Decision Making in Fire Risk Management

Tehler, Henrik

2001

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Decision Making in Fire Risk Management

Henrik Johansson

Department of Fire Safety Engineering
Lund University, Sweden

Brandteknik
Lunds tekniska högskola
Lunds universitet

Report 1022, Lund 2001

Decision Making in Fire Risk Management

Henrik Johansson

Lund 2001

Decision Making in Fire Risk Management

Henrik Johansson

Report 1022
ISSN: 1402-3504
ISRN: LUTVDG/TVBB--1022--SE

Number of pages: 56
Illustrations: Henrik Johansson

Keywords
Risk analysis, uncertainties, fire protection, decision analysis, Bayesian updating

Sökord
Riskanalys, osäkerheter, brandskydd, beslutsanalys, Bayesiansk uppdatering

Abstract
Various normative decision theories are discussed in the context of fire risk management. A method suitable for practical decision making in respect to fire safety investments is presented and exemplified. The method involves the use of second-order probabilities to represent uncertainty regarding probability values. A discussion on the use of Bayes theorem in combination with decision analysis is also included. The two case studies the thesis includes involve decisions regarding investments in water sprinkler systems for facilities belonging to the companies ABB and Avesta Sheffield, respectively. Calculations of the net present value of these investments are dealt with in the case studies.

Brandteknik
Lunds tekniska högskola
Lunds universitet
Box 118
221 00 Lund

brand@brand.lth.se
http://www.brand.lth.se

Telefon: 046 - 222 73 60
Telefax: 046 - 222 46 12

Department of Fire Safety Engineering
Lund University
P.O. Box 118
SE-221 00 Lund
Sweden

brand@brand.lth.se
http://www.brand.lth.se/english

Telephone: +46 46 222 73 60
Fax: +46 46 222 46 12

Summary

Investments in fire protection are characterised by their tending to not generate income but to only result in expenses, for example, those of the investment itself and of maintenance. Although an investment in fire protection might lead to the insurance premium being reduced, so that it could be seen as generating income, the question remains of how one should evaluate the reduction in fire risk which the investment involves. This is not easy to do, since both the occurrence and the spread of fire are highly uncertain, so that it is impossible to know in advance either how many fires will occur during the lifetime of the investment or, if a fire occurs, to what extent it will spread.

The attempt is made in the thesis to clarify the use of different normative decision models in the context of fire risk management. These decision models are used in order to be able to evaluate the reduction in fire risk which is achieved and at the same time to consider the more certain costs and benefits in the form particularly of the initial investment costs and the maintenance costs. Traditional Bayesian decision theory is presented briefly and its application in the area of fire risk is exemplified. Traditional Bayesian decision theory is regarded here as representing the basis for the decision of whether to invest in a fire protection system or not, but also as being in need of modification in order to be able to deal with decision problems related to fire. The reason for this is that traditional Bayesian decision theory does not allow probabilities and consequences to be expressed as being uncertain, but only as exact values. For some of the probabilities and consequences used in the quantitative analysis of the risk reduction achieved by fire safety investment, expressing them as an exact values is very difficult. The modified method employed for decision making here is referred to as the reliability-weighted expected utility (RWEU) model. This model involves a weighted average being used to represent all uncertain parameters (probabilities and consequences). The weighted averages with respect to the probability distribution describing the uncertainty of each of the uncertain parameters are used then to calculate the expected utility of the alternatives in question, the one with the highest expected utility being deemed the best in terms of this model.

For expressing the uncertainty regarding some specific probability, use is made of second-order probabilities, and for expressing the uncertainty in regard to frequencies, probability distributions representing one’s belief about what frequency values are most probable are employed. These distributions, which represent the uncertainty regarding the parameters used in the model of fire occurrence and of fire spread can be utilised then in Bayesian updating. Bayesian updating involves the subjectively estimated (prior) distributions for the frequencies and probabilities in question being updated by use of such information as statistics concerning the building at hand, for example. This provides a posterior distribution which is the result of both subjective and objective quantities or information. The thesis describes use of the Bayesian updating procedure both for reducing the uncertainty concerning the frequency of fire in the building and for updating the probabilities involved.

The thesis deals only with economic aspects of fire safety, other aspects such as those of human safety and of the flexibility of the safety system, not being dealt with explicitly. Since only economic matters are considered here, the results can be expressed in terms of assessments of the net present value of the fire protection investment considered. In calculating the net present value, use is made of the risk reduction which the investment provides in the form of the reduction in the expected costs due to fire. Two case studies, of Asea Brown Boveri (ABB) and Avesta Sheffield, respectively, are included to exemplify the use of the method suggested. Both case studies involves the calculation of the profitability of investments in water sprinkler systems for large industrial buildings. The case studies showed the net present value of investment in the sprinkler system for the ABB building to be 31.000.000 SEK and of that for the Avesta Sheffield building to be 156.000.000 SEK, which implies that the investments were profitable in both cases.

Sammanfattning (Summary in Swedish)
Att fatta beslut angående investeringar i brandskydd kan vara svårt. En anledning till detta är att den riskreducering som investeringen är tänkt att åstadkomma är svår att värdera. I denna rapport diskuteras olika beslutsteoretiska modeller med målet att kunna använda dem för att kunna ta hänsyn till riskreduceringen då man är intresserad av värdera olika investeringsalternativ angående brandskydd.

I rapporten presenteras traditionell Bayesiansk beslutsteori översiktligt, och tillämpningen av denna teori på beslutsproblem rörande brandskydd exemplifieras. Traditionell Bayesiansk beslutsteori utgör grunden för modellen som används för att avgöra huruvida det är lämpligt eller ej att göra en investering. Denna teori behöver dock modifieras för att praktiskt kunna användas i beslutsfattande angående brandskydd. Detta beror på att traditionell Bayesiansk beslutsteori förutsätter att beslutsfattaren är beredd att ange exakta sannolikheter och konsekvenser för de olika utfallen av beslutet. Detta är, i många fall, mycket svårt när det gäller brandskyddsrelaterade problem. Den modifierade modell som används i rapporten kallas *reliability-weighted expected utility*-modellen (RWEU) och innebär att sannolikhetsvärden kan anges på ett icke precist vis genom andra ordningens sannolikheter, d.v.s. flera värden betraktas som möjliga för en specifik sannolikhet. Den beslutsregel som används i RWEU-modellen är att det alternativ är bäst som har den högsta *viktade* förväntade nyttan. Beslutsregeln innebär att de sannolikheter som ingår i beräkningen av den förväntade nyttan viktas med hänsyn till de andra ordningens sannolikheter som antagits. Om man är osäker på konsekvensernas värden kan även dessa anges som flera värden och en sannolikhetsfördelning definieras för värdena i fråga.

Att icke precist uttrycka värden av parametrar (exempelvis sannolikheter) är lämpligt då man är intresserad av att uppdatera sin analys med hjälp av statistisk information från den byggnad i vilken man genomför analysen. Genom att använda Bayes sats i kombination med statistisk information angående bränder i en specifik byggnad kan den ursprungliga analysen uppdateras eller förbättras. Detta tillvägagångssätt är särskilt lämpligt att använda då man har knapphändig information angående en parameter eftersom man då kan kombinera en subjektiv bedömning från beslutsfattaren med objektiv information i statistiken. Detta sätt att uppdatera en analys kan även användas i långsiktigt riskhanteringarbete då man är intresserad av att registrera hur risken i en byggnad utvecklas över tiden. I rapporten ges exempel på hur både sannolikheter och frekvenser kan uppdateras.

Modellerna som används för att analysera en investering i rapporten beaktar endast ekonomiska aspekter av beslutet; eventuella andra aspekter som kan påverka beslutet, exempelvis säkerheten för personer i byggnaden, beaktas således inte. Detta gör det möjligt att uttrycka analysen av beslutet i form av en investeringskalkyl, vars resultat blir ett kapitalvärde för den aktuella investeringen.

För att demonstrera användningen av de modeller som presenterats i rapporten redovisas också två praktikfall som innebär att investeringskalkyler för heltäckande sprinklersystem upprättats. Investeringskalkylerna genomfördes i byggnader som tillhör ABB respektive Avesta Sheffield. Resultatet från analyserna är att kapitalvärdet för sprinklerinvesteringen i ABB-byggnaden är 31 Mkr och i Avesta Sheffield-byggnaden 156 Mkr, vilket innebär att investeringarna för respektive företag är lönsamma.

Acknowledgement

The work presented in the licentiate thesis has been supervised by Professor Sven-Erik Magnusson, who has provided valuable support and encouragement. Dr. Robert Goldsmith, Dr. Per-Erik Malmnäs and Lars Nilsson, Lic Eng, have all contributed to the thesis through their valuable advice in their areas of expertise.

During the work with the case studies included in the thesis I have received valuable help from Ingemar Grahn (Avesta Sheffield), Olle Österholm (Avesta Sheffield), Bo Sidmar (ABB) and Michael Zeeck (ABB).

I would also like to thank my friends at the Department of Fire Safety Engineering both for helping me with practical matters and for the inspiration they have given me during my writing of the thesis.

The financial support for the work on which the thesis is based that Brandforsk, The Swedish Fire Research Board, has provided is gratefully acknowledged. The work has been part of a project entitled *Ekonomisk optimering av det industriella brandskyddet* (Economic optimisation of industrial fire protection), having the project numbers 102-981 and 103-991. A reference group consisting of the following members has been connected with the project: Tommy Arvidsson (Swedish Fire Protection Association), Nils Fröman (Pharmacia & Upjohn), Ingemar Grahn (Avesta Sheffield), Anders Olsson (Zurich), Bo Sidmar (ABB), Michael Hårte (Saab Military Aircraft), Ola Åkesson (Swedish Rescue Services Agency), Sven-Erik Magnusson (Lund University), Lars Nilsson (Legal, Financial and Adminstrative Services Agency), Per Nyberg (Skandia), Björn Lindfors (SKF), Jan-Erik Johansson (Stora Risk Management), Liselotte Jonsson (Sycon) and Per-Erik Malmnäs (Stockholm University).

Contents

1. Introduction

Since the occurrence of fire is highly uncertain, one can never know how many fires will occur, if in fact any fires at all, in a given building or set of buildings during any specific period of time. What consequences a fire would have in a particular building, if it should occur, is also highly uncertain. This constitutes a problem when decisions are to be made concerning fire protection for a specific building, due to the uncertainties just described, making it extremely difficult to evaluate the decision alternatives that are available.

The present licentiate thesis attempts to clarify the use of certain normative decision theoretical models in the domain of fire risk management. A normative, as opposed to a descriptive, theory specifies how decision makers *should* make decisions rather than how they *actually make* them.

In addition to the present report, the licentiate thesis consist of two other reports (Johansson, 2000a and 2000b) written in Swedish.

1.1. Background

The management of an organisation has obligations towards various interest parties, the shareholders included, to manage effectively the risks that can threaten the organisation's goals. An initial step in achieving this is to assess the risks[1], determining whether the present risk is significant or not. If one determines that a significant risk is present, one has to decide how it can best be reduced. The complexity of the decision to be made can vary considerably, depending on the situation. Some decisions concerning risk are easy ones, due to relatively limited costs associated with reduction of the risk to an acceptable level. In order to make a decision of this kind extensive analyses of the problem are seldom needed. On the other hand, when decisions involve large uncertainties and there are considerable costs associated with the different risk-reducing alternatives, a thorough and time-consuming examination of the problem is often necessary in order that as satisfactory a decision as possible can be made.

The uncertainties that exist when one endeavours to model the occurrence and the spread of fire make it difficult to describe in precise terms the benefits in the form of increased safety (reduction of risk) that one receives when the fire protection in a building is improved. This, in turn, creates difficulties when different fire protection alternatives are to be evaluated so that a decision regarding them can be made. To find a solution to this problem, one has to create a model that can be used to assess the benefits one receives by choosing one fire protection alternative rather than another.

In such a model it is necessary to determine what goals the decision maker has in making the decision. For decision making with respect to fire protection, two frequent goals are economic ones and those of human safety. An economic goal, for example, could be that the sum of the fire protection investments not exceed a particular amount. A safety goal, in turn, could be that no individual be exposed to physical danger due to fire. A fire protection alternative that meets the demand of building regulations is commonly judged to fulfil the goals set for human safety, although it can well be the case that greater human safety than that which the building codes require is sought after. Evaluation of the amount of additional safety to be aimed at is not easy, however, since this can require that one assess the value of human life. Such evaluation is difficult to make and in the present thesis no attempt will be made to

[1] Throughout the thesis, if not noted otherwise, risk will be defined in accordance with Kaplan & Garrick (1981).

evaluate it. The evaluation of different fire protection alternatives will instead be based entirely on economic considerations.

Even if one restricts the evaluation of alternatives to consideration of economic goals alone, one has to deal nevertheless with the problem of how an increase in fire safety can be evaluated. Earlier investigations of economic aspects of fire protection suggested that general loss data from insurance companies be used for estimating the expected annual costs due to fire in a specific building (Ramachandran, 1998 and Shpilberg & Neufville, 1974). This approach implies, however, that any *specific* building can likewise be represented by information concerning the *general* case. If one wants a more accurate description of the expected annual costs in the building being analysed attributable to fire one needs to investigate the costs connected with the various fire scenarios that are possible in the building in question, which is also pointed out in Shpilberg & Neufville (1974). This can be problematic, however, since the often very limited amounts of information available make it difficult to estimate the probabilities of different events that could occur during a fire. Another difficulty in the evaluation of possible consequences of a fire in a particular building is that it is not certain that the direct and consequential (indirect) losses traditionally reported in the statistics of insurance companies (presented in for example Räddningsverket, 1999) are the losses that should be addressed in an analysis of a fire protection system from the building owner's perspective. The losses that the building owner wishes to evaluate are the losses he or she may need to defray. These are not the costs reported in the statistics of the insurance companies, which concern what the fires have cost the *insurance companies*. Those costs of a fire that the decision maker (in this case the building owner) needs to defray are called *uncompensated losses*. These could be such matters as deductibles, fines, costs of additional marketing campaigns, costs of postponed investments, and the like. Thus, some of these losses can be very hard to quantify. The uncompensated losses are dependent upon the type of building and type of firm being analysed. In the present thesis the question of what types of losses can be regarded as uncompensated losses will not be explored further. It will only be concluded that there can be other kinds of losses than direct and consequential losses and that in decision theoretic terms the uncompensated losses are the correct ones for a company to use in evaluating different fire protection measures.

As was indicated, there is often only a very limited amount of information available concerning the probabilities of occurrence of different events if a fire develops in a *specific* building. This can compel a person who performs a risk analysis or a decision analysis to use information from other sources than that of the actual building, such as expert judgements, general statistical information, and the like. The problem of how to combine information from different sources in order to make estimates pertaining to the building of interest arises then. Apostolakis (1988) and Kaplan & Garrick (1979) discussed this earlier. In the present thesis, the focus is on providing an overview of the methods involved in the use of information from sources other than that of the building in question, a number of additional examples also being provided.

The same method as that utilised in combining information from different sources can also be used in the continual updating and improvement of a risk analysis (which could serve as a basis for decision making) through use of monthly or yearly information about fires in a building or lack of fires. Although information on the occurrence and spread of fire in a specific building is often not sufficient to provide reliable estimates of the probability of different events that can occur during a fire, such information can be used to improve an analysis performed earlier, something which from a management perspective is very useful.

Thus by collecting statistics from a specific building and incorporating it into the existing analysis by use of the same approach as taken when information from different sources is combined so as to always have an up-to-date analysis of the fire risks in the building. It is also possible to monitor changes in the level of risk in a building, hopefully taking account of a possible increase in the risk of fire before a serious fire occurs.

Still a further important aspect of decision making here is how to present the basis for a decision to the decision maker. In the present context, where decisions regarding fire risk are involved, determining the basis for decisions can be complicated since not all decision makers are knowledgeable in the area of fire protection. In such a case it can be more favourable to present the basis for a decision in the form of an investment appraisal rather than of a risk analysis of the different alternatives. This can help the decision maker translate the increase in safety that an investment provides into monetary terms and can thus enable alternatives to be compared on a more rational economic basis.

1.2. Objective and purpose

The aim of the present thesis is to suggest a method for making practical use of investment appraisal as a decision aid in the area of fire-risk management. The thesis also aims at clarifying the manner in which statistical information can be combined with subjective estimates when management decisions concerning fire risk are made. To demonstrate use of the method suggested, two case studies are presented. These were carried out in two industrial facilities, the one belonging to Asea Brown Boveri (ABB) and the other to Avesta Sheffield.

1.3. Overview of the thesis

After this brief introduction, there follows a chapter introducing decision theory within the context of fire risk management. In that chapter, classical Bayesian decision theory is presented, along with criticisms of it and some of the newer models for decision-making that have been proposed. The focus is on models that can be of practical use in fire risk management when knowledge of the probabilities and consequences associated with fire are limited.

In chapter three the problem of how subjective estimates can be combined with objective statistics or measurements is discussed. This is in fact the same problem as that regarding how one should evaluate a previously performed risk analysis in the light of new information. The method employed is called Bayesian updating. The chapter contains an introduction to Bayes theorem (which is the basis for the updating process) as well as various examples illustrating the use of the theorem. In the same chapter, a survey of fire statistics from different industrial categories in Sweden is also presented. The survey is accompanied by estimates of the probabilities associated with different degrees to which a fire can be expected to spread. This information is intended to be used in performing a decision analysis pertaining to fire protection found in a specific building.

The fourth and fifth chapters are concerned with the more practical use of decision analysis in the context of fire risk management. Investment appraisal as it applies to safety systems, such as those for fire protection, is discussed, examples being given. Two case studies in which investment appraisals were performed at industrial facilities are also presented in chapter five.

The last chapter, finally, contains a discussion of the practical usefulness of the material which is presented in the thesis and of the conclusions that are drawn.

2. Normative decision analysis

In this chapter, traditional Bayesian decision theory will be presented briefly within the present context of fire risk management. Traditional Bayesian decision theory is a normative decision theory meaning that it *prescribes* how a decision maker *should* reach a decision. This is to be distinguished from a descriptive theory which *describes* how a person *actually* makes a decision. The thesis deals only with theories of the first type, i.e. with normative theories.

The chapter starts with a presentation of traditional Bayesian decision theory as described in Gärdenfors and Sahlin (1988). Following this, certain arguments that have been directed against this theory are presented. The interpretation of probability in risk analysis and decision analysis is also discussed in this chapter.

Since the problem faced in trying to decide between different fire protection alternatives can be quite complex, some of the estimates of probabilities (and of consequences) that one needs to make can be difficult to performe. It can be difficult, for example, to find a *unique* number to represent the probability of the event that the sprinkler will extinguish the fire. This is a not at all uncommon problem in fire risk analysis since one often has only very limited information, or no information at all, concerning a given probability in a specific building. This problem is discussed in the light of certain decision theories that suggest a solution to the problem of how to evaluate alternatives when large uncertainties are present.

2.1. Bayesian decision theory

We need to make decisions many times each day. Often, we are not even aware of being in a decision situation, but simply do what comes naturally. In opening a closed door, for example, one does not usually think of the decision situation of needing to choose between opening the door quickly or opening it slowly; one simply opens the door, often without paying any attention at all to the speed with which it was done. There can be other decision situations, however, where one needs to think a bit harder about the action to be taken, for example, whereas one should take the job one has just been offered or, perhaps more relevant in the present context, whether one should invest in a new sprinkler system. Decisions of this type usually require a more thorough analysis than when simply choosing the alternative that comes naturally.

