulties in establishing the proper on-site analysis of data, partly because of difficulties in defining and quantifying the denominators of error rates, the frequency of opportunities for error.

The OECD taxonomy (Rasmussen et al., 1981) was an attempt to describe 'errors' in terms of human 'malfunctions' due to normally successful psychological mechanisms that, when triggered under certain circumstances, led to failed task performance. In addition, the event and field studies threw light on the importance of 'latent errors' left over from inappropriate test, calibration and maintenance work and the need to use the assumptions made for the predictive risk analysis as a basis for active, operational risk management (Rasmussen, 1982).

One major problem with this approach is that it is only applicable for those installations and tasks for which empirical data are available. Consequently, several attempts have been made to get error data from other sources. One approach has been generation of data by 'expert judgement'. Subject matter experts are asked in a structured way to generate error frequencies and the data are then tested and verified in different ways. This approach has been advanced in particular by Embrey et al. (1984).

In spite of a widely used 'defense-in-depth' design strategy, most recent major accidents in large scale industrial systems appear to be caused by operation of the systems outside the predicted preconditions for safe operation during critical periods with excessive pressure from a competitive environment (Rasmussen, 1993b). In consequence, an important research issue is the adaptive behavior of organizations under pressure.

10.4.3. Models in terms of actual behavior of individuals and organisations

The analyses of accidents in this way raised a fundamental interest in two complex topics: (1) modelling operator behavior and understanding the difficulties faced by operators and maintenance staff in the existing systems and (2) the role of management practices on system safety.

1. *Operator Modelling and Human–Machine Studies.* The studies of NRC licensee event reports underlying the structure of the OECD taxonomy also contributed to a model of operators cognitive control of activities under routine and rare situations (Rasmussen, 1980). This approach aimed at a control theoretic model to explain the error modes found in the event reports in terms of failure of control processes at a signal processing, skill-based, a sign responding, rule-based level, and a symbolic, problem solving level. This model served also to facilitate communication with basic psychological research.

During the late 60s and 70s a widespread activity was emerging on modelling process operator performance during disturbed plant operation. These efforts were reflected in a series of NRC/IEEE workshops on operator modelling and control room design (see Schmall, 1979; Hall et al., 1981; Abbott, 1982) and in several NATO Advanced Study Institutes (for a review

see: Sheridan and Johannsen, 1976, on supervisory control, Moray (1977) on work load, Rasmussen and Rouse (1981) on diagnostics.

2. *Management Practices*. The success of the defence in depth design strategy depends on an effective risk management, a topic taken up by the NAS group of experts on 'Human Factors Research Needs in Nuclear Power', (Moray and Huey, 1988) and the World Bank workshops on 'Safety Control and Risk Management' (Rasmussen and Batstone, 1989). These cross disciplinary workshops were focused on the role of human error in management and a close interaction emerged with organisational and management research to understand the problem of systemic degradation of defences and the requirements for the actual functioning of organisations in high hazard systems (see the discussion also in the section on management and organisation research).

10.5. Implications for research and for control of human behaviour

As it is illustrated in the Fig. 8, convergent parallel changes of the paradigms of several disciplines are shaped by the increased interest in cognitive, intentional concepts, shown in the centre of the diagram, replacing the past focus on the mechanistic, normative approaches at its periphery. This is clearly a most promising precondition for the proposed cross-disciplinary approach to the problem of failure of socio-technical systems.

It is important to consider that the changing conceptions used for representation of human behaviour have important implications for the development of means to control the behaviour of individuals and organisations, and thus for effective risk management. In this respect, approaches to design of work systems based on principles derived from ecological psychology are important candidates (Rasmussen and Vicente, 1990; Vicente and Rasmussen, 1992).

11. Conclusion

Industrial risk management, environmental protection, and life-cycle engineering for a modern, dynamic, and highly integrated society raise some very basic problems related to the behaviour of adaptive socio-technical systems subject to a fast pace of change.

Specific research issues are development of a taxonomy of hazard sources and their control requirements, the vertical interaction among decision makers at all levels of society as it is found for the various hazard domains, and the influence from naturalistic decision making modes on their responses to change and to competitive pressure. These questions call for a cross-disciplinary research — which is not synonymous with a need only for applied research. We are faced with several basic research issues in this context.

Unfortunately, this mode of research matches poorly the constraints of academic research and the present funding policies of research councils. In this situation an issue is the creation of a cross-disciplinary research community that can cope with complex long term research issues without the constraints of academic institutes and their focus on short term, tenure strategies.