The reasons for some decisions being harder to make than others come from four different sources (Clemen, 1996). First, a decision may be difficult to make because of its *complexity*. The complexity can be due, for example, to there being several different issues one has to deal with in making the decision. The decision of where a new airport is to be built, for example, may be influenced by the travel time from the nearest city, the level of noise that people nearby are exposed to, the construction costs, etc. As the number of factors affecting a decision grows, the decision maker has increasing difficulties in keeping the different issues in mind. A decision analysis can help the decision maker to structure the problem and to keep track of the different issues that affect the decision.

Second, a decision may be difficult to make because of the *uncertainties* associated with it. For example, the uncertainty of whether there will be a large fire in a specific building during the next 30 years could make the decision of whether to install a fire protection system there difficult. In another, more extreme case, the uncertainties regarding the consequences of different alternatives for the storage of spent nuclear fuel from now until the year 2080 can make the decision of which type of storage to choose difficult. A thorough decision analysis can help the decision maker deal with the uncertainties involved by finding important sources of uncertainty and letting him/her quantify these in order to assess their influence on the decision.

Third, a decision may be difficult because of the decision maker having differing and partly opposing objectives to be met. This can force him/her into making a *trade-off* between the benefits in one area and costs in another, for example. A trade-off in the area of fire protection could be between the cost of a particular fire protection system and the benefits of the increased safety the system provides. There can be other trade-offs as well, such as between the attributes of flexibility and of cost. Trade-offs of this sort compel the decision maker to make judgements concerning how important the different attributes involved are to him/her. These are judgements which are not always easy to make. Decision analysis gives the decision maker a tool for dealing with the trade-offs here in a quantitative way.

Fourth, a decision may be difficult because of *different perspectives* on the decision problem leading to differing conclusions. This is especially true when the decision maker is not a single person but a group of persons. There may be various persons in the group who differ in their preferences regarding the possible consequences. This can result in a complicated situation. Even if the decision maker is alone in his or her decision, differing perspectives can constitute a problem, since it is not always easy to know one's own preferences exactly. The use of decision analysis can help the decision maker to produce an accurate description of his/her preferences.

In conclusion, there can be different sources of difficulties in making a decision, but a decision analysis of appropriate comprehensiveness can help the decision maker sort things out and hopefully decide which alternative is best.

In terms of traditional Bayesian decision theory, a decision maker has different *alternatives* to choose from, the alternatives differing in the *consequences* they have for the decision maker, depending on which of several possible *states of the world* occurs. Thus, the uncertainty which the problem contains is represented by the different possible states of the world, to each of which the decision maker has to assign a probability. In traditional Bayesian decision theory, it is assumed that the decision maker can represent his/her belief regarding the different possible states of the world by a unique probability distribution. This means that each uncertain state that can occur and that affects the outcome of the decision must be assigned a specific probability value such that the sum of all the probability values involved is equal to 1.

The decision problem can be described by use of a decision matrix. An example of a decision matrix, expressed in general terms, is shown in Figure 1, where $s_1, s_2,\ldots,s_m$ are different states of the world, $a_1, a_2,\ldots,a_n$ are the different alternatives and $o_{n,m}$ are the consequences that occur if alternative n is chosen and state m is the one that occurs.

Alternative	State of the world s_1	s_1	.	.	s_m
a_1	$o_{1,1}$	$o_{1,2}$			$o_{1,m}$
a_2	$o_{2,1}$	$o_{2,2}$			$o_{2,m}$
.	.	.			.
.	.	.			.
a_n	$o_{n,1}$	$o_{n,2}$			$o_{n,m}$

Figure 1 General model of a decision matrix.

To exemplify the use of a decision matrix, an example from Savage (1954) will be employed:

> "Your wife has just broken five good eggs into a bowl when you come in and volunteer to finish making the omelet. A sixth egg, which for some reason must either be used for the omelet or wasted altogether, lies unbroken beside the bowl. You must decide what to do with this unbroken egg. Perhaps it is not too great an oversimplification to say that you must decide among three acts only, namely, to break it into the bowl containing the other five, to break it into a saucer for inspection, or to throw it away without inspection. Depending on the state of the egg, each of these three acts will have some consequence of concern to you" (Savage, 1954)

The problem posed by Savage can be represented by a decision matrix, one which is shown in Figure 2 (Gärdenfors and Sahlin, 1988).

	State	
Act	Good	Rotten
Break into bowl	six-egg omelet	no omelet, and five good eggs destroyed
Break into saucer	six-egg omelet and a saucer to wash	five-egg omelet and a saucer to wash
Throw away	five-egg omelet and a good egg destroyed	five-egg omelet

Figure 2 Decision matrix of Savage's omelet problem (Gärdenfors and Sahlin, 1988).

Looking at the matrix in Figure 2, one can see the three possible acts: break the egg into the bowl; break it into the saucer; throw it away. Depending on the state of the egg, i.e. whether it is rotten or not, and the act chosen, some one of the following consequences will occur: six-egg omelet; no omelet and five good eggs destroyed; six-egg omelet and a saucer to wash; five-egg omelet and a saucer to wash; five-egg omelet and a good egg destroyed; five-egg omelet.

2.1.1. The principle of maximising expected utility

This is a general way of presenting the decision problem, but it provides no advice on which alternative should be chosen. According to traditional Bayesian decision theory, in order to establish which alternative is best, the decision maker should follow a set of axioms[2] in his/her decision making. By following these axioms, the decision maker *will then act in accordance* with the principle of maximising expected utility (PMEU) (Gärdenfors and Sahlin, 1988). PMEU implies that the decision maker should calculate the expected utility of each of the available decision alternatives and choose the alternative that has the highest

[2] There are several axiomatic systems of this sort, that formulated by Ramsey (1931) and that formulated by Savage (1954) being two of the more famous ones (Gärdenfors and Sahlin, 1988).

expected utility value. Thus, in traditional Bayesian decision theory the maximising of expected utility is treated as the result of the decision maker having followed this set of axioms. Malmnäs (1994) has shown, however, that the principle of maximising expected utility does not follow from the axioms[3]. He also concludes that the chances of supporting the principle by a formal justification in terms of an axiomatic system are very slight. It thus appears that the logical foundations of PMEU are weak. As will be explained below, this does not mean, however, that PMEU is a poor decision criterion for use in fire protection.

Support for the adequacy of approach, such as PMEU as a *choice rule generator*, i.e. as a way of identifying the optimal choice between alternatives (e.g. between different lotteries), can come either from above or from below. Support from above means showing that the approach in question yields a result which is the only one which satisfies certain desirable properties, e.g. those described by some axiomatic system. Support from below, in contrast, means that one can show that the solution arrived at does not entail counterintuitive choices to any appreciable extent. As just indicated, there appears to be little chance of supporting PMEU from above. In contrast, as regards support from below, Malmnäs (1999) has concluded that PMEU is a better choice rule generator than any simpler choice rule generators such as the Minimax or Maximax rule, and that no other choice rule generator appears to be better than PMEU. There appears to be little reason, therefore, to turn to any other approach for evaluating of alternatives in fire protection engineering, although it could be desirable to improve PMEU in order to be able to deal with decision situations involving imprecise probabilities (see section 2.3).

2.1.2. Example of an analysis using traditional Bayesian decision theory

The part here dealing with traditional Bayesian decision theory will be concluded by an analysis of Savage's omelet problem to find the alternative with the highest expected utility. One has first to assign a utility value to each of the two possible consequences for each of the three decision alternatives available. The possible consequences that the omelet problem involves and the corresponding utility values are shown in Table 1. On the basis of the information there, it can be concluded that the consequence "Six-egg omelet" is best, followed by "Six-egg omelet and a saucer to wash" and the others in the order shown.

Table 1 Illustration of the consequences and the corresponding utility values in Savage's omelet problem.

Consequence ($o_{i,j}$)	**Description**	**Utility ($u_{i,j}$)**
$o_{1,1}$	Six-egg omelet	1
$o_{1,2}$	No omelet and five good eggs destroyed	0
$o_{2,1}$	Six-egg omelet and a saucer to wash	0,95
$o_{2,2}$	Five-egg omelet and a saucer to wash	0,65
$o_{3,1}$	Five-egg omelet and a good egg destroyed	0,6
$o_{3,2}$	Five-egg omelet	0,7

Note, regarding the consequences, that for the i-values, 1=break the egg into the bowl, 2=break the egg into the saucer, and 3=throw the egg away, and that for the j-values, 1=egg is good, and 2=egg is rotten.

It can also be seen that the utility values assigned to the different consequences range from 1 to 0. The scale can be chosen arbitrarily; what is of interest is the ratio between the different utility numbers. A scale ranging from 1 to 0 is often used since it is easy to work with and, as

[3] Malmnäs (1994) has examined the axiomatic systems suggested by Herstein and Milnor (1953), by Savage (1954) and by Oddie and Milne (1990).

will be shown later in this chapter, offers certain mathematical simplifications when imprecise probabilities are being dealt with.

To calculate the expected utility for each of the three alternatives, one needs to estimate the probability for each of the two possible states of the egg, i.e. the probability that the egg will be good and the probability that it will be rotten, where these two probabilities complement each other (add up to 1). In terms of traditional Bayesian decision theory, estimating the one probability or the other is done subjectively by the decision maker and thus represents his/her *belief* in the event in question. This constitutes a significant difference as compared with the more common frequentistic interpretation of probability, according to which a probability of a given event can be defined as the limiting value of the ratio of the number of successful trials (trials in which it occurred) to the total number of trials (trials in which it could either occur or not occur).

Although a more thorough discussion of a Bayesian interpretation of probability will be presented shortly, suppose in the example considered that the decision maker's belief in the egg being rotten can be represented by the numerical value of 0,2. This allows one to calculate the expected utility for each of the three decision alternatives listed in Figure 2. The first alternative (a_1) was to break the egg into the bowl, putting it in contact with the other five eggs; the second alternative (a_2) was to break the egg into a saucer for inspection; the third alternative (a_3) was to throw the egg away without inspection. In the calculations shown below of the expected utility of each of these three alternatives – $E(U_1)$ etc. – $p_S(s_1)$ is the probability (subjective probability) that the egg will be in state 1, that of its being good, and $p_S(s_2)$ is the probability of its being in state 2, that of its being rotten. In the designations of the utilities with which these probabilities are linked – $u_{1,1}$, etc. – the first subscript refers to the alternative and the second subscript to the state of the egg. The calculations are as follows:

$$E(U_1) = p_S(s_1) \cdot u_{1,1} + p_S(s_2) \cdot u_{1,2} = 0,8 \cdot 1 + 0,2 \cdot 0 = 0,8$$
$$E(U_2) = p_S(s_1) \cdot u_{2,1} + p_S(s_2) \cdot u_{2,2} = 0,8 \cdot 0,95 + 0,2 \cdot 0,65 = 0,89$$
$$E(U_3) = p_S(s_1) \cdot u_{3,1} + p_S(s_2) \cdot u_{3,2} = 0,8 \cdot 0,6 + 0,2 \cdot 0,7 = 0,62$$

It follows from the expected utility calculated for the different alternatives that alternative 2, first breaking the egg into a saucer for inspection before putting it into the bowl with the rest of the eggs, has the highest expected utility and should thus be chosen. This illustrates certain basic principles of how a decision analysis can be conducted using traditional Bayesian decision theory.

2.2. Subjective probability

There are different ways in which probability can be understood, depending on what type of decision theory one employs. In traditional Bayesian decision theory, probability is perceived as a subjective probability (also called a personal probability), one that can be uniquely determined through betting rates. A subjective probability is thus a reflection of a decision maker's belief concerning a particular event. It could for example be the probability represented by the statement, "The odds are 3 to 1 that it will rain tomorrow".

A decision maker's subjective probability can be determined by letting him/her choose between fictitious lotteries in which the event that occurs determines the outcome. Assume, for example that you are asked to assign a probability to the event that it will rain in London day after tomorrow. This probability can be derived by letting you choose between two alternatives. The first alternative could be that you are going to draw a ball from an urn with 50 blue and 50 red balls. If you draw a red ball you receive 100 SEK[4] and if you draw a blue you receive nothing. The second alternative is that you receive 100 SEK if it rains in London day after tomorrow and nothing if it does not. Which alternative should you choose? If you choose alternative 1, your subjective probability that it will rain in London day after tomorrow is less than 0,5. If you choose alternative 2, on the other hand, your probability of the event is higher than 0,5. Finally if you are indifferent between the two alternatives, your subjective probability is 0,5. This is a first step. To go on then, if it is alternative 1, for example that you chose, a new decision situation can be created, such as having an urn with 45 red balls and 55 blue and again letting you decide whether you prefer alternative 1 or 2. Decision situations of this sort can be continued until an urn is found with a proportion of red and blue balls such that you are indifferent between the two alternatives. When this point has been reached, your subjective probability for the event can be derived by knowing the proportions of red and blue balls in the urn.

Such an approach can be used to reveal a person's subjective probability regarding a particular event. What if the person is indifferent between the two alternatives, however, not only when the urn contains 30 red balls, but also when it contains 35 and when it contains 40 red balls? The person might state that although he is *definitely* indifferent between the alternatives when there are 35 red balls in the urn, he also feels indifferent between them, both when there are 30 and when there are 40 red balls in the urn. Ambiguity of this sort is not accepted in traditional Bayesian decision theory. There, a decision maker cannot assign more than one probability to a given event but he must assign a *specific* numerical probability value to it.

In problems involving decision-making in a fire protection context, one is frequently forced to assign probabilities to various rather uncertain events contained in a fire scenario. This is often difficult to do since it is not uncommon for there to be very little information concerning the probabilities involved. Thus, it may be more helpful here to assign a *set of plausible probability values* than to have to settle for a precise value. It is important to recognise that the set of values finally selected are not objective quantities despite there being a set of values rather than a single value. If the person serving as decision maker should change, the probability values regarded as plausible might change as well. Nevertheless, if one assign a set of probability values rather than a single value, it is more likely that different persons can agree to their being reasonable than if only a single value is employed. In the context of the thesis, the subjectivity of probabilities need not pose a serious problem, since the aim is to provide a recommendation to the decision maker, which is often a company. In order to prepare the recommendation, one needs to use the "company's" subjective probabilities, which could be interpreted as those of the persons responsible for decisions there. Since these persons are very likely lacking in knowledge of fire protection and may thus not be considered able to provide meaningful estimations themselves, they are likely to have to rely instead on an expert or a group of experts to provide the estimates. Although they should provide whatever reasonable estimates as they can, they may very well have to simply declare, "We believe in this expert (or group of experts) and accept his/her (their) estimates as our own".

In the present context, the fact that traditional Bayesian decision theory is unable to deal with the ambiguity of various of the probabilities of interest could mean its being of limited

[4] SEK is an abbreviation for Swedish crowns (kronor).

applicability. There are also theories, however, that are able to deal with imprecise probabilities and some of them, which can be combined with a Bayesian approach, can be seen as useful in the present context. These will be taken up next.

2.3. Imprecise probabilities

As already indicated, it can sometimes be very difficult to specify a precise probability for an event. What is the probability, for example, that there will be a bus strike in Verona, Italy, next month (the example comes from Goldsmith and Sahlin, 1983)? You would probably consider it impossible to assign an exact probability value to this event. Also, if someone asked you for your preferences concerning a bet (see above) involving the occurrence or non-occurrence of this event, you would probably end up with a relatively large probability interval for the event occurring in which you would be indifferent between the two alternatives. Thus, you would not be able to assign a specific probability value to the event and you have to consider the probability in question as being imprecise.

Imprecise probabilities are common in decision situations concerned with fire protection. This is due to the uncertainties relating to the occurrence of a fire and how it would then develop. What is the probability, for example, that the staff extinguishing a fire in the storage area of a particular factory if such a fire in fact occurs? Probabilities of this kind are difficult to estimate since very little information about the parameters is available and since it is difficult to create any useful general model to use in estimating the probability in question.

Since in traditional Bayesian decision theory one cannot regard probabilities (or the utilities of possible consequences) as being imprecise, one needs to modify the theory so as to be able to deal with the kinds of uncertain estimates needed in decision making regarding fire protection. In the thesis two models that can handle imprecise probabilities and inprecise consequences are presented. The first model, called the reliability-weighted expected utility model (RWEU), represents a slight modification of traditional Bayesian decision theory. This model is used in the case studies included here. The other model, called Supersoft Decision Theory (SSD), is included because it can be used to deal with problems involving access to very little information. The use of only very limited information means that the calculations that need to be performed are more complex and more difficult to carry out. Nevertheless, through the use of computers, SSD can become a very useful tool in decision analyses in which information about the problem is very limited.

2.3.1. Reliability-weighted expected utility

A practical and easy way of dealing with the problem of handling imprecise probabilities is to perform a traditional Bayesian decision analysis but, instead of assigning specific probability measures to events assigning a set of probability measures to a given event. This means replacing a single probability measure by a *set* of "plausible" probability measures, which in turn results in a *set* of expected utility measures. To each plausible probability value, a reliability value is assigned. Although the reliability value is a second-order probability[5], the term reliability is used in order to avoid dealing with the rather complex concept of "the probability of a probability". The decision criterion used here involves choosing the alternative with the highest reliability-weighted expected utility (RWEU) (see Hansson, 1991, for example). This in turn means that, for each probability in the model describing possible fire scenarios that are deemed uncertain, one specifies a second-order probability distribution

[5] It has been argued that second-order probabilities are not needed to express the uncertainty concerning a probability value (see Savage, 1954, for example). In Appendix G this issue is discussed with respect to the use of second-order probabilities in decision making concerning fire safety, using an example from Pearl (1988).

extending over the range of probability values that are considered plausible. This second-order probability distribution could look like the one shown in Figure 3, for example. In that figure, *P(Ext)* is the probability that a particular fire protection system will extinguish a fire. As can be seen, the person who has created the probability distribution has judged the value of 0,5 to be the most reliable and the values of 0,4 and 0,6 to be less reliable but still plausible. The first-order probability here thus describes how likely a fire protection system is to extinguish a fire and the second-order probability how likely it is that a given value of the first-order probability is certain.

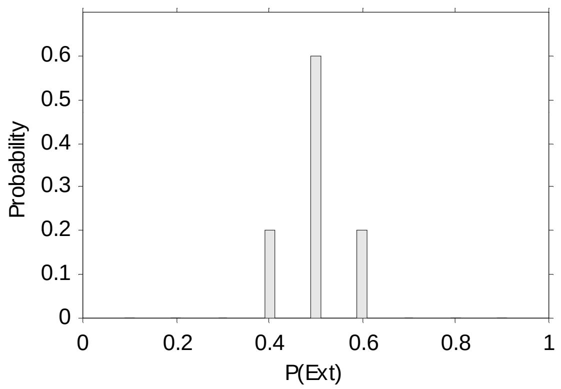

Figure 3 Illustration of a second order probability distribution.

In order to find the RWEU, one has to calculate *a single* value for the first-order probability of each event of interest. The value for this first-order probability is found by weighting each possible first-order probability value by its reliability (its second-order probability) and taking the weighted average of these. Thus, the probability value representing *P(Ext)* would be 0,5 since 0,2*0,4 + 0,6*0,5 + 0,2*0,6 = 0,5. Such weighted values are used then to calculate the expected value of the decision alternatives, the alternative with the highest RWEU-value being the optimum alternative according to the model.

A danger with use of this decision criterion is that it could result in decision situations appearing to be more clear than they actually are. Assume, for example, that you have calculated the RWEU for two alternatives and that, in terms of absolute values, the difference between the two expected utilities is insignificant. Although you would probably say that the two alternatives are equally good or that you cannot decide between them, according to the RWEU-criterion one of the alternatives is the best, since the alternatives differ in their RWEU-values. This problem can be avoided by calling a decision *robust* if the decision recommended by the RWEU-criterion has the highest expected utility *for most* of the combinations of plausible probability values weighted in terms of their reliability values. No clear definition of what “for most” means can be given, however. Rather, it is up to the individual decision maker to decide what is meant by “robust”. In the present thesis, however, a decision will be treated as robust if the decision recommended according to the RWEU-criterion is the one with the highest expected utility in more than 95% of the reliability weighted combinations of probabilities and consequences.

In the present context, reliability-weighted expected utility analysis could be employed in the following way: first, create a model in accordance with traditional Bayesian decision theory and assess all the probabilities and utilities exactly. Then assess the class of probability measures you consider appropriate for each of the probabilities you are uncertain about. In practice, you cannot assign a class of such probability measures to all of the (first-order) probabilities involved since the problems with which fire protection is concerned are often so complex that it is impractical to do so. Rather, one needs to identify the probabilities which affect the results of the analysis most and assess classes of probability measures for these probabilities only. One way of doing this is to estimate a lower and an upper limit enclosing each of the plausible values for the probabilities in question and to then note how the result is affected when the probabilities are adjusted, one at a time, each from its lowest to its highest value. Classes of probability measures should then be assessed for those probabilities that affect the result the most.

Since in practise an analysis of this sort would probably be done in the form of a Monte Carlo-simulation, assessing a distribution of a probability could be done by specifying a (second-order) probability distribution corresponding to the initial single probability measure. The result of the Monte Carlo-simulation could be displayed in the form of a histogram showing what expected utility values are most probable. This method is employed in the thesis in connection with analyses in practical terms carried out at the firms ABB and Avesta Sheffield. By studying the distribution of the differences in expected utility between the alternatives, one can determine how robust the decision in question is.

In both the analysis carried out at ABB and that carried out at Avesta Sheffield, the decision recommended by the reliability-weighted expected utility criterion was found to be robust.

2.3.2. Supersoft Decision Theory

Another way of dealing with imprecise probabilities (and also imprecise consequences) is to use a method that can handle problems in which vague statements concerning the probabilities and consequences are allowed. Two such methods are the Delta-method (Danielsson, 1995) and Supersoft Decision Theory (Malmnäs, 1995), which allow the decision maker to use vague assessments of the different probabilities and the values of the different outcomes. Such vague assessments might be, for example, "The probability must be between 0,2 and 0,8" or "The consequence c_1 is at least twice as good as consequence c_2". These vague expressions are interpreted as inequalities, which for the probability just mentioned could be in the form of $0{,}2 < p < 0{,}8$. In this thesis, only Supersoft Decision Theory (SSD) will be dealt with.

The first thing one needs to do in evaluating a decision situation in terms of SSD is to create the representation of it in a *decision frame*. This representation consists of the following: the different alternatives that can be chosen ($a_1,\ldots, a_n$), for each a_i a description of the possible consequences C_i, a list L_1 of conditional probability statements, and a list L_2 containing utility statements concerning the consequences.

To create a representation of the decision frame, pairwise disjoint trees ($T_1,\ldots, T_n$) are created such that the events contained in L_1 and L_2 are associated with disjunctions of elements in the trees.

Assume you have to make a decision of whether or not you should install a sprinkler system in a factory. The decision frame in this case could consist of the two alternatives: that you do not install a sprinkler system (a_1) and that you install one (a_2). The descriptions of the respective consequences could in the case referred to have been those presented in Figure 4. The time limit of 40 years was selected as representing the economical lifetime of the sprinkler system.

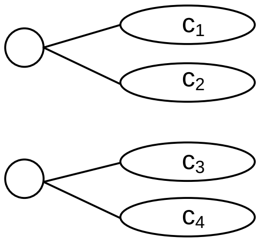

c_1 = No major fire will occur within the next 40 years (no sprinkler system)

c_2 = A major fire will occur within the next 40 years (no sprinkler system)

c_3 = No major fire will occur within the next 40 years (sprinkler system installed)

c_4 = A major fire will occur within the next 40 years (sprinkler system installed)

Figure 4 Illustration of the possible consequences in the form of two trees.

In this example, the lists L_1 and L_2 could look as they do in Table 2 and Table 3.

Table 2 Conditional probability statements.

$\mathbf{L_1}$
The probability of a major fire occurring during the next 40 years, if we *do not* install a sprinkler system, is between 0,01 and 0,015 ($p_{1,Fire}$)
The probability of a major fire occurring during the next 40 years if we install a sprinkler system is between 0,01 and 0,05 ($p_{2,Fire}$)

Table 3 Utility statements.

$\mathbf{L_2}$
The consequence c_2 is better than c_4
The consequence c_3 is better than c_2
The consequence c_1 is better than c_3
The value distance between c_1 and c_3 is equal to the distance between c_2 and c_4
The value distance between c_3 and c_2 is more than 20 times larger than the distance between c_1 and c_3

The two trees T_1 and T_2 in this case can look as they do in Figure 5.

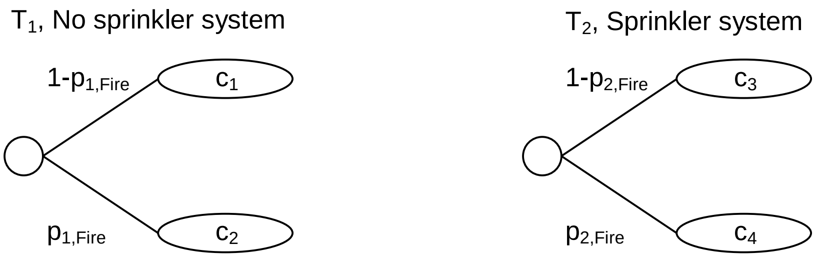

Figure 5 Illustration of the trees.

The next step is to represent the statements in L_1 and L_2 in the decision frame by the inequalities *S(p)* for representation of the probability statements and *U(v)* for the utility statements. In producing this numerical representation of the statements, one needs to check that a solution to *S(p)* and *U(v)* exists, i.e. that there are a combination of values such that all inequalities in *S(p)* and *U(v)* are satisfied.

In the above example, *S(p)* and *U(v)* could be as they are in Table 4 and Table 5, respectively.

Table 4 Representation of the probability statements.

S(p)
$0{,}01 < p_{1,Fire} < 0{,}15$
$0{,}01 < p_{2,Fire} < 0{,}05$

Table 5 Representation of the utility statements.

U(v)
$u(c_2) > u(c_4)$
$u(c_3) > u(c_2)$
$u(c_1) > u(c_3)$
$u(c_1) - u(c_3) = u(c_2) - u(c_4)$
$20(u(c_1) - u(c_3)) < u(c_3) - u(c_2)$

Using the representations of the probabilities and the utilities as shown in Table 4, Table 5 and the trees in Figure 5, it is possible to create a probability/value part of the decision frame, as shown in Table 6.

Table 6 The probability/value part of the decision frame.

B(p)	B(v)
$p_{1,Fire}$	$u(c_1)$
$1\text{-}p_{1,Fire}$	$u(c_2)$
$p_{2,Fire}$	$u(c_3)$
$1\text{-}p_{2,Fire}$	$u(c_4)$

The evaluation of the alternatives is based then on expected utility. In using the decision frame $F = (a_1,\dots,a_n, T_1,\dots T_n, B(p), B(v))$, three criteria for the evaluation of the alternatives are employed: *Min(E(U))*, *Max(E(U))*, *Mean(E(U))*. The first term refers to the lowest value for the expected utility, the second term to the highest value for it and the last term to the mean value for the expected utility.

In the evaluation process, one starts with the original decision frame *F* and examines the difference between the alternatives according to the criteria just mentioned (*Min*, *Max* and *Mean*). If alternative 1 is compared with alternative 2, for example, the result will be expressions consisting of elements of *B(p)* and *B(v)* representing $Min(E(U_{alt1}))\text{-}Min(E(U_{alt2}))$, $Max(E(U_{alt1}))\text{-}Max(E(U_{alt2}))$ and $Mean(E(U_{Alt1}))\text{-}Mean(E(U_{Alt2}))$.

If all the criteria indicate one and the same alternative to be the best, that is the alternative that is best according to SSD. It is possible, however, that the different criteria do not all point to the same alternative, the decision frame needs, in that case, to be contracted. A decision frame could be contracted by using one of the following three contraction procedures: contraction of the probability part, contraction of the value part, and contraction of both the value and probability part. An example of a contraction of the value part will be shown in the following example.

To simplify the calculations necessary in SSD one should always try to minimise the number of variables used. For example, instead of using the utilities $u(c_1)$, $u(c_2)$, $u(c_3)$ and $u(c_4)$, in the example just discussed, one could express them in two new variables *x* and *y* and thus reducing the number of variables from four to two. Assume that the worst consequence, that of the building being equipped with a sprinkler system but nevertheless there occurs a major fire, represents a utility value of 0 ($u(c_4) = 0$). If the difference, in utility value, between that consequence and the consequence of the building not being equipped with a sprinkler system and there occurs a major fire (consequence 2) is termed x, then the absolute utility value of the consequence just mentioned is x ($u(c_2) = x$). Furthermore, assume that the difference in utility value between consequence 2 and the consequence of the building being equipped with a sprinkler system and that it does not occur a major fire (consequence 3) is termed y, then the utility value of consequence 3 is equal to $x + y$ ($u(c_3) = x + y$). As a consequence of the utility statement $u(c_1) - u(c_3) = u(c_2) - u(c_4)$ shown in Table 5 the utility value of consequence 1 is $2x + y$ ($u(c_1) = 2x + y$). It is convenient to have a utility scale between 0 and 1 and since the worst consequence (c_4) has the utility value 0 then the best consequence (c_1) should have the utility value 1. Thus, $u(c_1) = 2x + y = 1$. The utilities associated with the different outcomes are shown in Figure 6.

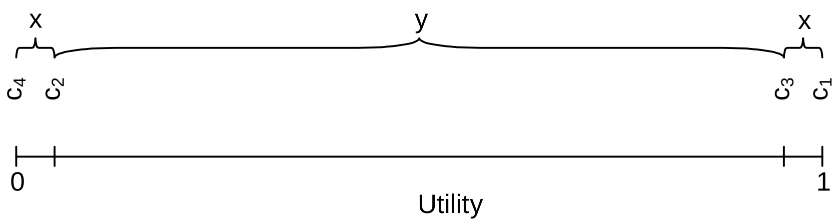

Figure 6 Illustration of the utilities associated with the different consequences.

Using this new way of expressing the problem and starting by evaluating only the probability part of the decision frame, one gets (see appendix A):

Max(E(Alt1)) - Max(E(Alt2)) = 0,95 – 0,94x – 0,99y [2.1]

Min(E(Alt1)) - Min(E(Alt2)) = 0,85 – 0,8x – 0,95y [2.2]

Mean(E(Alt1)) - Mean(E(Alt2)) = 0,9 – 0,87x – 0,97y [2.3]

Since the difference between the two alternatives according to the three criteria is thus determined by how one evaluates the different consequences and one assumes the best consequence to have the utility value of 1 ($2x + y = 1$), it is possible to calculate the value of y required in order for the alternative 2 to be best in terms of the different criteria. Use of equations [2.1], [2.2] and [2.3] indicates that y must be higher than 0,923 according to the first criterion (equation [2.1]), higher than 0,818 according to the second (equation [2.2]) and higher than 0,869 according to the third criterion (equation [2.3]) (see Appendix A).

From the decision frame one know that $20(u(c_1) - u(c_3)) < u(c_3) - u(c_2)$, which is equivalent to $20x < y$. Thus, the statements here require that $y > 0{,}909$ (by use of $2x + y = 1$). The requirements are placed on y because of the decision frame shown in Figure 7, where it can be seen that alternative 2 (installing a sprinkler system) is the best alternative in two of the three criteria (criteria 2 and 3). It can also be seen that not much is lacking for alternative 2 to be the best also according to criterion 1. This implies that one should try to contract the decision frame in order to determine whether it is possible that alternative 2 can also be the best one according to criterion 1.

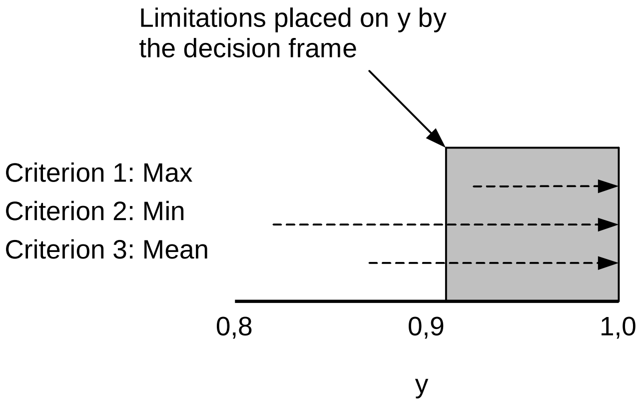

Figure 7 *Illustration of the requirements placed on y. For alternative 2 to be the best alternative, y must be in the range represented by the arrows. The grey area shows where y must be on the basis of the statements contained in the decision frame.*

In this case the most appropriate contraction can be judged to be a contraction of the value part of the decision frame. This is performed by looking at the last statement in Table 5, concerning the difference in utility value between the different consequences and noting that the last statement is the one that results in $20x < y$ and thus satisfies the requirements placed on y (Figure 7). If one can contract the limitations so that $y > 0{,}923$, then one can say that alternative 2 is the best alternative according to SSD. In order to do so, one needs to accept the condition that y is approximately 24 times as high as x. This translates into a statement such that, "The value distance between c_3 and c_2 is more than 24 times as large as the distance between c_1 and c_3". In practice, this means agreeing to the following: One evaluates the difference between a) investing in a sprinkler system and *not* having a major fire and b) *not* investing in a sprinkler system and *having* a major fire, a difference that can be called *A*. One also evaluates the difference between c) *not* investing in a sprinkler system and *not* having a major fire and d) investing in a sprinkler system and *not* having a major fire, a difference that can be called *B*. If one considers the difference *A* to be at least 24 times as large as *B*, then the alternative of installing a sprinkler system (alternative 2) is the one which is best.

This example shows how one can use SSD in assessing a decision situation concerning fire protection. Although we know rather little about the problem, we can make some vague statements and translate these into inequalities that can be used to analyse the problem. The results of the analysis conducted may mean one's needing to adjust one's initial statements so as to be able to arrive at a solution to the problem. SSD indicates how much one needs to change one's initial statement so as to be able to draw conclusions.

For now, applying SSD to practical problems is cumbersome and can only be done with problems of limited character. A computer program that could handle SSD would be of great help and, if such a program were available, the method would be highly useful. Use of the SSD method would have several advantages then. The most obvious is that the evaluation time would probably be short, and it would also be advantageous to be able to use vague expressions regarding the consequences of fires and of the probabilities concerning fire spread.

2.4. Utility functions

In practical applications of decision theory to problems of fire protection, it is cumbersome to have to estimate the utility values of the different consequences, since the number of possible consequences can be considerable. It is easier instead to assign monetary consequences and then translate these into utilities. The function that does this is called the utility function.

The form of the utility function determines whether the decision maker is risk-neutral, risk-seeking or risk-averse (Clemen, 1996). A decision maker who is characterised as risk-neutral evaluates consequences in accordance with their monetary value, which means using the expected monetary value of the uncertain situation and comparing it with that of some other situation (if he/she uses PMEU). A person characterised as risk-seeking assigns a greater utility value to a positive monetary consequence than a risk-neutral person does; if this risk-seeking person behaves in accordance with PMEU, he/she would always agree to participate in a lottery in which the expected monetary value was as high as the cost of participating in the lottery. A person characterised as risk-averse, in turn, would assign a lower utility value to a positive monetary consequence than a risk-neutral person would, and would thus not participate in a lottery if the price of participating was the same as the expected monetary outcome.

If one diagrams the utility functions, one notes that the risk-seeking function is concave, that the risk-neutral function is linear and that the risk-averse function is convex (Figure 8).

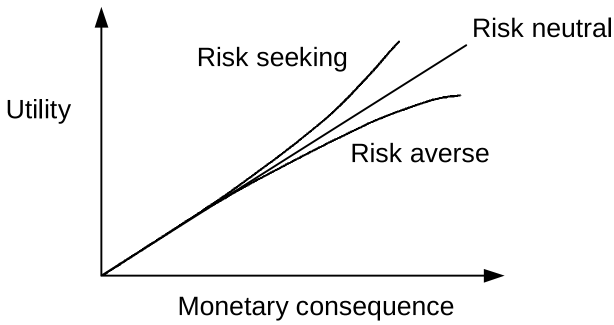

Figure 8 Diagram of the utility functions of persons characterised as risk-seeking, risk-neutral and risk-averse, respectively.

In the present context the reason for not choosing the risk-neutral utility function, i.e. using the expected monetary value as the basis for a decision, can be that one wants to assign a greater "weight" to more serious consequences. In practice, this would mean one's not evaluating a loss of 100 thousand SEK say as being a hundred times as bad as a loss of 1 thousand SEK. If one is risk-averse, one in fact, evaluates a 100 thousand SEK loss as being *more* than a hundred times as bad as a 1 thousand SEK loss.

A risk-averse utility function can be a *help* when one wants to assign utilities to the different outcomes. Instead of having to estimate the utility of every monetary consequence, one can simply let the utility function do it. All one has to do is to assign the utility function to be used.

In practise, it is best to use a utility function that is easy to work with; one such function is the exponential utility function, shown in equation [2.4] (Clemen, 1996). The exponential utility function is determined by only one variable R, the risk tolerance. In equation [2.4] $U(x)$ is the utility value assigned to the monetary consequence x. Utility functions will also be discussed in chapter 5.

$$U(x) = 1 - e^{-x/R} \tag{2.4}$$

Another concept closely linked to the utility function is the *certainty equivalent* (CE). Take as a point of departure the lottery shown in Figure 9. If one assumes that the probability of winning (p_{Win}) is 0,5, then the expected monetary outcome of the lottery is 50 SEK. If one evaluated the lottery using a risk-averse utility function, one might end up with a utility value of 0,63 for the outcome of 100 SEK and with the utility value 0 for the consequence 0 SEK. What does it mean to say that the utility value of the consequence 100 SEK is 0,63? Although this utility value seen in isolation has no specific meaning, comparing it with other utility values allows one to rank different outcomes in order of preference and also to use these various utility values to determine how much more desirable one particular consequence is than another.

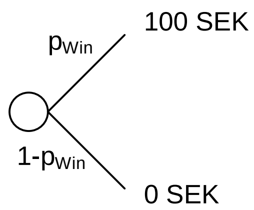

Figure 9 Description of a lottery.

The *CE* for a lottery is defined in equation [2.5], p_i there being the probability of consequence i, $U(x_i)$ is the utility value of the monetary amount x_i (received if consequence i occurs) and n is the number of consequences.

$$U(CE) = \sum_{i=1}^{n} p_i \cdot U(x_i) \tag{2.5}$$

The CE is the price one would pay for participating in the lottery, given that one makes decisions in accordance with PMEU and that one evaluates the monetary outcomes in accordance with the utility function. Using the risk-tolerance value (*R*) of 100 SEK indicated in equation [2.4], the *CE* for the lottery shown in Figure 9 is 38 SEK, which means that the decision maker would agree to paying 38 SEK to participate in the lottery.