Hopefully, the convergence of paradigms found within several technical and academic disciplines will facilitate cross-disciplinary research in the future and an inspiration could be the approach the MacArthur Foundation has successfully applied for research in public health care (Kahn and Prager, 1994). A good foundation for this is the world-wide multi-disciplinary network created by the annual Bad Homburg workshops on 'New Technology and Work' (Wilpert, 1987-).

References

- Abbott, L.S. (Ed.) (1982) *Proceedings of Workshop on Cognitive Modelling of Nuclear Plant Control Room Operators*, Dedham, Massachusetts, August 1982. NUREG/CR-3114. U.S. Nuclear Regulatory Commission, Washington, D.C.
- Albert, D. (1985) Performance and Paralysis: The Organisational Context of the American Research University. *Journal of Higher Education* **56** (3), 243–280.
- Altman, J.W. (1970) Behavior and Accidents. *Journal of Safety Research*, September.
- Amendola, A. (1989) Planning and Uncertainties. In *Proceedings of the 2nd World Bank Workshop on Risk management and Safety Control*. Karlstad, Sweden, Rescue Services Board.
- Aschenbrenner, K.M., Biehl, B. and Wurm, G.M. (1986) Antiblockiersystem und Verkehrssicherheit: Ein Vergleich der Unfallbelastung von Taxen Mit und Ohne Antiblockiersystem. (Teilbericht von die Bundesanstalt für Strassenwesen zum Forschungsproject 8323: Einfluss der Risikokompensation auf die Wirkung von Sicherheitsaussnahmen). Mannheim, Germany. Cited in Wilde, G.S. (1988) Risk Homeostasis Theory and Traffic Accidents: Propositions, Deductions, and Discussion in Recent Reactions. *Ergonomics* **31**, 441–468.
- Baram, M. (1996) Generic Strategies for Protecting Worker Health and Safety: OSHA's General Duty Clause and Hazard Communication Standard. *Occupational Medicine: State of the Art Reviews* **11** (1), January–March.
- Barley, S.R. (1988) On Technology, Time, and Social Order: Technically Induced Change in the Temporal Organization of Radiological Work. In *Making Time*, ed. F.A. Dubinskias. Temple Univ. Press, Philadelphia.
- Barnard, C.I. (1938) *The Function of the Executive*. Harvard University Press, Cambridge, MA.
- Bøgetoft, P. and Pruzan, P. (1991) *Planning with Multiple Criteria*. North-Holland, Amsterdam.
- Brehmer, B. (1987) Models of Diagnostic Judgements. In *New Technology and Human Error*, ed. J. Rasmussen, K. Duncan and J. Leplat. Wiley and Sons, New York.
- Brehmer, B. (1992) Dynamic Decision Making: Human Control of Complex Systems. *Acta Psychologica* **81**, 211–241.
- Brunswik, E. (1952) *The Conceptual Framework of Psychology*. Chicago University Press, Chicago.
- Chapanis, A. (1970) The Error Provocative Situation. *Symposium on Measurement of Safety*. National Safety Council.
- Christensen, J. (1972) Overview of Human Factors in Design. *National Safety Congress*.
- Clark, B.R. (1956) *Adult Education in Transition*. University of California Press, Berkeley.
- Cyert, R.M. and March, J.G. (1963) *A Behavioral Theory of the Firm*. Prentice Hall, Englewood Cliffs, NJ.
- Edwards, W. and Tversky, A. (1976) *Decision Making*. Penguin Books, Baltimore.
- Ekner, K.V. (1989) On: Preliminary safety related experiences from establishment of bicycle paths in Copenhagen, 1981–83. (Technical Report, in Danish).: Stadsingniørens Direktorat, Copenhagen.
- Embrey, D.E., Humphreys, P., Rosa, E.A., Kirwan, B. and Rea, K. (1984) Slim-Maud: an approach to assessing human error probabilities using structured expert judgement. NUREG/CR-3518 BNL-NUREG-51716.
- Engwall, L. (1986) Newspaper Adaptation to a Changing Social Environment: A Case Study of Organizational Drift as a Response to Resource Dependence. *European Journal of Communication* **1**, September, pp. 327–341.
- Estonia (1995) Accident Investigation Report; Part Report covering technical issues on the capsizing on 28 September 1994 in the Baltic Sea of the ro-ro passenger vessel MV ESTONIA. The Joint Accident Investigation Commission. Stockholm: Board of Accident Investigation.
- Flach, J., Hancock, P., Caird, J. and Vicente, K. (Eds.) (1994) *Ecology of Human–Machine Systems: A Global Perspective*. Lawrence Erlbaum, Hillsdale, NJ.
- Fujita (1991) What Shapes Operator Performance? JAERI Human Factors Meeting, Tokyo, November, 1991. To appear in *International Journal of Man–Machine Studies*: Data, Keyholes for the Hidden World of Operator Characteristics.
- Gibson, J.J. (1961) Contribution of Experimental Psychology to Formulation of the Problem of Safety. In *Behavioural Approaches to Accident Research*. Association for the Aid of Crippled Children, London.
- Gibson, J.J. (1966) *The senses considered as perceptual systems*. Houghton Mifflin, Boston.
- Gibson, J.J. (1979) *The ecological approach to visual perception*. Houghton Mifflin, Boston.
- Gibson, J.J. and Crooks, L.E. (1938) A Theoretical Field Analysis of Automobile Driving. *The American Journal of Psychology* **LI** (3), 453–471.
- Green, E. and Bourne, A.J. (1972) *Reliability Theory*. Wiley, London.
- Green, E., Marshall, J. and Murphy, T. (1968) Preliminary Investigation into the Time Response of Operators. (Internal Document).: UKAEA.
- Gulick, L. and Urwick, L. (Eds.) (1937) *Papers on the Science of Administration*. Institute of Public Administration, New York.