The concept of *CE* will be used later when the investment appraisal of fire protection systems is discussed.

3. Bayesian updating

In the previous chapter it was concluded that lack of information can affect people's preferences regarding different alternatives in a decision situation. For decision making in fire risk management, it can be useful to know how one can *improve* one's knowledge of the specific probabilities or frequencies to be used in the decision model through utilising new information concerning these probabilities/frequencies.

The chapter will start by introducing Bayes theorem, which is the central part of the Bayesian updating process, and then go on to present various examples of Bayesian updating when new information is obtained.

Updating a probability means that one first uses the information already available (relevant statistics, and the like), including one's own assessments made earlier, so as to produce a *prior-estimation*. The prior-estimation is used then in combination with new information, for example concerning the building of interest, to form the *posterior-estimate*. The posterior-estimate could thus be a representation of general information, of expert judgement and of specific information concerning the building of interest. This posterior-estimate can be used then in decision analysis concerning problems in which this probability is involved.

A section presenting the results of an investigation of probabilities concerning the extent to which a fire in a factory can be expected to spread will conclude the chapter. The information from this study can be used as the basis for the initial subjective estimations performed, i.e. the prior-estimations.

3.1. Bayes theorem

In any situation to which risk analysis or decision analysis can be applied, new information may be received making it necessary to revise the belief one had regarding some parameter in one's model of the problem. How should this new information be used to revise one's old belief in a logical way? This is the question that Bayes theorem provides an answer to. Bayesian updating is a formal way of combining both subjective and general information with objective information pertaining to a specific building.

Basically speaking, one starts with some belief one has about a specific parameter, such as a probability pertaining to some possible event during a fire. This prior belief about the probability in question may have originated from the judgement of experts combined with the use of general statistics pertaining to the type of building, factory or whatever involved, from visual inspection of the premises, or whatever. What is important about this initial probability is that it is subjectively estimated (see the previous chapter) and that objective information can be used to revise it. Revising it properly means that if two persons start out with two completely different probability estimates, they should nevertheless end up with approximately the same final estimate if the amount of new information they receive is sufficiently large.

The problem is basically that of having an initial estimate of the probability of a particular state, of having received new information pertaining to this probability and of wanting to update this initial probability in a logically and consistent way on the basis of this new information. Let $P(S_1)$ denote the initial probability that State 1 is the true state. In the present context this state could be a particular value of a probability in the model of fire spread employed. State 1 could be, therefore, that the probability is 0,1, for example. Note that one does not need to actually observe which state is true. If one did, the probability for the state in question would then be either 1 or 0. However, one does need to observe indirect information, i.e. information affected by the true state. What one wants to obtain is $P(S_1/NS)$, the probability that State 1 is the true state, given that some new information, termed here New Statistics (NS), has been received.

Consider two products, on the left- and right-hand side of the following equation, each representing the probability of the truth of State1 *and* of the existence of New Statistics (equation [3.1]). The equation follows from elementary probability concepts.

$$P(S_1/NS)P(NS) = P(NS/S_1)P(S_1) \quad [3.1]$$

By rearranging the elements in equation [3.1], one can create the expression shown in equation [3.2].

$$P(S_1/NS) = \frac{P(NS/S_1)P(S_1)}{P(NS)} \quad [3.2]$$

The total probability theorem allows one then to replace the *P(NS)* of equation [3.2] by the sum of the probabilities of the New Statistics in question having been observed given all states that are possible. The result, shown in equation [3.3], is called Bayes theorem.

$$P(S_1/NS) = \frac{P(NS/S_1)P(S_1)}{\sum_{All\ States\ i} P(NS/S_i)P(S_i)} \quad [3.3]$$

Equation [3.3] only shows the calculation of the probability that State 1 is the true state. If one wants to calculate the probability of any one of the other states that are possible one simply replaces State 1 in equation [3.3] by whatever state one wishes to do the calculations for.

In Bayes theorem, $P(S_1)$ is called the *prior probability* of the event that State 1 is the true state; $P(S_1/NS)$ is the *posterior probability*, i.e. the probability that State 1 is the true state as assessed after one has observed the evidence contained in New Statistics. $P(NS/S_1)$ is a likelihood-function expressing the probability of the evidence that New Statistics contains being observed *given* that the true state of the world is State 1.

If one considers the initial probability of each of the possible states, one obtains the prior probability *distribution*. The result of Bayesian updating will then be the posterior probability distribution. For example, if there are only five possible states and one deems them to be equally likely, then the prior distribution should look as it does in Figure 10.

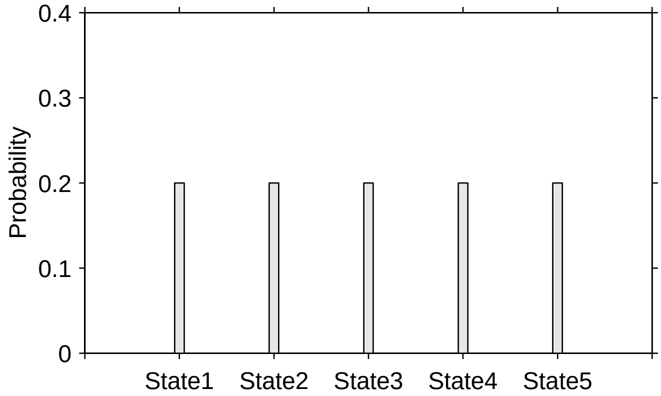

Figure 10 Example of a prior distribution.

To demonstrate the use of the Bayesian updating process, two examples will be given. The first is an example of the estimation of the probability of a release of radioactive material during transport. The example shows how statistics concerning 4000 accident free transports can be used to update one's belief concerning the release frequency. The second example concerns the updating of frequency of fire in a factory with the help of fire statistics pertaining to the previous five years.

The first example, one of Kaplan and Garrick (1979), concerns the frequency of release of radioactive material in the transport of spent nuclear fuel by train. Public discussion was underway concerning the frequency of release, some people claiming that transport of spent radioactive fuel by rail was extremely dangerous and that the 4000 release-free transports there had been thus far did not constitute any meaningful evidence for the safety of such transport since 4000 was a very small number as compared with the 10^8 to 10^{10} transports in which, according to official estimates one, accident could be expected to occur.

Kaplan and Garrick argued that it was not the issue of whether the frequency of a release would be once in 10^8 or 10^{10} transports that was important but whether the spent nuclear fuel could be transported "safely" or not, i.e. if the probability of a release would be in the order of once in 10^2 to 10^3 transports or once in 10^8 to 10^{10}. They also showed that with respect to that question the 4000 release free transports constitute very important evidence. This was done using Bayesian updating technique.

Kaplan and Garrick (1979) define their terms as follows (p. 233):

> Let
>
> B stand for the statement "we have 4000 shipments with no releases".
>
> Let
>
> A_1 stand for the statement "the frequency rate is 10^{-3}"
> A_2 stand for the statement "the frequency rate is 10^{-4}"
> A_3 stand for the statement "the frequency rate is 10^{-5}"
> A_4 stand for the statement "the frequency rate is 10^{-6}"
> A_5 stand for the statement "the frequency rate is 10^{-7}"
> A_6 stand for the statement "the frequency rate is 10^{-8}"

Using the same notation as just presented, the prior probability distribution for the frequency rate can be written as $P(A_i)$, i = 1,2,…,6.

The prior distribution was, by use of expert judgement, assumed to have the form shown in Figure 11.

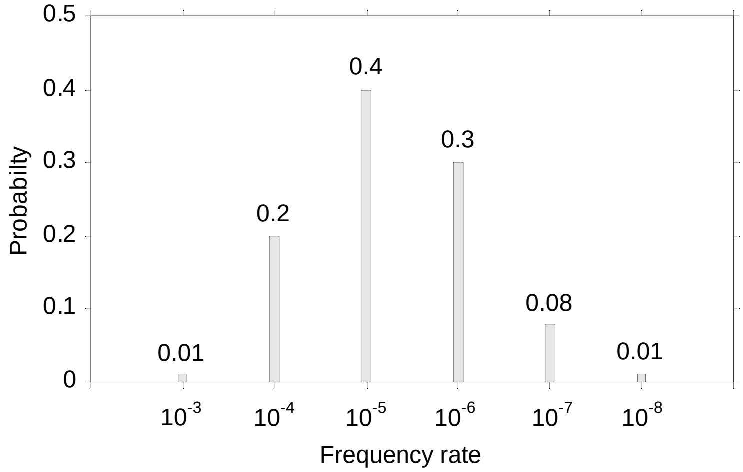

Figure 11 Prior distribution for the frequency of radioactivity release per shipment of spent atomic fuel.

To obtain the result, i.e. the posterior distribution of $P(A_i|B)$, i = 1,2,…,6, one uses the same reasoning as described for the derivation of Bayes theorem. Note that equation [3.4] is the same as equation [3.1] but is written with the notation used for this problem.

$$P(A_i|B)P(B) = P(B|A_i)P(A_i) \quad [3.4]$$

Rearranging the terms in equation [3.4] and using the total probability theorem yields equation [3.5], which is Bayes theorem.

$$P(A_i/B) = \frac{P(B/A_i)P(A_i)}{\sum_{i=1}^{6} P(B/A_i)P(A_i)} \quad [3.5]$$

In equation [3.5], $P(A_i)$ for all i is shown in Figure 11 and $P(B|A_i)$ is the probability that 4000 release-free transports would have been observed given the frequency rate of A_i. This probability can be calculated for A_1 by use of equation [3.6].

$$P(B/A_1) = (1-10^{-3})^{4000} = (0,999)^{4000} = 0,0183 \quad [3.6]$$

Using the same method of calculation as in equation [3.6] allows one to create Table 7.

Table 7 *The probability that 4000 transports would have been observed to be release-free, given a particular release frequency (A_i).*

A_i	$P(B \mid A_i)$
A_1	0,01828
A_2	0,67031
A_3	0,96079
A_4	0,99601
A_5	0,99960
A_6	0,99996

As can be seen in Table 7, the probability is very low that 4000 release-free transports would have occurred given that the release frequency is 10^{-3}, i.e. one accident in 1000 transports. One can see at the same time that the probability is very high that the 4000 release-free transports would have occurred if the release frequency rate had been quite low, 10^{-5} (A_3) or still lower.

It is now possible to calculate the posterior distribution for the accident frequency rate using Bayes theorem (equation [3.5]). The resulting posterior distribution is shown in Figure 12.

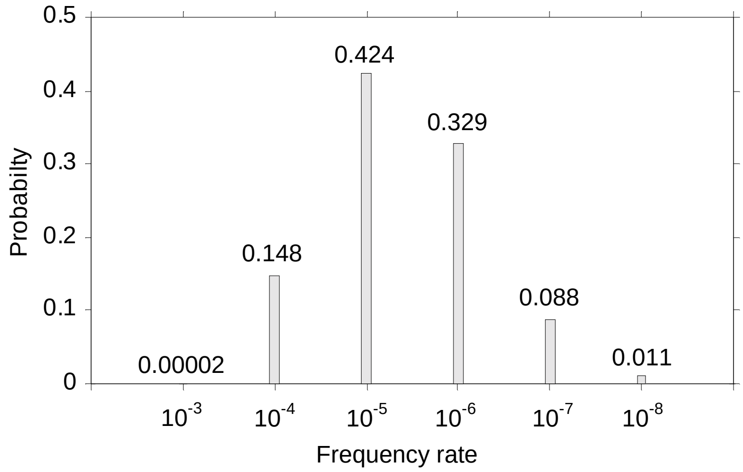

Figure 12 *Posterior distribution for the accident frequency rate.*

From Figure 12 it can be concluded that the 4000 release-free transports indeed constitute valuable evidence concerning the safety of transporting radioactive material. The evidence virtually eliminated the possibility that the release frequency is in the order of once in every 1000 years (10^{-3}) and it considerably lowered the probability that the release frequency is in the order of once in every 10000 years (10^{-4}). Regarding still lower frequency rates, below 10^{-4}, there was not much of a change, since the number of release-free transports was not high enough to strongly influence those frequencies.

This example shows how Bayes theorem can be used to adjust a subjectively estimated prior distribution with the help of new information, here of a statistical character. Thus, Bayes theorem represents a logical way of combining subjective judgements with objective statistics or measurements, one which is very useful for decisions concerning different issues related to fire protection.

Consider the following example of how Bayesian updating can be used in fire risk management so as to provide a basis for a decision that an engineer has been asked to make in connection with the fire risk analysis of a specific factory belonging to the metalworking industry. A highly important factor in such an analysis is the fire frequency. Assume that the engineer wishes to obtain as good an estimate of the fire frequency as possible. The information available to the engineer is his/her own general experience, his/her own subjective judgements concerning the specific building and general information showing the fire frequency in other buildings within the metalworking industry.

Since the information available is not specific to the building at hand, despite its applying to the category of industry involved, the engineer needs to subjectively adjust the information to fit the conditions present in the building at hand. Assume that the engineer has difficulties in determining a specific fire frequency for the building, considering it highly likely, for example, that the fire frequency is somewhere between 1 and 5 fires per year, but is unwilling to assign a specific value to the parameter and desires more information so as to be able to make a better estimate. The engineer can represent his/her estimate of the fire frequency using a prior probability distribution, prior inasmuch as information the engineer receives or takes account of later may lead to this estimate being revised. The prior probability distribution could look as that does in Figure 13, for example. As can be seen in the figure, the engineer has assigned no preference to any value in the range of 1 to 5 fires per year but rather considers it just as likely that the fire frequency is 1 per year as that it is 2,5, 4 or any other of the possible values.

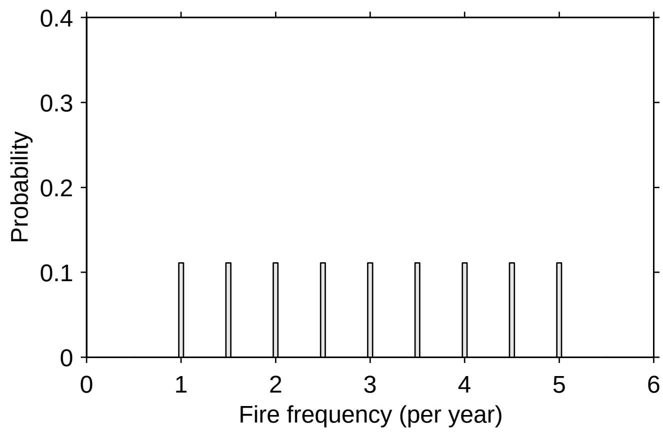

Figure 13 Prior probability distribution for the fire frequency.

In chapter 2 the reliability weighted expected utility (RWEU) method was discussed. In that method, a probability distribution is used to represent the uncertainty concerning some “true” parameter value. The RWEU method fits remarkably well with the concept of Bayesian updating, since the assigned probability distribution, defined over a range of different values of the uncertain parameter (such as a probability) which is involved, can be used as a prior probability distribution in the Bayesian updating method. This means that if one starts out by using the RWEU method and assigns a probability distribution to each of the uncertain probabilities and uncertain frequencies in the model, one can use the Bayesian updating technique in combination with new information in order to produce new and updated (posterior) probability distributions for the values of the parameters.

Assume that in order to adjust his/her initial belief regarding the fire frequency shown in Figure 11 the engineer wishes to use statistics pertaining to the specific building of interest. Assume that there have been nine fires in the building during the past five years. This information can now be incorporated into the previous body of knowledge (the prior distribution) by use of Bayes theorem. Let λ_1 stand for the statement “the fire frequency is 1 per year on average”, λ_2 for “the fire frequency is 1,5 per year on average” and so on in accordance with Figure 13. Bayes theorem can be expressed then as in equation [3.7], in which NS refers, as earlier, to New Statistics.

$$P(\lambda_j / NS) = \frac{P(NS/\lambda_j)P(\lambda_j)}{\sum_{i=1}^{9} P(NS/\lambda_i)P(\lambda_i)} \qquad j = 1,2,\ldots,9 \qquad [3.7]$$

$P(\lambda_j)$ is 1/9 for all j (see Figure 13). $P(NS|\lambda_j)$ in turn can be calculated by use of a Poisson distribution.

Assume that there have been nine fires in the building during the past five years. The Poisson distribution can be used then to calculate the probability that nine fires would have occurred in five years *given* some specific value for the fire frequency. The values for the fire frequency to be used here are shown in Figure 13. In using the Poisson distribution, one assumes that the fires in the building occur randomly and are independent of each other.

Calculation of the probability that nine fires would have occurred in five years, given that the fire frequency was 1 fire per year, will now be shown. Calculation of the other fire frequencies is in principle the same, but involves different frequency values. In the calculations a Poisson distribution is employed in which λ is the fire frequency per year, t is the period of time in years and k is the number of fires that occurred during those years. Use of equation [3.8] indicates the probability to be 0,0363 of 9 fires occurring in 5 years, given that the fire frequency is 1 per year.

$$P(NS/\lambda_1) = e^{-(\lambda t)}(\lambda t)^k / k! = e^{-(1\cdot 5)}(1\cdot 5)^9 / 9! \approx 0,0363 \qquad [3.8]$$

The remaining probabilities $P(NS|\lambda_j)$ are shown in Table 8.

Table 8 $P(NS|\lambda_j)$ as a function of fire frequency.

λ_j	$P(NS|\lambda_j)$
λ_1	0,0363
λ_2	0,1144
λ_3	0,1251
λ_4	0,0765
λ_5	0,0324
λ_6	0,0107
λ_7	0,0029
λ_8	0,0007
λ_9	0,0001

It is now possible to calculate the posterior probability for each of the fire frequencies here, using Bayes theorem. The calculation for the fire frequency of 1 fire per year is shown in equation [3.9]. The posterior probability for each of the fire frequencies, i.e. the posterior probability distribution, is shown in Table 9 and in Figure 14.

$$P(\lambda_1/NS) = \frac{P(NS/\lambda_1)P(\lambda_1)}{\sum_{i=1}^{9} P(NS/\lambda_i)P(\lambda_i)} = \frac{0,0363 \cdot 0,1111}{0,0443} = 0,0910 \quad [3.9]$$

Table 9 Posterior probabilities for the different fire frequencies.

λ_j	$P(\lambda_j \mid NS)$
λ_1	0,0910
λ_2	0,2867
λ_3	0,3133
λ_4	0,1917
λ_5	0,0812
λ_6	0,0267
λ_7	0,0073
λ_8	0,0017
λ_9	0,0004

The posterior probability distribution of the fire frequency shows that the probability is very small that the fire frequency in the building in question is greater than four fires per year. The result obtained, in the form of the posterior probability distribution, can be employed in a decision analysis using the RWEU method.

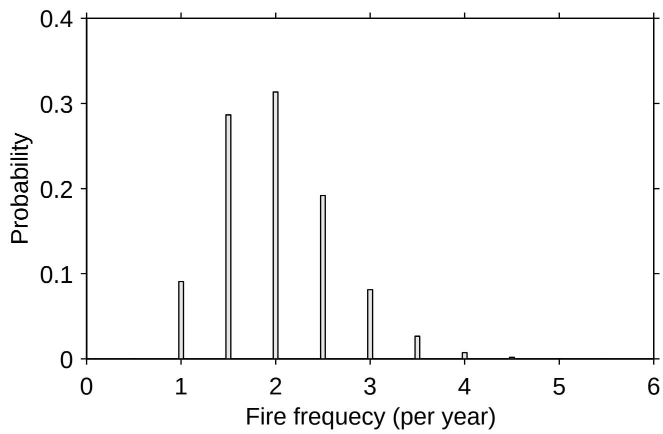

Figure 14 The posterior probability distribution of the fire frequency.

The Bayesian updating process constitutes a highly useful way of improving previous estimates concerning events affecting a decision and it fits well with use of the RWEU method. The updating process is also very practical since a posterior distribution from one year can be used as a prior distribution the next year. Thus, the updating process can be used not only as the basis for a decision but also for continuously monitoring the fire risk in a given building, for example (see Johansson 2000a).

3.2. The probability of different fire spread

In order to produce the first prior distribution of a probability that applies to a particular building, it is desirable to have an estimate of the probability that would apply to the category of buildings to which the specific building can be assumed to belong. This probability can be used as the basis for estimating the probability of interest. For a more thorough discussion of this, see the paper included in Appendix F.

By the use of statistics of fires that occurred in Sweden, it is possible to identify approximately the extent to which a particular fire would spread before being extinguished. This information can be used then in order to calculate the probability for different degrees of fire spread. To do this, one needs a model of how a fire can develop in a building. The model used here is shown in Figure 15, where four possible fire scenarios that can occur are shown.

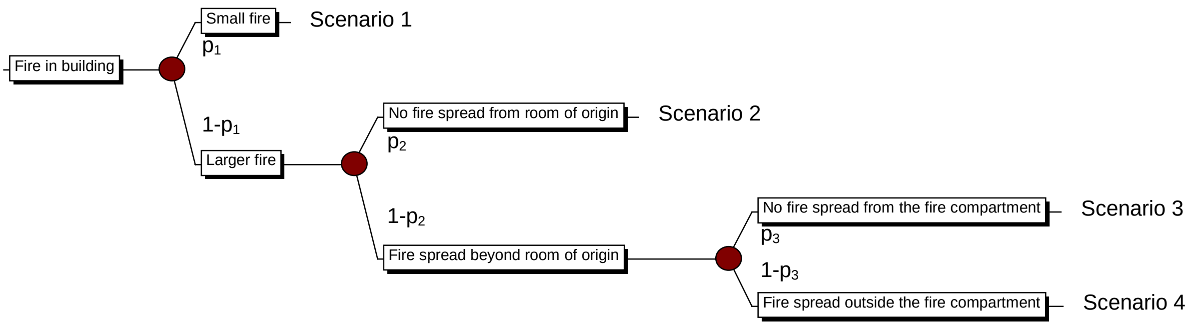

Figure 15 Description of the model for fire spread in a building.

Fire statistics collected by Räddningsverket (the Swedish Rescue Service Agency) during 1996, 1997 and 1998 are summarised in Table 10, showing for different industries and for conditions of a building being with and without sprinklers.

Table 10 Extent of fire spread in buildings in different industries in Sweden during 1996, 1997 and 1998.

	Scenario 1	Scenario 2	Scenario 3	Scenario 4	Total number of fires
Buildings without sprinklers					
Metalworking and machine industry	425	357	31	39	**852**
Chemical industry	124	106	3	18	**251**
Food manufacturing industry	91	81	4	12	**188**
Textile industry	26	20	2	6	**54**
Warehouses	39	79	5	35	**158**
Forest-product industry	220	283	24	78	**605**
Other branches of manufacturing	276	241	8	36	**561**
Repair shops	50	113	17	67	**247**
Buildings with sprinklers					
Metalworking and machine industry	94	40	3	2	**139**
Chemical industry	24	22	1	5	**52**
Food manufacturing industry	18	20	1	1	**40**
Textile industry	7	11	0	1	**19**
Forest-product industry	94	82	7	27	**210**
Other branches of manufacturing	115	59	12	7	**193**

Knowing in this way how many of the fires have resulted in spread of a certain degree allows one to estimate the different probabilities shown in Figure 15, together with 95% confidence intervals. A description of how these estimates were obtained is found in Appendix F. Note that the probabilities are conditional probabilities, meaning for example that p_2 is conditional upon that a fire having started and upon the fire not having developed according to Scenario 1 (see Figure 15).

The result obtained in analysing the fire statistics for the different industries is an estimate followed in each case by a 95% confidence interval of the estimate. These results are presented in Table 11, where $I_{p1,min}$ is the lower boundary and $I_{p1,max}$ the upper boundary of the confidence interval.

Table 11 Estimates of the probabilities contained in the fire spread model together with the 95% confidence interval for each estimate.

	$I_{p1,min}$	**p_1**	$I_{p1,max}$	$I_{p2,min}$	**p_2**	$I_{p2,max}$	$I_{p3,min}$	**p_3**	$I_{p3,max}$
Buildings without sprinklers									
Metalworking and machine industry	0,47	**0,50**	0,53	0,80	**0,84**	0,87	0,33	**0,44**	0,56
Chemical industry	0,43	**0,49**	0,56	0,77	**0,83**	0,90	-	**0,14**	-
Food manufacturing industry	0,41	**0,48**	0,56	0,76	**0,84**	0,91	-	**0,25**	-
Textile industry	0,35	**0,48**	0,61	-	**0,71**	-	-	-	-
Warehouses	0,18	**0,25**	0,31	0,58	**0,66**	0,75	-	**0,13**	-
Forest-product industry	0,33	**0,36**	0,40	0,69	**0,74**	0,78	0,15	**0,24**	0,32
Other branches of manufacturing	0,45	**0,49**	0,53	0,80	**0,85**	0,89	-	**0,18**	-
Repair shops	0,15	**0,20**	0,25	0,50	**0,57**	0,64	0,12	**0,20**	0,29
Buildings with sprinklers									
Metalworking and machine industry	0,60	**0,68**	0,75	-	**0,89**	-	-	-	-
Chemical industry	0,33	**0,46**	0,60	-	**0,79**	-	-	-	-
Food manufacturing industry	-	**0,45**	-	-	**0,91**	-	-	-	-
Textile industry	-	**0,37**	-	-	**0,92**	-	-	-	-
Forest-product industry	0,38	**0,45**	0,51	0,62	**0,71**	0,79	-	**0,21**	-
Other branches of manufacturing	0,53	**0,60**	0,67	0,66	**0,76**	0,85	-	**0,63**	-

One should remember that the estimates presented in Table 11 are estimates, for each of the categories, of the *mean* value of the probability. Even if one had exact knowledge of the value of this parameter (which we do not) there would still be uncertainty concerning the parameter value in a *specific* factory or building belonging to a particular industrial category. Therefore, one can only use the estimates presented here as a point of departure for estimating the parameter value in any *given* industrial building.

For example, assume one is employing the same model of fire spread as shown in Figure 15 and that one wants to estimate the parameters of the model in a specific building (building *A*) that has no sprinklers, one that belongs to the forest-product industry. Suppose that since one has no other information than that presented in Table 11 one decides to use the estimate of the mean value of the parameters for the forest-production industry. Since one knows that the values for a specific company and a specific building are likely to deviate from the mean value for the industrial category in question, one can choose to represent one's estimate in the case at hand by a prior-distribution (second order probability distribution) that represents one's beliefs concerning what values are most probable. The prior-distributions of the parameters are shown in Figure 16, Figure 17 and Figure 18. The probability $p_{1,A}$ is the probability that a fire in the *specific* building A will develop according to Scenario 1.

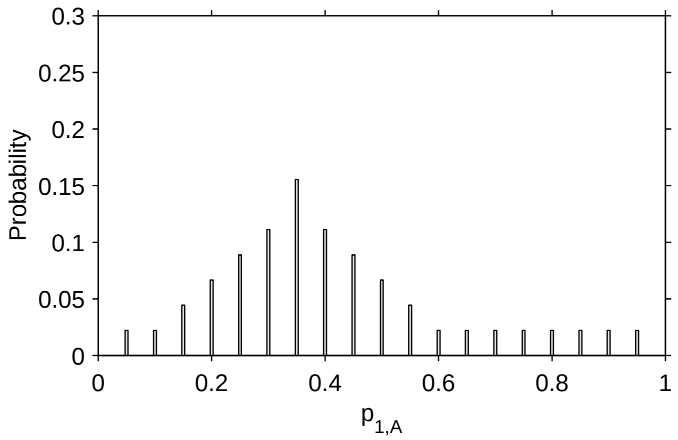

Figure 16 The prior-distribution of $p_{1,A}$.

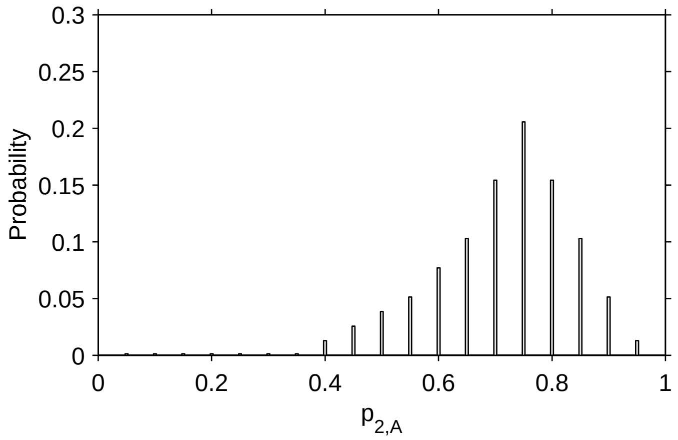

Figure 17 The prior-distribution of $p_{2,A}$.

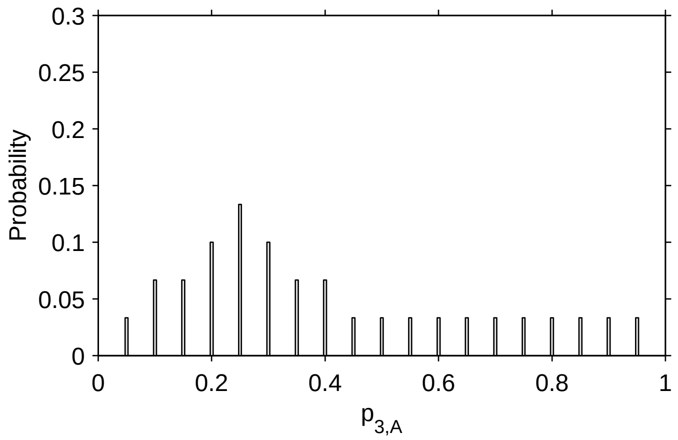

Figure 18 The prior-distribution of $p_{3,A}$.

The choice of a prior distribution is influenced here by the mean value of the parameters as presented in Table 11. For example in Table 11 the estimate of the mean value of p_1 for the forest-product industry is 0,36, the confidence interval being fairly small (lying between 0,33 and 0,40). According to the prior-distribution (for $p_{1,A}$), as shown in Figure 16, 0,35 is the value with the highest probability, but there is a relatively high probability too that the value can be as high as 0,55. It should also be noted that it is the prior distribution for $p_{3,A}$ which has the broadest spread that parameter thus being the one about which one is most uncertain. This is because not many observations of fires are available that could be used to estimate this parameter.

When a prior distribution has been selected, the next thing to do is to investigate the fire statistics for the building in question. Assume that there have been 17 fires in the building and that 5 of these developed according to scenario 1 (Figure 15) and the rest according to scenario 2. This information can be used to improve the original prior-distributions for $p_{1,A}$ and $p_{2,A}$. By use of the Bayesian updating procedure as described earlier in the chapter, one obtains the posterior distributions shown in Figure 19 and Figure 20. In those figures are the prior distributions displayed as well, so as to make the comparison easier.

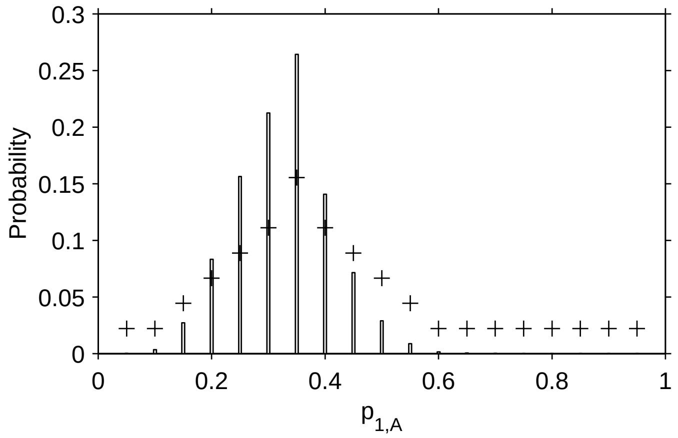

Figure 19 Posterior-distribution of $p_{1,A}$. The prior-distribution is plotted as +.

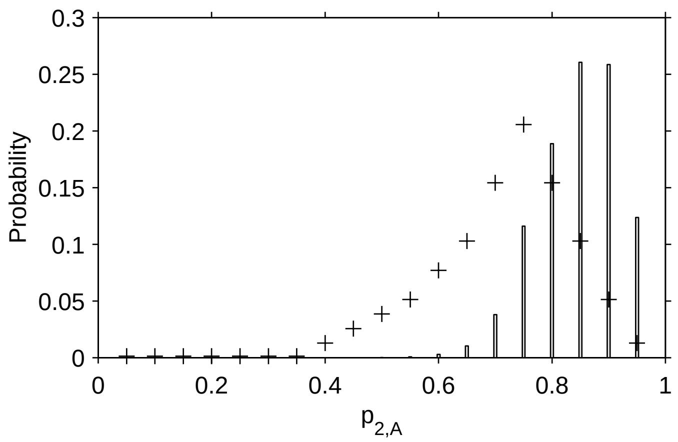

Figure 20 Posterior-distribution of $p_{2,A}$. The prior-distribution is plotted as +.

This example shows how general estimates concerning the mean value for the probability for a particular type of event connected with fire in a building of a given industrial category can be used in combination with Bayesian updating. On the basis of the resulting posterior probability distribution one can conclude that the updating procedure reduced the probability of occurrence of $p_{1,A}$-values that were higher than 0,45 and significantly reduced the probability of occurrence of $p_{2,A}$-values that were below 0,75.

The reason for not updating $p_{3,A}$ is that it was assumed that no fire had been classified as belonging to scenario 3 or 4, making it impossible to update $p_{3,A}$.

4. Decision making concerning fire protection

A number of different methods for risk assessment in connection with fire protection are available, ranging from checklists and index methods (e.g. the Gretener Method (BVD, 1980) and the Dow Fire and Explosion Index (Watts, 1995)) to quantitative methods of risk analysis (e.g. Frantzich (1998)). However, these methods focus on the estimation of risk and do not deal with the evaluation of different fire protection measures. Although it would be possible, of course, to analyse various alternatives by use of such methods and to calculate the differences between them (expressed in whatever unit the method in question employs). This difference does not provide an answer to the question of whether a particular investment, of say 100.000 SEK, was worth what it cost but can only be used to assess whether or not the risk was reduced. In contrast, a quantitative risk analysis could be extended so as to provide the basis for a decision analysis concerning different fire protection measures.

This chapter will focus on how to evaluate different fire protection alternatives. The methods employed will be based on traditional Bayesian decision theory as presented in chapter 2, although with modifications so that imprecise probabilities and imprecise consequences can be dealt with explicitly (see section 2.3.1). The type of decision analysis considered will also be one involving a Bayesian updating process such as discussed in chapter 3 used for reducing uncertainty concerning the frequency of fire[6] in a specific building.

Although there can be many different objectives for a decision concerning fire protection, economic objectives are the only type of objectives that will be dealt with explicitly in the present chapter. The reason for not addressing the safety of the people involved, for example, is that if the fire protection of a building complies with building regulations, a matter which is taken for granted in the thesis, the safety of these people is often assumed to be adequate. If, on the other hand, a building owner wants to increase the safety of people above and beyond the demands made by the building code, one can end up with the difficult task of setting a value on human life. The valuation of human life is not a part of this thesis and will thus be disregarded in the decision analysis.

The scope of the decisions considered in this chapter is limited. The methods of decision making suggested are suitable for decisions on fire protection for one or several buildings with the primary objective of minimising the sum of the costs of fire protection and the costs due to fire. The reason for this restriction is that if the decision concerns larger safety investments where there are other objectives than the ones just mentioned, the decision process can become considerably more complicated, exceeding the intentions of the thesis. Although there are methods for dealing with decisions involving multiple objectives, using these methods in the present context would greatly increase the extent of the decision analysis called for and the workload required to produce it. Using only economic objectives here appears reasonable inasmuch as the objective of the thesis is to suggest a *practical* method using normative decision analysis in fire risk management. The fact that only economic objectives are explicitly treated in the decision analysis implies the decision maker's needing to make a non-explicit subjective evaluation of all the other possible objectives, such as flexibility, environmental issues, and the like when comparing the available alternatives and making a decision.

In the remainder of this section, the economic losses that a decision analysis deals with are discussed. The section focuses mainly on different levels of analysis that are possible

[6] The Bayesian updating method can also be used for probabilities.

depending on how much effort one can put into investigating the economic consequences of a fire.

A method for evaluating the risk reduction achieved by fire protection measures is then presented briefly. The method involves the use of the annual expected costs due to fire. Suggestions of how to use this measure in assessment pertaining to a specific building are provided.

4.1. Economic losses

As just indicated, only economic objectives are explicitly dealt with in the approach to decision analysis to be discussed. The question is what kind of economic objectives are to be involved. Should long-term strategic effects of a serious fire be a part of the analysis, or is it more practical to consider instead simply the direct and consequential (indirect) losses, those traditionally reported by the insurance companies.

The strategy in the decision analysis here will be to take account simply of the costs of fire that the company eventually has to defray. These costs, will be termed uncompensated costs, probably depend to a considerable extent on the type of business the company is in and the type of insurance it has; in the thesis no attempt will be to made to analyse exactly how the various uncompensated losses are determined. Uncompensated costs that could be important are those of additional marketing campaigns to retake customers after a production stop, and of the reduction in sales due to production stop, deductibles, etc. Since some companies may have insurance against various types of losses that another company does not have, it is difficult to provide any general advice concerning uncompensated losses, other than that an analysis of the specific situation with which the company is faced is often the best way to determine what the company's uncompensated losses in case of a fire are.

An analysis of fire safety investment can be performed on at least three levels, those of (1) ignoring the increase in safety and of basing the evaluation of the investment on parameters one is basically certain about, such as investment costs, reduction in insurance premiums, maintenance costs etc., (2) taking account of all costs at level 1 and adding to this the valuation of the risk reduction achieved by analysing the reduction in the annual expected costs due to fire, using a subset of the uncompensated costs or of any other costs for which the relation they have to the uncompensated costs can be estimated, or (3) taking account of all costs at level 1 and attempting to estimate the uncompensated costs directly.

An analysis at level 2 could for example be that the costs due to fire are both the direct and consequential losses that the company receives compensation for from insurance. Conclusions regarding the investment can then be drawn by estimating the size of the uncompensated losses in comparison with the costs taken account of. For example, in a case study performed at a cold-rolling mill belonging to Avesta Sheffield, the uncompensated losses were estimated to be of approximately the same size as the direct and the consequential losses compensated for by insurance (see chapter 5). In a case study performed at facilities belonging to Stora Enso (a major producer of forest products in Sweden), on the other hand, the uncompensated losses were estimated to be negligible, despite the direct and the consequential losses being of considerable magnitude (Wikström, 2000). The fact that the uncompensated losses in the Avesta Sheffield case were estimated to be in the order of several billion SEK for some of the more severe fire scenarios but to be practically neglectible in the Stora Enso case is an indication of the large variations found between different companies. This emphasizes how difficult it is to provide any general advice regarding the magnitude of uncompensated losses.