- Haddon, W. Jr. (1966) *The Prevention of Accidents. Preventive Medicine*. Little, Brown and Co., Boston, MA.
- Hale, A.R., Kirwan, B. Guldenmund, F. and Heming, B. (1996) Capturing the River: Multi-level Modelling of Safety Management. In *Second ICNPO Conference on Human Factors and Safety*. Berlin, November. To be published.
- Hall, R.E., Fragola, J.R. and Luckas, J.L. (Eds.) (1981) *Proceedings of the 1981 IEEE Standards Workshop on Human Factors and Nuclear Safety*, Myrtle Beach, August–September. IEEE, New York.
- Hammond, K.R., McClelland, G.H. and Mumpower, J. (1980) *Human Judgment and Decision Making*. Hemisphere Publishing, Frederick A. Praeger, New York.
- Jacobs, I.N. (1957) Safety Systems for Nuclear Power Reactors. AIIE-Pacific General Meeting. Paper 57–906.
- Johnson, W. (1973) MORT The Management Oversight and Risk Tree Analysis. (Technical Report SAN 8212).: Atomic Energy Commission, Washington, US.
- Kahn, R.L. and Prager, D.J. (1994) Interdisciplinary Collaborations are a Scientific and Social Imperative. *The Scientist*, July 11.
- Keeney, R.L. and Raiffa, H. (1976) *Decisions with Multiple Objectives, Preferences and Value Trade-offs*. John Wiley and Sons, New York.
- Klein, G., Orasanu, J., Calderwood, R. and Zsombok, C.E. (Eds.) (1994) *Decision Making in Action: Models and Methods*. Ablex, Norwood, NJ.
- Leveson, N.G. (1995) *Safeware: System Safety and Computers*. Addison-Wesley, Reading, MA.
- Lewin, K. (1951) *Field Theories in Social Science*. Harper and Row, New York.
- March, J.G. and Simon, H.A. (1958) *Organizations*. Wiley, New York.
- Moray, N. (Ed.) (1977) *Mental Workload*. Plenum Press, New York.
- Moray, N. and Huey, B. (Eds.) 1988. *Human Factors Research and Nuclear Safety*. National Academies Press, Washington, DC.
- Parsons, T. (1960) *Structure and Process in Modern Society*. The Free Press of Glencoe, New York.
- Rasmussen J. and Rouse, W.B. (Eds.) (1981) *Human Detection and Diagnosis of System Failures*. Plenum Press, New York.
- Rasmussen, J. (1969) Man–Machine Communication in the Light of Accident Record. Presented at *International Symposium on Man–Machine Systems*, Cambridge, September 8–12. In *IEEE Conference Records*, 69C58-MMS, Vol. 3.
- Rasmussen, J. (1979) Notes on Human Error Analysis and Prediction. In *Synthesis and Analysis Methods for Safety and Reliability Studies*, ed. G. Apostolakis and G. Volta. Plenum Press, London.
- Rasmussen, J. (1980) What can Be Learned from Human Error Reports. In *Changes in Working Life*, ed. K. Duncan, M. Gruneberg and D. Wallis. John Wiley and Sons, New York.
- Rasmussen, J. (1982) Human Factors in High Risk Technology. In *High Risk Safety Technology*, ed. E.A. Green. John Wiley and Sons, London.
- Rasmussen, J. (1983) Skill, Rules and Knowledge; Signals, Signs, and Symbols, and other Distinctions in Human Performance Models. *IEEE Transactions on Systems, Man and Cybernetics SMC-13* (3).
- Rasmussen, J. (1990) Human Error and the Problem of Causality in Analysis of Accidents. *Phil. Trans. R. Soc. Lond. B* **327**, 449–462.
- Rasmussen, J. (1990) The role of Error in Organizing Behavior. *Ergonomics* **33** (10/11), 1185–1190.
- Rasmussen, J. (1992) Use of Field Studies for Design of Work Stations for Integrated Manufacturing Systems. In *Design for Manufacturability: A Systems Approach to Concurrent Engineering and Ergonomics*, ed. M. Helander and M. Nagamachi. Taylor and Francis, London.
- Rasmussen, J. (1993) Deciding and Doing: Decision Making in Natural Context. In *Decision Making in Action: Models and Methods*, ed. G. Klein, J. Orasano, R. Calderwood, and C.E. Zsombok. Ablex Publishing, Norwood, NJ.
- Rasmussen, J. (1993) Market Economy, Management Culture and Accident Causation: New Research Issues? In *Proceedings Second International Conference on Safety Science*. Meeting Budapest Organizer Ltd, Budapest.
- Rasmussen, J. (1993) Perspectives on the Concept of Human Error. Keynote address at: *Human Performance and Anaesthesia Technology*. Society for Technology in Anaesthesia. Conference: New Orleans, February.
- Rasmussen, J. (1994) Complex Systems, Human Factors, and Design of Teaching Curricula. Invited contribution to Festschrift for Professor Bernotat. In *Mensch Maschine Systeme und Neue Informationstechnologien*, ed. K.P. Gärtner, W. Stein and H. Widdel. Verlag der Augustinus Buchhandlung, Aachen.
- Rasmussen, J. (1994) Risk Management, Adaptation, and Design for Safety. In *Future Risks and Risk management*, ed. N.E. Sahlin and B. Brehmer. Kluwer, Dordrecht.