4.2. A model for the estimation of the expected annual cost due to fire

One of the most difficult aspects of decision making with respect to fire protection is that of evaluating the increased safety that different fire protection investments can bring about. This is due to the large uncertainties concerning both how often a fire will occur during the period for which planning is carried out and what will happen if a fire should occur.

As its title implies, this section will be concerned with estimation of the expected annual cost due to fire in a specific building. One can wonder why the focusing now is on the expected costs due to fire and not on the different alternatives and the consequences of each, given a certain state of the world, as was the case in chapter 2. This is due to its being a very cumbersome job to create a decision matrix containing each possible outcome in the case of a fire, together with its probability. Instead, this is done implicitly in the model describing the possible scenarios that can occur given that a fire has started in the building. The reason for having a problem analysed in terms of a decision matrix divided up into various consequences that depend on the state of the world is that this allows one to calculate the expected utility by summing for all possible consequences the various products of in each case the utility of a particular consequence and the probability associated with it. Doing this is difficult and time-consuming. An easier approach is to calculate the expected utility by using event trees instead of matrixes, the result of the event tree calculations just as of matrix calculations being expected utilities.

One may also wonder why the focus here has moved from that of expected utilities to expected costs due to fire. The reason for this is that in the present context only economic objectives are being considered. Also, if one assumes that the decision maker is risk-neutral (see section 2.4) the expected utility and the expected monetary outcome are the same. There is really no problem, however, in calculating the expected utilities given occurrence of fire in the building in question, since this requires simply that one translate the monetary outcomes to utilities by use of a utility function (see section 2.4) prior to evaluating the event tree.

In estimating the annual expected costs due to fire, use will be made both of the expected costs due to fire given that a fire has occurred in the building in question and the annual frequency of fire in the building. The product of these two values represents the expected annual costs due to fire.

4.2.1. Fire frequency

The average number of fires per year that is estimated as occurring in a building is called the fire frequency. The fire frequency is not the same thing as the probability of fire. Whereas the fire frequency is the average number of fires occurring during the time period in question, the probability of fire is a measure of how likely it is that *one or more* fire(s) will occur during this time period.

Since the fire frequency has a strong influence on the annual expected costs due to fire if the fire frequency is doubled, for example, the annual expected costs are also doubled it is of interest to endeavour to minimise the uncertainties concerning this parameter. To do this, the Bayesian updating method described in chapter 1 is utilised for combining subjective judgements with objective statistics.

In using the Bayesian concept of updating, one should first create in a subjective way a prior distribution representing one's uncertainty regarding the fire frequency in the building. This can be done by using information on fire frequencies in similar buildings, if that kind of information is available. One can also conduct a visual inspection in order to obtain an overall impression of the building, noting how many possible sources of fire there are, for example, such as machines running high temperatures, and the like. A hint of the fire sources to look for can be obtained by looking at statistics of what kinds of equipment are the most frequent sources of fire in the type of building considered.

If it is difficult to estimate the prior distribution, one can use a so-called diffuse prior, which is a prior distribution expressing no or virtually no preference for any particular value (Benjamin & Cornell, 1970), for example a uniform distribution. In doing this one expresses no preference for any particular fire frequency but instead estimates each possible fire frequency as being equally likely. In the present context, the most practical approach is to represent this with a discrete distribution. One can also express the prior distribution as a continuous distribution, but unless one chooses a so-called conjugate distribution, the updating process is cumbersome. Johansson (2000a) presents various conjugated distributions that can be useful in the context of fire protection, but in the present thesis continuous distributions will not be considered further, discrete distributions being used instead.

After having expressed ones belief about the fire frequency in the form of a prior distribution, statistics concerning the number of fires in the present building can be used to update or improve the initial estimate. This is done in the same manner as was described in chapter 3. The resulting posterior distribution can be used in the analysis of the different alternatives.

4.2.2. Expected cost due to fire

In order to calculate the expected costs due to *one* fire, one needs to create a model of what can happen if a fire breaks out and how likely each type of event is. The model used in the thesis will be described in this section briefly. For a more comprehensive account, see Johansson 2000a and Johansson 2000b.

The technique used to visualise the model is an event tree technique, meaning that the outcome of a fire is seen as being determined by a set of uncertain events as described in a tree. For example, if a fire occurs in a building it may occur in different areas of the building. This is modelled by a probability node in the event tree. Figure 21 shows part of an event tree, where an example of the probability node just mentioned is shown.

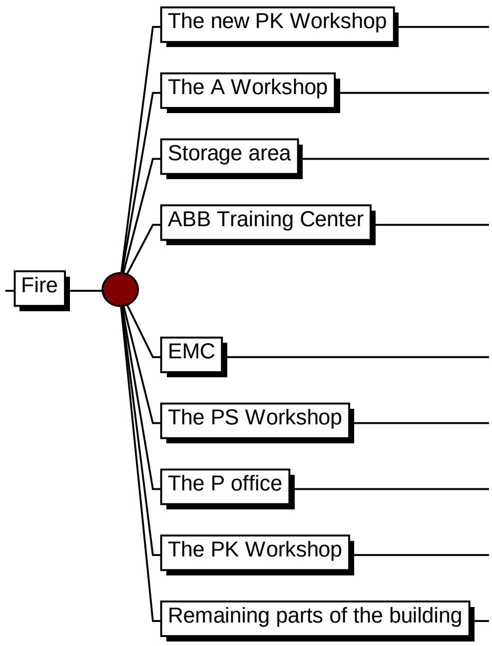

Figure 21 Probability node representing the probability of a fire starts in a specific area. The example is taken from the case study performed at ABB.

As can be seen in the figure, a fire that occurs in the building may have started in any one of nine different areas. This is represented by a node and by the events leading from the node. It is important that the probabilities of the events leading from a node sum to 1. Following the probability node concerning where the fire started, come such nodes as those representing the probability that the staff will extinguish the fire, that the sprinkler system will extinguish the fire (if such a system is present), that the fire cells will contain the fire, etc. Since what kinds of probabilities and events are considered in the event tree is dependent on the specific building involved, the modelling of a fire must always be done in an individual way that takes account of this. In the case studies performed at ABB and Avesta Sheffield, the following protection "systems" were considered:

Active systems. Systems designed to actively extinguish a fire should be considered in the analysis. These could be water sprinkler systems, for example, CO_2-systems, light water systems, etc. In estimating the conditional probability that an active system will extinguish the fire (i.e. conditional on all preceding events in the event tree), one can use as a point of departure whatever investigations are available concerning the reliability of the system (see National Fire Protection Association, 1976, for example). However, one should remember that the numbers presented in such investigations are estimates of the mean values found for a whole group of systems and that the reliability of a particular system may differ from this. One should best use the values obtained in investigations of this sort simply as a point of departure in estimating the reliability of a *specific* system.

Passive systems. System (such as a wall) designed to *stop* a fire from *spreading* further in a building but not designed to actively extinguish the fire were likewise considered. Investigations regarding the reliability of fire-rated walls or fire-rated windows, for example, appears to not be as common as those concerning active systems. This makes it more difficult for the decision maker to estimate the conditional probabilities involved. However, since one can tolerate probabilities being given in an imprecise way, one can accept the decision maker is representing a conditional probability by an interval or by a probability distribution.

Fire department. Since the fire department can affect the outcome of a fire, it can be represented by the conditional probability that it will succeed in extinguishing a fire. This probability could be estimated in cooperation with representatives of the fire department in question and would probably be estimated in terms of a large probability interval (an imprecise probability). Särdqvist (2000) presents useful information regarding the performance of the fire department in manual fire fighting operations.

Fire type. If a fire starts where there is not much combustible material, it may go out by itself. This could be modelled as the probability of a fire spreading beyond the initial, limited phase.

Staff. If staff members detect a fire and have the appropriate equipment, they may succeed in extinguishing the fire before it grows to any significant size. The probability of staff members extinguishing a fire would be expected to depend on their training, the amount of fire fighting equipment they have access to, etc.

In looking at different fire protection “systems”, one realises that some of the conditional probabilities involved are very difficult to estimate. This is why one is unable to use a traditional Bayesian approach to estimate these probabilities as precise values. For some of the probabilities, one would simply not be able to settle on any precise value. Instead, one can use the concept of reliability-weighted expected utility (RWEU) explained in the previous chapter. In employing RWEU one estimates a set of plausible values for the different probabilities and consequences and then assigns a second-order probability distribution to this set. Since one assumes that utility and monetary value are the same thing, one has to calculate a set of values for the expected costs due to fire.

The calculation of one value for the expected costs due to fire involves summing the product of the probability and the monetary outcome for each consequence in the tree describing the possible fire scenarios. For example, consider the event tree shown in Figure 22. To calculate the expected consequence (which can be the expected costs if one uses costs as consequences), one first calculates the product of the correct conditional probability and the monetary consequence for each of the consequences shown in Figure 22 and then sums these products.

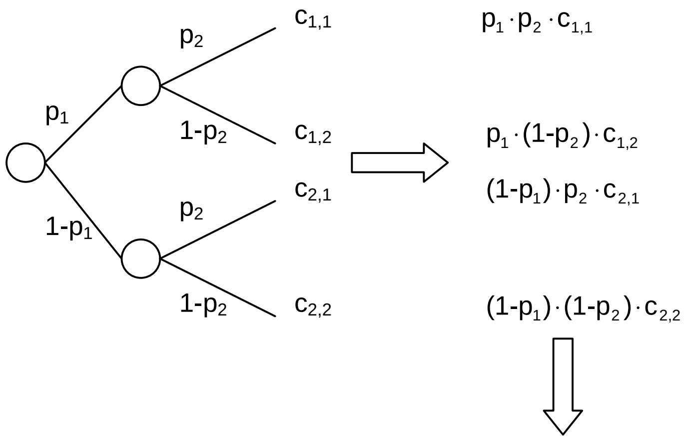

$$E(C)=p_1 \cdot p_2 \cdot c_{1,1} + p_1 \cdot (1-p_2) \cdot c_{1,2} + (1-p_1) \cdot p_2 \cdot c_{2,1} + (1-p_1) \cdot (1-p_2) \cdot c_{2,2}$$

Figure 22 Illustration of the calculation of the expected consequence from an event tree.

One can note in Figure 22 that if there is uncertainty concerning one or more of the probabilities (or consequences), resulting in a set of plausible values, then the effect will be a set of plausible expected consequences. This is how one deals with the uncertainty concerning the probabilities and consequences; one represents each of them by a set of values that are sufficiently spread to correspond to the uncertainty regarding one's belief about the parameter involved. In practice this is done by specifying a probability distribution for the parameters and running a Monte Carlo-simulation, resulting in a histogram showing which values for the expected costs due to fire are most probable. As was stated in section 2.3.1 the decision criterion states that the alternative with the highest reliability-weighted expected utility is the best alternative. However, if the decision is not robust (see section 2.3.1) then one should continue one's analysis in order to obtain more information regarding various of the uncertain parameters.

4.2.3. Annual expected cost due to fire

By multiplying the expected cost due to a single fire by the fire frequency per year, one can calculate the annual expected costs due to fire. Although this value can serve as an indicator of the risk present in the building, the thesis is primarily concerned with the *change* in the annual expected costs due to fire when some particular kind of fire protection system is installed. Thus, two different models for the expected costs due to fire are employed, one for the building with the fire protection system of interest installed and the other for the building without such a fire protection system, the *difference* between the two in expected annual costs being calculated. This difference in annual expected costs is useful since it can be employed directly in an investment appraisal, in which case the investment calculations will also result in a set of values.

5. Investment appraisal

This chapter concerns in particular presentation of the results of a decision analysis regarding possible fire protection measures. Since the persons actually making the decision to install a new fire protection system do not always have a background in fire engineering, it is important to consider how the results can be presented in as understandable a way as possible. It will be shown how an investment appraisal can be used for presenting the results arrived at. An investment appraisal is often suitable for this since it is a tool a company commonly makes use of, a decision maker thus usually being familiar with how to interpret the results. The investment appraisal method employed is the net present value method, also called the present worth method (see, Northcott 1995, for example).

5.1. Net present value method

When one invests in a fire protection system, a large part of the negative effects of the investment (such as its initial costs) become manifest in the initial stages of the system's lifetime. In contrast, one receives the positive effects of the investment (i.e. the reduction in risk) during the entire economic lifetime of the investment. For comparing the current negative effects with the positive effects that will accrue in the future, a cash-flow diagram can be employed. Figure 23 provides an example of a cash-flow diagram. As can be seen, an initial investment of 10.000 SEK is expected to yield a yearly positive cash flow of 3.000 SEK. In evaluating the investment, one should not simply look at the numbers in the cash-flow diagram and compare them directly. It is more appropriate to think of the investment problem as representing a decision analysis. This means that one needs to compare the investment in question with the other alternatives that are available. One alternative assumed to always be present is to invest the money at a predetermined level of return, such as in a bank account. Looking at the problem in this way emphasises the fact that 3.000 SEK received after five years is not the same as 3.000 SEK received immediately. If one received 3.000 SEK immediately one could have considerably more than this after five years later through having invested the money at a given level of return. Another reason to why 3.000 SEK obtained after a period of five years is not as good as 3.000 SEK received immediately is inflation, which results in a decrease in the purchasing power of money.

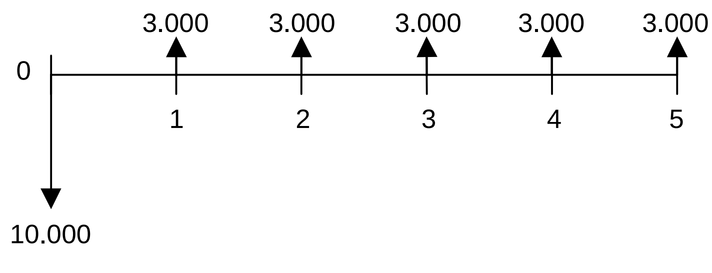

Figure 23 Example of a cash-flow diagram. An arrow pointing upwards designates a positive cash flow and an arrow pointing downwards a negative cash flow. Values are given in SEK.

Although there can be other sources of the problems just discussed, the important thing is that in making an investment appraisal one can deal with the problems involved by discounting future cash-flows to the point in time when the investment was made. Discounting means that one determines a present amount that yields a specified future sum, given a particular interest rate. This allows one to compare cash flows occurring at different times. In discounting a future cash flow, one needs to determine an interest rate that is "certain", i.e. the interest rate of the alternative one is comparing the investment with (such as investing in some other project). This interest rate is called the *discount rate*. When all the cash flows have been discounted, they are summed, the result being termed the net present value of the investment. If this value is positive, the investment is profitable; otherwise it is not.

The current value of a future cash flow can be calculated using equation [3.10], where K_0 is the cash flow discounted to year 0, K_n is the cash flow in year n and i is the discount rate.

$$K_0 = \frac{K_n}{(1+i)^n} \qquad [3.10]$$

For example, suppose one is interested in knowing whether an investment in a smoke extraction system is profitable. Assume that the initial cost is 10.000 SEK, the reduction in the annual expected costs due to fire is 3.000 SEK, and the annual maintenance costs 300 SEK. Assume too that the economic lifetime of the system is 30 years and that the discount rate is 15%. As can be seen in Table 12, in which only the first 10 years are shown, the net present value of the investment is 3551 SEK.

Table 12 Cash-flow table for a smoke extraction system. All values are given in SEK.

	Year										
	0	1	2	3	4	5	6	7	8	9	10
Initial investment	-10000										
Maintenance		-300	-300	-300	-300	-300	-300	-300	-300	-300	-300
Damage costs		3000	3000	3000	3000	3000	3000	3000	3000	3000	3000
Cash-flow surplus		2700	2700	2700	2700	2700	2700	2700	2700	2700	2700
Present worth	-10000	2348	2042	1775	1544	1342	1167	1015	883	768	667

Discount rate	15%
Net present value	**3551**

In Table 12 the line termed Damage costs is the *reduction* in annual expected costs due to fire, Cash-flow surplus is the sum of all cash-flows during the year in question and Present worth is the discounted Cash-flow surplus.

In the investment appraisal just presented, the fact that prices may change during the lifetime of the system is ignored. If one wants to consider price-changes, there are basically two ways of doing this: (1) estimating all the positive and negative cash-flows in terms of constant prices and using a discount rate that represents the opportunity costs of capital in the absence of inflation (real discount rate) or (2) estimating all future cash-flows in current or inflated prices and using a discount rate that includes an allowance for inflation (nominal discount rate).

Assume one selects method (1) for taking account of price changes and that one estimates the real annual increase in maintenance costs to be 2%, the real increase in the reduction of the expected annual cost due to fire to be 4% and that the real discount rate is 15%. This would result in the investment being considered to be more profitable since according to Table 13 the net present value would increases then to 6341 SEK.

Table 13 Cash-flow table for a smoke extraction system with price changes. All values are given in SEK.

	Year										
	0	1	2	3	4	5	6	7	8	9	10
Initial investment	-10000										
Maintenance		-306	-312	-318	-325	-331	-338	-345	-351	-359	-366
Damage costs		3120	3245	3375	3510	3650	3796	3948	4106	4270	4441
Cash-flow surplus		2814	2933	3056	3185	3319	3458	3603	3754	3911	4075
Present worth	-10000	2447	2218	2010	1821	1650	1495	1355	1227	1112	1007

Discount rate	15%
Net present value	**6341**

The general way of conducting an investment appraisal for a fire safety system is thus to estimate the annual reduction in the expected costs due to fire and to use that reduction in the investment appraisal in the manner illustrated in the examples presented in Table 12 and Table 13. Although it is possible, of course, to take account of inflation, taxes, the residual value of the investment, and the like, in the present thesis only this simple demonstration of how to utilise an investment appraisal in connection with fire protection investments will be provided. For a more comprehensive discussion of investment appraisals, see Northcott (1995), for example.

5.2. Uncertainties

Since as emphasised in previous chapters the reduction in the annual expected costs due to fire is uncertain, one cannot identify a specific value for a given parameter but needs instead to use a set of values. This means that the net present value as calculated in the investment appraisal is likewise represented by a set of values. When making an investment appraisal in practical terms it is probably best, *if* the decision to make the investment has been found to be robust, to only use *one* net present value. Thus, one first calculates the difference in terms of reliability-weighted expected utility (monetary value) between the alternative of making the investment under consideration and that of letting the building stay in its present condition. If when using the resulting difference the net present value of the investment is found to be positive, one needs then to determine whether the investment decision is robust and, if this is the case, one can then provide a recommendation concerning the investment, along with an estimate of its net present value.

5.3. Risk adjusted net present value

Traditionally, when one carries out an investment appraisal, one uses the expected monetary outcome as the basis for decision. Yet when a fire protection system is what one is dealing with there may be reason to consider basing one's decision on the expected utilities of the different alternatives and to not assume that one is risk-neutral in the sense discussed in chapter 2.

The reason for not employing risk-neutrality is that one might want larger accidents to be assigned greater weight in one's assessments. For example, one might want a loss of 1.000.000 SEK to count not as a thousand times the loss of 1.000 SEK but as more. This could be modelled by a risk averse utility function. If one decides to evaluate a decision using a risk averse utility function, one has to choose an appropriate value for risk tolerance (R), a matter discussed in chapter 2 (assuming that one uses the exponential utility function described there). The value of R could be estimated by asking top executives questions about their risk taking and deriving from their answers an approximate *R*-value. This may not be possible, however, since one may not have access to the top executives at the company.

Howard (1988) has investigated decision making in three companies and found an exponential utility function with an *R*-value amounting to approximately 6% of the companies' net sales to seem appropriate. Although the scope of Howard's investigation is limited, it provides a hint of the size of the risk tolerance here. As a first approximation, when one is unable to conduct an investigation of the risk taking propensities the company's top executives, it seems reasonable to use an *R*-value of 6% of the net sales.

If one uses a risk-averse utility function to calculate the expected utility of an alternative, one can also calculate the *expected disutility* by changing the sign of the losses in the exponential utility function presented in chapter 2. This results in the disutility function shown in equation [3.11]. Expected disutility appears easier to work with when fire protection problems are being considered since all the consequences one deals with are negative, i.e. their involving losses.

$$DU(x) = 1 - e^{x/R} \qquad [3.11]$$

In equation [3.11], *DU(x)* is the disutility associated with the negative monetary outcome (loss) *x* and *R* is the risk tolerance.

Cozzolino (1978) defines Risk Adjusted Cost (RAC) in terms of equation [3.12], where *E(DU(x))* is the expected disutility of a decision alternative, which in turn is equal to the disutility of a loss of a particular amount of money. This amount of money is the RAC.

$$DU(RAC) = E(DU(x)) \qquad [3.12]$$

In the present context, RAC can be interpreted as the monetary amount the decision maker considers to be equal in value to being faced with a particular situation involving uncertainty. It is the monetary amount that the decision maker would be willing to pay for having the risk removed completely. Consider, for example, the event tree shown in Figure 24, which depicts the uncertain outcomes of a fire.

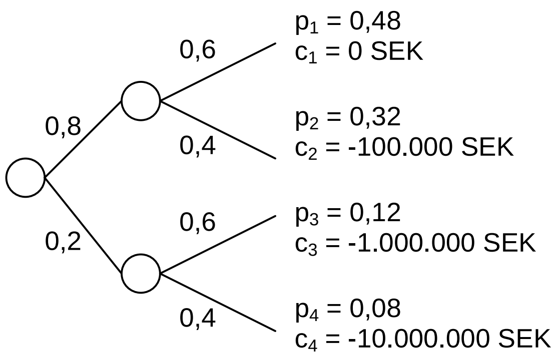

Figure 24 Illustration of an event tree describing different outcomes of a fire.

If one assumes that the decision maker is risk-neutral, the expected utility and the expected monetary outcome of a fire is -952.000 SEK. If, on the other hand, one assumes that the company owning the building to which the event tree applies has an annual net sales of

17.000.000 SEK and thus (according to Howard) a R-value of approximately 1.000.000 SEK the expected disutility (using the disutility function contained in equation [3.11]) of a fire is 1762. Translated back into monetary value by use of the disutility function, this represents approximately -7.500.000 SEK. This is the RAC according to equation [3.12].

In performing an investment appraisal for a new fire safety system in the building in question one starts by calculating the difference in expected utility between the alternative of keeping the building in its present condition and of investing in the fire safety system. Assume that the investment would bring about a reduction in the first probability (0,2 in Figure 24) to 0,1. This would yield new *p*-values of 0,54 (p_1), 0,36 (p_2), 0,06 (p_3) and 0,04 (p_4), which in turn would result in the expected monetary consequence being -496.000 SEK and the RAC being approximately -6.900.000 SEK. In the investment appraisal, one is interested in the *difference* in expected utility between the two alternatives. This means that 456.000 SEK (-496.000 -(-952.000)) should be used as an annual income (assuming the fire frequency is 1 fire per year) if the investment appraisal is to be conducted according to the method presented earlier in the chapter. If instead one uses the difference in RAC between the two alternatives as an annual income, the value becomes 600.000 SEK (-6.900.000-(-7.500.000)). Thus, using the RAC instead of the expected monetary value results in an increase of the net present value for the investment. The net present value obtained using RAC as an income could be termed the *risk-adjusted net present value*.

Using RAC instead of the expected monetary value emphasises the effect of the most serious accidents on the decision, or more formally, one assumes that the decision maker is risk-averse. A good way of conducting the analysis here is to assume that the decision maker is risk-neutral and to calculate the net present value of the investment accordingly. If the value obtained is positive, one considers the investment to be desirable, whereas if the value is negative one goes on to analyse the risk-adjusted net present value.

5.4. An investment appraisal for ABB Automation Products

ABB Automation Products is a company within the ABB group that develops and produces products that monitor, control and protect different types of processes in industrial plants and electric power plants. The turnover is approximately 2,4 billion SEK and the company has 1.400 employee in Västerås and in Malmö (Sweden).

The present investment appraisal deals with the investment in a sprinkler system for a building called building 358. In building 358, ABB Automation Products assembles circuit cards and automation products and produces force-measurement equipment. The activities in the building constitute a major part of the company’s total turnover and represents a very important segment of the ABB group, since they provide other companies within the group with circuit cards, for example.

The building is situated in an industrial area in Västerås. The building is approximately 55.000 m^2 in size and it is divided up into eleven different fire compartments. The nearest fire department is that in Västerås, which has 6 to 10 minutes driving time to reach the building. The building is equipped with a smoke detection system linked to the fire department. It is also equipped with a sprinkler system that covers the entire building, although this was not the case in the middle of the nineties. Since the current activities in the building are similar to what they were then, the present investment appraisal will be done for the building *without* a sprinkler system, so as to see whether an analysis carried out using this method in the middle of the nineties would have shown the sprinkler system to be a profitable investment.

Since the monetary losses to be used are losses that affect ABB Automation Products only, and not other companies within the group, and only direct and consequential losses will be addressed, the analysis is performed at level 2 (see section 4.1). The analysis being performed at level 2 means that many possible costs, such as costs of additional marketing campaigns to retake market shares, have not been addressed in the analysis. In addition, since the direct- and consequential losses included in the analysis are ones that will be compensated for by the insurance company, a subjective evaluation of the ratio of the costs used in the analysis to the uncompensated costs to ABB needs to be performed when the results are made use of.

To estimate the profitability of the sprinkler system, a model is created for estimating the reduction in the annual expected costs due to fire. The model is created in the manner described in chapter 4, as well as in Johansson (2000a). In the model, many of the probabilities and consequences are uncertain which means that use of the reliability-weighted expected utility model is called for. This model involves each uncertain probability or consequence being represented by a probability distribution that encompasses the values judged to be most plausible. This results, by use of a 5000 Monte Carlo simulation, in a set of values for the reduction in the annual expected costs due to fire. In Figure 25, these values are presented in the form of a histogram. The economic losses used in the analysis are presented in Appendix B.

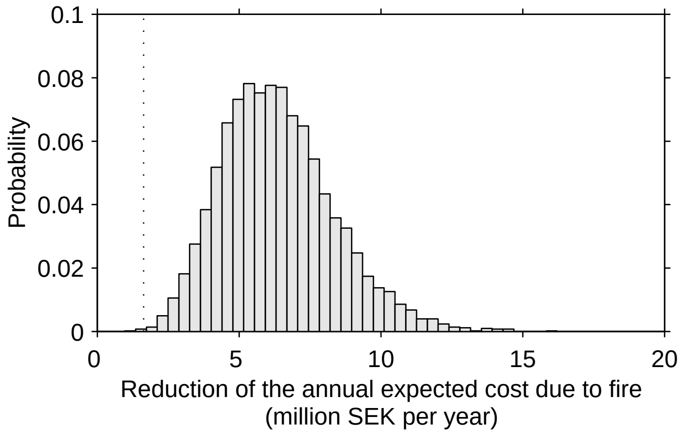

Figure 25 Reduction in the annual expected costs due to fire.

As can be seen in the figure, if the sprinkler system is installed the reduction in the annual expected costs due to fire is considerable. The mean value of the expected reduction in the annual expected costs due to fire is 6,4 million SEK per year, an approximately 90% confidence interval for this lying between 3,5 million SEK per year and 10,0 million SEK per year.

The costs for the sprinkler system amount to approximately 10 million SEK, the costs for maintenance being estimated to be some 0,1 million SEK per year. These are the only other economic aspects considered, except for the reduction in expected damage costs, that are taken account of in the analysis. The economic lifetime of the sprinkler system is assumed to be 40 years, and the discount rate to be 15%. No price changes are considered in the analysis.

An investment appraisal using the reliability-weighted expected utility (monetary cost)[7] is shown in Table 14, where the “Damage costs” are the reduction in the annual expected costs due to fire and the “Yearly surplus” consists of “Damage costs” minus “Maintenance”. The row entitled “Present value” is the present value of the “Yearly surplus” row. Only the first 10 years are presented in Table 14. All of the 40 years are shown in appendix C.

In Table 14 it can be seen that the investment has a net present value of **31 million SEK**.

Table 14 Investment appraisal for a sprinkler system in building 358 (ABB). Only the first ten years are presented here. All of the 40 years are presented in Appendix C. All values are in millions of SEK.

ABB	Year										
	0	1	2	3	4	5	6	7	8	9	10
Sprinkler costs	-10										
Maintenance		-0,1	-0,1	-0,1	-0,1	-0,1	-0,1	-0,1	-0,1	-0,1	-0,1
Damage cost		6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36
Yearly surplus		6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26
Present value	-10,0	5,44	4,73	4,12	3,58	3,11	2,71	2,35	2,05	1,78	1,55

Discount rate	15%
Net present value	31,1

In order to assess how robust the decision to install a sprinkler system is, one can run a series of Monte Carlo simulations to determine in how many of the simulations the net present value is positive. It turns out that the net present value is positive in all but one of the 5000 Monte Carlo simulations. This can be seen in Figure 25, where the dotted line represents the lower boundary for the reduction in the annual expected costs due to fire under conditions of the investment’s being profitable. Thus, the decision to install a sprinkler system can be deemed to be robust under the assumptions made in the model.

Figure 26 can be used to show the net present value for variations in the reduction in the annual expected costs due to fire. The dotted line represents the values chosen for the analysis in Table 14.

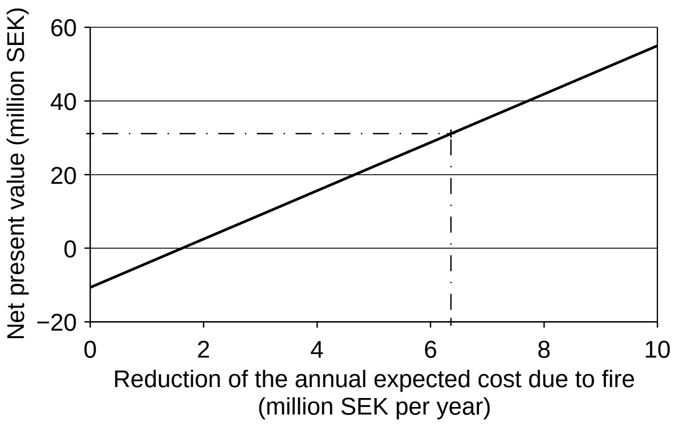

Figure 26 The net present value of the sprinkler system shown as a function of the reduction in the annual expected costs due to fire.

[7] It is assumed that the decision maker is risk neutral.

Figure 26 can also be used by a decision maker who wants to estimate the reduction in the uncompensated expected losses due to installation. If he/she considers the reduction in the uncompensated expected losses to be half as great as the reduction in expected losses used in the analysis, for example, the net present value of the investment is then approximately 10 million SEK, as Figure 26 indicates. One needs to bear in mind, however, that the lower one considers the uncompensated losses to be in comparison with the losses used in the analysis, the less robust the decision becomes. This means that once the decision maker has estimated the ratio of the uncompensated losses to the losses used in the analysis, it is best that he/she check the robustness of the decision again. No explicit estimation of the uncompensated losses in comparison with the losses used in the analysis has been made in the ABB case, but if one assumes that they are of the same magnitude as the costs used in the analysis the investment in a sprinkler system is profitable.

5.5. An investment appraisal for Avesta Sheffield

Avesta Sheffield, one of the world's leading suppliers of stainless steel, has 6.600 employees worldwide. During the financial year 1998/1999, the annual sales of the Avesta Sheffield groups was 5,8 billion SEK.

Avesta Sheffield has a cold-rolling mill in Nyby (Sweden) that produces approximately 160.000 tons (in 1998 and 1999) of cold rolled steel per year. This constitutes a major part of Avesta Sheffield's annual steel production of approximately 1.000.000 tons. The investment appraisal concerns the possible investment in a sprinkler system for the entire cold-rolling mill, which is approximately 15.000 m^2, in size.

The analysis here is conducted in the same way as the ABB analysis, that is, through estimating the reduction achieved by the sprinkler system in the annual expected costs due to fire. This reduction can be used then in an investment appraisal to determine the profitability of the investment. The losses associated with each fire scenario being analysed are the direct losses and the consequential losses for the entire Avesta Sheffield group. If a fire destroys the cold-rolling mill, the consequential losses for other facilities owned by Avesta Sheffield will be substantial. Their losses as well as the consequential losses in the cold-rolling mill itself need to be accounted for. In summarising the consequential losses for the group, one finds them to amount to approximately 1,1 billion SEK per year. A more comprehensive presentation of the costs associated with different degrees of fire spread can be found in Appendix D.

To illustrate the effects of the uncertainties in the probabilities and consequences used in the model of fire spread in the cold-rolling mill, 5000 Monte Carlo simulations were performed according to the method of reliability-weighted expected utility (see section 2.3.1). In the simulations, values were assigned to the uncertain parameters as probability distributions. The results (reduction in the annual expected costs due to fire) are shown in Figure 27. The mean value of the reduction in the annual expected costs due to fire is approximately 32 million SEK. An approximately 90% confidence interval lies between 21 million SEK per year and 46 million SEK per year.

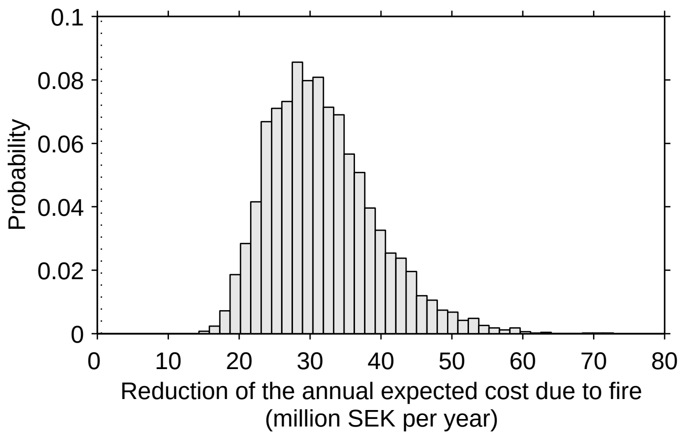

Figure 27 The reduction in the annual expected costs due to fire in the cold-rolling mill.

The cost of the sprinkler system is estimated to be about 2,5 million SEK and it is assumed here an economic lifetime of 40 years. The annual maintenance costs are estimated to be 50.000 SEK. The discount rate is assumed to be 20%. The net present value of the investment, based on the conditions just presented, is **156 million SEK**. The cash flow for the first ten years is shown in Table 15.

Table 15 Investment appraisal of a sprinkler system in a cold-rolling mill (Avesta Sheffield). Only the first ten years are presented here. All of the 40 years are presented in Appendix E. All values are in millions of SEK.

Avesta Sheffield	Year										
	0	1	2	3	4	5	6	7	8	9	10
Sprinkler costs	-2,5										
Maintenance		-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05
Damage cost		31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80
Yearly surplus		31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75
Present value	-2,5	26,46	22,05	18,37	15,31	12,76	10,63	8,86	7,38	6,15	5,13

Discount rate	20%
Net present value	156

Since the net present value of the investment is positive, the decision to install the sprinkler system is recommended, but because of the annual reduction in the damage cost being uncertain, one needs to check on the robustness of the decision. This can be done by examining at what reduction in the expected yearly damage costs the investment is no longer profitable and then comparing Figure 27 with this value.

The boundary between when the investment is profitable and when it is not is approximately the point at which the reduction in the annual expected costs due to fire is 0,6 million SEK per year (the dotted line in Figure 27). This means that the decision is robust, given the conditions at hand and the assumptions that have been made.

Figure 28 can be used to illustrate the net present value of the investment for different reductions in the annual expected costs due to fire. The dotted line in the figure represents the value of the reduction in the annual expected costs due to fire used in the analysis shown in Table 15.

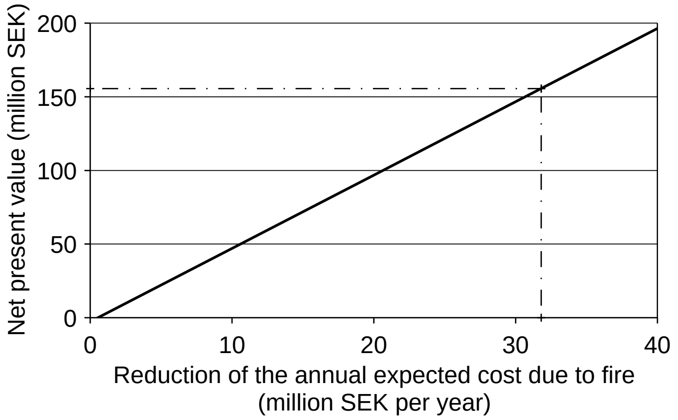

Figure 28 *Illustration of the net present value of the sprinkler system as a function of the reduction in the annual expected costs due to fire.*

In the present analysis, the reduction in the uncompensated losses was estimated to be of approximately the same magnitude as the reduction in losses used in the analysis[8]. Thus, the calculated net present value can be used directly as a basis for decision.

5.6. Comments on the investment appraisals

The results of the investment appraisals show investment in a sprinkler system to be profitable both for ABB and for Avesta Sheffield. One should bear in mind, however, (i) that the reduction in the expected costs due to fire is uncertain and (ii) that even if one had a precise number for the reduction involved there would still be uncertainty about whether the system would “pay off” during its lifetime. Point (ii) implies that even if one had a “perfect” analysis, meaning that one had calculated the “true” expected reduction in the annual costs due to fire, one could not be certain that a fire that the sprinkler system would have extinguished and that would otherwise have caused very serious damage would in fact have occurred.

In comparing the analysis of the ABB-building with that of the Avesta Sheffield-building, one notes the difference in net present value. In the ABB case, the net present value of the sprinkler system is 31 million SEK, whereas in the Avesta Sheffield case it is 156 million SEK. The large difference can be accounted for by the following factors:

- In the analysis of the ABB-building only losses that affect ABB Automations Products have been considered, not losses that affects the remaining ABB group. In the Avesta Sheffield case, losses for the entire group have been considered.

- General fire protection (fire rated walls, fire alarms etc.) is much better in the ABB-building. This implies that the marginal effect of the sprinkler system is greater in the Avesta Sheffield case.

[8] This estimation was done by personnel from Avesta Sheffield.

6. Summary and discussion

In order to make decisions on fire protection in buildings, one needs a way of evaluating different alternatives and of comparing them. In the present thesis, a method of doing this has been proposed.