- Rasmussen, J. (1994) Taxonomy for Work Analysis. In *Design of Work and Development of Personnel in Advanced Manufacturing. Human Factors in Advanced Manufacturing*, ed. G. Salvendy and W. Karwowski. Wiley-Interscience, New York.
- Rasmussen, J. and Batstone, R. (Eds.) (1989) *Why do Complex Organizational Systems Fail? Summary Proceedings of a Cross Disciplinary Workshop on "Safety Control and Risk Management"*. Word Bank, Washington, DC.
- Rasmussen, J. and Vicente, K.J. (1990) Ecological Interfaces: A Technological Imperative in High tech systems? *International Journal of Human Computer Interaction* 2 (2) 93–111.
- Rasmussen, J. and Timmermann, P. (1962) Safety and Reliability of Reactor Instrumentation with Redundant Instrument Channels. (Risø Report No. 34). January.
- Rasmussen, J., Pejtersen, A.M. and Goodstein, L.P. (1994) *Cognitive Systems Engineering*. Wiley, New York.
- Rasmussen, J., Brehmer, B. and Leplat, J. (Eds.) (1991) *Distributed Decision Making: Cognitive Models for Cooperative Work*. John Wiley and Sons, London.
- Rasmussen, J., Pedersen, O.M., Mancini, G., Carnino, A., Griffon, M. and Gagnolet, P. (1981) Classification System for Reporting Events Involving Human Malfunction. Risø-M-2240.
- Reason, J.T. (1990) *Human Error*. Cambridge University Press, Cambridge.
- Rees, S. and Rodley, G. (Eds.) (1995) *The Human Costs of Managerialism: Advocating the Recovery of Humanity*. Pluto Press of Australia, Leichhardt, NSW.
- Rigby, L.W. (1970) Nature of Error. (Technical Report): Sandia Lab.
- Rochlin, G.I., La Porte, T.R. and Roberts, K.H. (1987) The Self Designing High Reliability Organization: Aircraft Carrier Flight Operations at Sea, *Naval War College Review*, Autumn.
- Roethlisberger, F.J. and Dickson, W.J. (1939) *Management and the Worker*. Harvard University Press, Cambridge, MA.
- Rook, L.W. (1962) Reduction of Human Error in Industrial Production. (Technical Report): Sandia Lab, June.
- Savage, C.M. and Appleton, D. (1988) CIM and Fifth Generation Management. In *Fifth Generation Management and Fifth Generation Technology*. SME Blue Book Series. Society of Manufacturing Engineers, Dearborn, Michigan.
- Schiavo, M. (1997) *Flying Blind, Flying Safe*, New York: Avon Books. For reviews see *TIME Magazine*, 31 March, 1997. pp. 38–48 and 16 June, pp. 56–58, 1997.
- Schmall, T.M. (Ed.) (1979) *Proceedings of the 1979 IEEE Standards Workshop on Human Factors and Nuclear Safety*. Myrtle Beach, December. IEEE, New York.
- Selznick, P. (1949) *TVA and the Grass Roots*. University of California Press, Berkeley, CA.
- Senge, P.M. (1990) *The Fifth Discipline: The Art and Practice of The Learning Organization*. Doubleday Currency, New York.
- Senge, P.M. (1990) The Leader's New Work: Building Learning Organizations. *Sloan Management Review* 7, Fall.
- Shell, (1992) A Study of Standards in the Oil Tanker Industry. Shell International Marine Limited, May.
- Sheridan, T.B. and Johannsen, G. (Eds.) (1976) *Monitoring Behaviour and Supervisory Control*. Plenum Press, New York.
- Siddall, E. (1954) *A Study of Serviceability and Safety in the Control System of the NRU Reactor*. (Technical Report AECL 399) (CRNE 582). AECL, Toronto.
- Simon, H.A. (1957) *Administrative Behavior*. Macmillan, New York.
- Simon, H.A. (1957) *Models of Man, Social and Rational*. John Wiley and Sons, New York.
- Status (1994) What Antilocks Can Do, What they Cannot Do. *Status* 29 (2), January, pp. 1–5. Insurance Institute for Highway Safety, Arlington, VA.
- Stenstrom, B. (1995) What Can We Learn from the ESTONIA Accident? Some observations on technical and human shortcomings. In *The Cologne Re Marine Safety: Seminar*. Rotterdam, 27–28 April.
- Svedung, I. and Rasmussen, J. (1996) Representation of Accident Scenarios. To be published.
- Swain, A.D. (1963) A Method for Performing Human Factors Reliability Analysis. (Monograph-685): Sandia Corp., Albuquerque, NM.
- Swain, A.D. and Guttman, H.E. (1983) *Handbook on Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications*. NUREG/CR1278, USNRC.
- Taylor, D.H. (1981) The Hermeneutics of Accidents and Safety. *Ergonomics* 24 (6), 487–495. Also in: *New Technology and Human Error*, ed., J. Rasmussen, K. Duncan and J. Leplat. Wiley and Sons, New York.
- Taylor, D.H. (1987) The role of Human Action in Man Machine Systems. In *New Technology and Human Error*, ed. J. Rasmussen, K. Duncan and J. Leplat. Wiley and Sons, New York.
- Taylor, F.W. (1911) *Scientific Management*. Harper and Row, New York.
- Taylor, J.R. (1994) *Risk Analysis for Process Plant, Pipelines, and Transport*. E and FN Spon, London.
- Thompson, J.D. (1967) *Organizations in Actions*. McGraw-Hill, New York.

- Tversky, A. and Kahneman, D. (1974) Judgment under uncertainty: Heuristics and biases. *Science* **185**, 1123–1124.
- Vicente et al. (1995) A Field Study of Operator Cognitive Monitoring at Pickering Nuclear Generating Station. (Technical Report CEL 9504).: Cognitive Engineering Laboratory, University of Toronto.
- Vicente, K.J. and Rasmussen, J. (1992) Ecological Interface Design: Theoretical Foundations. *IEEE Trans. SMC* **22** (4), 589–607, July/August.
- Visser, J.P. (1991) Development of Safety management in Shell Exploration and Production. Contribution to '91 Bad Homburg Workshop on Risk Management. Published in: B. Brehmer and J.T. Reason (Eds.), *In Search of Safety*. Lawrence Earlbaum, Hove, UK.
- von Neuman, J. and Morgenstein, O. (1944) *Theory of Games and Economic Behavior* (Reissued 1980).
- Waldrop, M.M. (1987). Computers Amplify Black Monday. *Science* **238**, 602–604.
- Weber, M. (1947) *The Theory of Social and Economic Organization*. A.M. Henderson and Talcott Parsons (trans.) and Talcott Parsons (ed.). The Free Press of Glencoe, New York.
- Weick, K. (1977) Organization Design: Organizations as Self-designing Systems. *Organizational Dynamics*, Autumn, pp. 32–46.
- Wilde, G.J.S. (1976) Social Interaction Patterns in Driver Behaviour: An Introductory Review. *Human Factors* **18** (5) 477–492.
- Wilde, G.J.S. (1985) Assumptions Necessary and Unnecessary to Risk Homeostasis. *Ergonomics* **28** (11) 1531–1538.
- Wilpert, B. (Ed.) (1987-) *New Technology and Work Series*. Wiley, London.