It was suggested that to provide a basis for a decision a model of the possible outcomes of a fire be created. Although a method for creating such a model is proposed in the thesis, it is important to recognise the fact that every building is unique and that a model suitable in one building may not be suitable in another. It seems reasonable to assume nevertheless that the features described in chapter 4 be incorporated into such a model. Whatever structure the model has, it should enable one to estimate the probability of different consequences through the probability of different events, such as that the sprinkler system will extinguish the fire, and the like, being estimated. Since the thesis focuses on economic aspects of fire, the consequences are expressed as economic losses. This enables one to calculate the expected costs given that a fire has occurred. Multiplying the expected costs given a fire by the fire frequency during a year yields an estimate of the annual expected costs due to fire.

Using the expected annual cost due to fire as a basis for decision involves use of the decision rule found in traditional Bayesian decision theory, that of maximising the expected utility. In this case, it is assumed that utility and monetary outcome are equivalent. There is a problem in the use of traditional Bayesian decision theory in the present context, however, in that it assumes that the decision maker can assign a unique probability distribution here defined over the possible states of the world. This implies that each possible outcome of a fire, and each of the events leading to any particular outcome, be assigned a specific probability measure. In the present context this is very difficult to do since the information available concerning a specific probability is often very limited. This makes it difficult to assign precise values to each of the probabilities. The method proposed for use here instead is based on a modification of traditional Bayesian decision theory, one in which probabilities and consequences are assigned not fixed values but a *set* of plausible values. A plausible value should be considered here as a value that does not contradict the decision maker's knowledge. To capture the effect of inaccurate or imprecise information, one can then use some form of second-order measurement defined over the set of plausible values. In the present context the concept of second-order probabilities (see for example Goldsmith & Sahlin, 1982) is suggested for representing the decision maker's beliefs concerning the first-order probabilities. There are other methods too, however, for handling imprecise probabilities, Supersoft Decision Theory (SSD), presented here briefly, being one of these. An advantage of SSD is that the decision maker can use vague statements such as "the consequence c_1 is at least ten times as bad as consequence c_2" in the decision analysis. This means that not particularly much information is needed to analyse the problem since very broad intervals for the probabilities and very vague statements about the relation between the various consequences can be used. At the moment, SSD is somewhat cumbersome for such applications as those discussed in the present thesis. For now, therefore, the method of using second-order probabilities seems to be the most attractive approach.

Since traditional Bayesian decision theory is used here in modified form, one cannot use the decision rule of that theory but has to modify that as well. By allowing the different probabilities to be represented by a set of probability measures, one also ends up with a set of expected utility values for each alternative. If one uses the concept of second-order probabilities here to express how reliable the different probabilities are, one ends up with a probability distribution for the expected utility of an alternative. If one has two alternatives, one thus has two such probability distributions. Through calculating the difference in expected utility between the two alternatives, one can produce a probability distribution for the difference in expected utility and use this distribution in making the decision. The decision criterion suggested in the thesis is called the reliability-weighted expected utility criterion. Use of this criterion means, to start with, that if a specific probability value is uncertain one assigns a second-order probability distribution to the different probabilities that could be correct. One then calculates a weighted probability value using the reliabilities (second-order probabilities) of the different possible (first-order probabilities) values as weights. The probability values calculated in this fashion should be used finally together with the principle of maximising expected utility. This could be taken to mean, however, that if one has 100 possible combinations of probability values and in 49 of these combinations alternative 1 has the highest expected utility whereas in 51 the alternative 2 has the highest expected utility, alternative 2 should be the most preferred alternative. This would seems somewhat unreliable, on the other hand, since a very small change in the plausible probability values could have the effect of alternative 1 suddenly being the most preferred instead. Use of the term robust decision is suggested here, such that a decision is regarded as robust if the preferred alternative has the highest expected utility for *most* of the combinations of plausible probability values. Although no definition of how many "most " represents is given, it appears reasonable enough to state that if one alternative is found to be better in more than 95% of the combinations, the decision can be regarded as robust. In the practical cases presented in chapter 5, the result of both of the analyses were regarded as robust.

Sometimes, one wants to combine one's subjective estimates of the probabilities in the model with actual measurements or statistics. One might want, for example, to combine one's estimate of the probability that the staff will extinguish a fire in a particular factory with the information that the staff extinguished two out of four fires during the past year in that factory. To combine the subjective estimates with the objective information here, one can use Bayes theorem. This is what is called Bayesian updating, an approach dealt with in the thesis. Bayesian updating fits remarkably well with the concept of second-order probabilities since it provides a way of adjusting second-order probabilities in the light of new information. In the case studies presented in chapter 5, Bayesian updating was only applied to fire frequency, the fire frequency first being estimated as a prior distribution. The prior distribution expresses the decision maker's prior knowledge about what fire frequency values are most reliable. The prior distribution can then be updated using information concerning how many fires have occurred during a certain period.

In presenting the results of a decision analysis of a particular fire protection system, it is desirable to use a method that the decision maker is familiar with and that is consistent with his/her way of thinking. Accordingly, an investment appraisal appears to be a sensible way of presenting the results of the analysis. In the thesis two examples of how an investment appraisal can be performed are given. In both cases the method follows what is outlined in the thesis. First a model of the fire spread in the building is created on the basis of different events, such as those concerning the reliability of the sprinkler system, the staff, fire compartmentation, and the like. The model requires then that the probabilities of the events in the model that affect the outcome of a fire be estimated. These estimates are made in terms of the method of reliability-weighted expected utility. The results of both analyses were that the decision arrived at was robust, meaning that the alternative preferred in terms of the model was the same for most of the combinations of plausible values of the probabilities. Instead of presenting the entire distribution of plausible values for the expected utility, the mean value of the resulting probability distribution was employed. This value was used then in an investment appraisal, the results indicates the net present value of the investment in a sprinkler system in the ABB building to be **31 million SEK** and in the Avesta Sheffield building to be **156 million SEK**.

Although presenting the net present value of an investment represents a good way of describing its profitability, it could be wise to report not only this value but also to what extent possible the variation in the set of values employed for the reduction in the annual expected costs due to fire. It is suggested that an approximately 95% confidence interval be employed here.

The approach of using the reliability-weighted expected utility criterion together with the concept of robustness appears to be useful when difficult risk management decisions concerning fire protection are to be made. The method can help the decision maker analyse whether or not to make a certain investment, an analysis which can be carried within a reasonable period of time. Furthermore, through using the Bayesian updating technique it is possible to not only reduce the uncertainty present in a decision analysis but also to produce a risk assessment tool that can be used to monitor the risk in a building over time. The only thing needed for monitoring is a computer program based on the model and that one continually provide the program with new fire statistics relevant to the object of interest.

7. References

Ang, A.H-S. & Tang, W.H., 1975: *Probability Concepts in Engineering Planning and Design, Volume I – Basic Principles*, John Wiley & Sons, New York.

Apostolakis, G., Kaplan, S., Garrick, B.J. & Duphily, R.J., 1980: Data Specialization for Plant Specific Risk Studies, *Nuclear Engineering and Design*, 56, pp 321-329.

Apostolakis, G., 1988: Uncertain Data and Expert Opinions in the Assessment of the Unavailability of Suppression Systems, *Fire Technology*, Vol. 24, No. 2.

Apostolakis, G., 1978: Probability and Risk Assessment: The Subjectivistic Viewpoint and Some Suggestions, *Nuclear Safety*, Vol. 19, No. 3.

Apostolakis, G., 1986: Some Probabilistic Aspects of Fire Risk Analysis for Nuclear Power Plants, *Proc. 1st International Symposium on Fire Safety Science*, pp 1039-1046.

Benjamin J., R. & Cornell, C., A., 1970: *Probability, Statistics, and Decision for Civil Engineers*, McGraw-Hill Book Company.

Blom, G., 1989: *Sannolikhetsteori och statistikteori med tillämpningar,* Studentlitteratur, Lund. (Swedish)

Clemen, R.T., 1996: *Making Hard Decisions – An Introduction to Decision Analysis,* Second edition, Duxbury Press.

Cozzolino, J.M., 1978: A Method for the Evaluation of Retained Risk, *The Journal of Risk and Insurance*, Vol. XLV, No. 3, September 1978.

Danielsson, M., 1995: *Computing Best Choises Using Imprecise Information*, Stockholm University, Department of Computer and Systems Sciences, Lund.

Goldsmith, R.W. & Sahlin, N-E., 1983: The role of second-order probabilities in decision making, *Analysing and Aiding Decision Processes*, ed. by P.C. Humphreys, O. Svenson, and A. Vari, North-Holland, Amsterdam.

Gärdenfors, P. & Sahlin, N-,E., 1988: Introduction: Bayesian decision theory – foundations and problems, *Decision, probability, and utility,* ed. by Sahlin, N-E. & Gärdenfors, P., Cambridge University Press, Cambridge.

Hansson, S.O., 1991: *An Overview of Decision Theory*, SKN Report 41, Department of Philosophy, Uppsala University.

Herstein, I.N. & Milnor, J., 1953: An Axiomatic Approach to Measurable Utility, *Econometrica*, 21, 291-297.

Howard, R.A., 1988: Decision analysis: Practise and promise, *Management Science*, Vol 34, No.6.

Johansson, H. 2000a: *Osäkerhetshantering i riskanalyser avseende brandskydd*, Report 3113, Lund University, Department of Fire Safety Engineering, Lund.

Johansson, H. 2000b: *Beslusanalys och investeringskalkylering avseende brandskydd*, Report 3118, Lund University, Department of Fire Safety Engineering, Lund.

Kaplan, S., & Garrick, B., J., 1979: On the Use of a Bayesian Reasoning in Safety and Reliability Decisions – Three Examples, *Nuclear Technology*, Vol. 44, July.

Kaplan, S., & Garrick, B., J., 1981: On the Quantitative Definition of Risk, *Risk Analysis*, Vol. 1, No. 1.

Malmnäs, P.E., 1994: Axiomatic Justifications of the Utility Principle: a Formal Investigation, *Synthese*, 99, 233-249.

Malmnäs, P.E., 1995: *Methods of Evaluation in Supersoft Decision Theory*, Department of Philosophy, Stockholm University.

Malmnäs, P.E., 1999: *Foundations of Applicable Decision Theory*, Department of Philosophy, Stockholm University.

National Fire Protection Association, 1976: *Fire Protection Handbook*, 14th Edition, Boston.

Northcott, D., 1995: *Capital Investment Decision Making*, Dryden, London.

Oddie, G. & Milne, P., 1990: *Act and Value*, unpublished.

Pearl, J., 1988: *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann Publishers, Inc., San Francisco, California.

Ramsey, F.P., 1931: Truth and Probability, *The Foundations of Mathematics and Other Logical Essays*, ed. by R.B. Braithwaite, Routledge and Keagan Paul, pp. 156-198.

Räddningsverket, 1999: *Räddningstjänst i siffror*, I99-074/99. (In Swedish)

Savage, L.J. 1954: *The Foundations of Statistics*, John Wiley, New York.

Shpilberg, D. & Neufville, R., 1974: Best Choice of Fire Protection: An Airport Study, *Fire Technology*, No.10, pp.5-14.

Särdqvist, S., 2000: *Demand for Extinguishing Media in Manual Fire Fighting*, Report 1021, Lund University, Department of Fire Safety Engineering, Lund.

Watts, J.M., 1995: Fire Risk Ranking, SFPE Handbook of Fire Protection Engineering 2nd ed., ed. P.J. diNenno, NFPA.

Wikström, F., 2000: Kostnad-nyttaanalys av brandskyddsinvesteringar vid Stora Enso pulp, Skutskärs bruk, Report 5050, Lund University, Department of Fire Safety Engineering, Lund.

Appendix A: Supersoft Decision Theory

The trees representing the alternatives that are considered in the example in section 2.3.2 can be illustrated as follows:

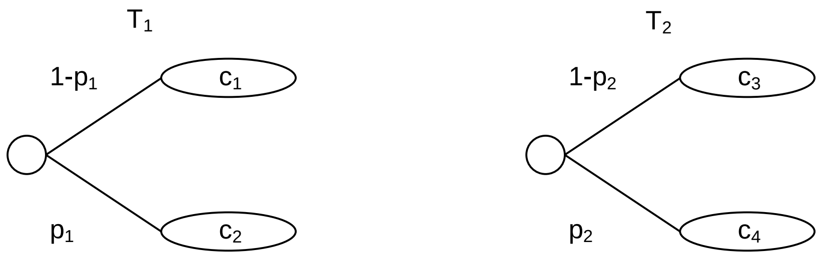

The first thing to do in evaluating the alternatives according to SSD is to represent the utilities associated with the different consequences with some suitable parameters. To do this the parameters *x* and *y* are introduced.

The resulting expressions when the utilities are expressed in the form of *x* and *y* follow:

$u(c_1) = 2x + y = 1$ **[1]**
$u(c_2) = x + y$
$u(c_3) = x$
$u(c_4) = 0$

The inequalities representing the probabilities given in the example in the text are the following:

$0{,}05 < p_1 < 0{,}15$
$0{,}01 < p_2 < 0{,}05$

The statement concerning the value (utility) distance between the alternatives (see Table 3 in section 2.3.2) yields

$y > 20x$ **[2]**

One can then express the expected utility of the different alternatives in the form of probabilities (p_1 and p_2) and utilities (using *x* and *y*).

$E(Alt1) = p_1 * u(c_3) + (1 - p_1) * u(c_1)$
$E(Alt1) = p_1 * x + (1 - p_1)$

$E(Alt2) = p_2 * u(c_4) + (1 - p_2) * u(c_2)$
$E(Alt2) = (1 - p_2) * (x + y)$
$E(Alt2) = x + y - p_2 * (x + y)$

In evaluating the alternatives one starts with the probability part of the decision frame and calculate *Max(E(Alt)), Min(E(Alt))* and *Mean(E(Alt))* for both alternatives. Since one only evaluates the probability part, the result will be linear forms of *x* and *y*.

Before evaluating *Max(E(Alt))*, *Min(E(Alt))* and *Mean(E(Alt))* one can analyse the constraints posed on *y* by statements regarding the values distances between the consequences, which yields

[1] and [2] $\Rightarrow$ 11y/10 > 1
y > 0,909

Thus, *y* has to be higher than 0,909.

One now concentrates on the evaluation of the differences between the alternatives in terms of *Min*, *Max* and *Mean*.

MAX
The maximum expected utility for each alternative, with respect to the probabilities expressed in the form of *x* and *y*, is given by the following equations[9]:

$\mathrm{Max}_P(E(Alt1)) = 0{,}05x + 0{,}95$
$\mathrm{Max}_P(E(Alt2)) = 0{,}99x + 0{,}99y$

Calculating the difference between the maximum expected utility of the two alternatives yields

$\mathrm{Max}_P(E(Alt1)) – \mathrm{Max}_P(E(Alt2)) = 0{,}05x + 0{,}95 - 0{,}99x - 0{,}99y = 0{,}95 – 0{,}94x – 0{,}99y$

Now, if the difference calculated is greater than 0, alternative 1 is the best according to this criterion.

If: $0 < 0{,}95 – 0{,}94x – 0{,}99y$ $\Rightarrow$ Alternative 1 is the best.

$0 < 0{,}95 – 0{,}94x – 0{,}99y$

[1] $\Rightarrow$ $0 < 0{,}48 – 0{,}52y$
$y < 0{,}923$

For alternative 1 to be the best, the *y* value has to be smaller than 0,923. This means that it is not possible to determine which alternative is the best in terms of the Max-criterion since the only limitations one has on *y* is that its value can not be smaller than 0,909 (see the top of the page). If *y* has a value of 0,910, for example, and thus *x* has a value of 0,045 (by the use of [1]) then alternative 1 is the best. On the other hand, if y has a value of 0,930 and thus *x* has a value of 0,035 then alternative 2 is the best.

Conclusion: It is not possible to determine which alternative is the best in terms of the Max-criterion.

[9] The subscript *P* indicates that the equations are the maximum expected utility with respect to the probabilities.

MIN
In the same way as for Max, the expressions for the minimum expected utility of the two alternatives (Min) are:

$Min_P(E(Alt1)) = 0{,}15x + 0{,}85$
$Min_P(E(Alt2)) = 0{,}95x + 0{,}95y$

This results in the following expression for the difference:

$Min_P(E(Alt1)) – Min_P(E(Alt2)) = 0{,}15x + 0{,}85 - 0{,}95x - 0{,}95y = 0{,}85 – 0{,}8x – 0{,}95y$

If: $0 < 0{,}85 – 0{,}8x – 0{,}95y \Rightarrow$ Alternative 1 is the best.

If the difference between the two alternatives just described is greater than 0, alternative 1 is best.

Using [1], one can conclude that *y* has to be smaller than 0,818 in order for alternative 1 to be the best according to this criterion.

$[1] \Rightarrow 0 < 0{,}45 – 0{,}55y$
$y < 0{,}818$

This, together with the previous result that, due to the statements made concerning the value distances between the alternatives, *y* had to be greater than 0,909, it can be concluded that *x* and *y cannot* possess values such that alternative 1 is the best.

Conclusion: Alternative 2 is the best when using the Min-criterion.

Mean
The mean value of the expected utilities associated with the probabilities here can be expressed in terms of the integrals below. Evaluating the integrals results in linear expressions of *x* and *y*.

$$\text{Mean}(E(Alt1)) = 10\int_{0,05}^{0,15} (P_1 * x + (1 - P_1))\, dP_1 = 10\left[\frac{P_1^2 x}{2} + P_1 - \frac{P_1^2}{2}\right]_{0,05}^{0,15} =$$
$$= 10\left(\frac{0{,}15^2 x}{2} + 0{,}15 - \frac{0{,}15^2}{2} - \left(\frac{0{,}05^2 x}{2} + 0{,}05 - \frac{0{,}05^2}{2}\right)\right) =$$
$$= 10\left(0{,}01x + 0{,}09\right) = 0{,}1x + 0{,}9$$
$$\text{Mean}(E(Alt2)) = 25\int_{0,01}^{0,05} x + y - P_2 * (x + y)\, dP_2 = 25\left[P_2(x+y) - \frac{P_2^2(x+y)}{2}\right]_{0,01}^{0,05} =$$
$$= 25\left(0{,}05(x+y) - \frac{0{,}05^2(x+y)}{2} - \left(0{,}01(x+y) - \frac{0{,}01^2(x+y)}{2}\right)\right) =$$
$$= 25\left(0{,}0388(x+y)\right) = 0{,}97(x+y)$$

The difference in the mean values can be expressed as an function of *x* and *y*:

Mean(E(Alt1))-Mean(E(Alt2)) = $0,1x + 0,9 - 0,97(x+y) = 0,9 - 0,87x - 0,97y$

Thus, if alternative 1 is best according to this criterion, the difference between the alternatives in terms of mean value must be positive.

If: $0 < 0,9 - 0,87x - 0,97y \Rightarrow$ Alternative 1 is the best.

Using the relation between *x* and *y* [1] yields

[1] $\Rightarrow 0 < 0,465 - 0,535y$
$y < 0,869$

Since it has already been concluded that *y* has to be greater than 0,909, one can conclude that alternative 2 is the best alternative according to this criterion.

Conclusion: Alternative 2 is the best according to the Mean-criterion.

Appendix B: Damage costs in the ABB building

In the analysis of the ABB Automation Products activities in building 358, the costs which destruction of the different fire compartments would result in were estimated. The estimates were made by ABB personnel, who in each case were asked to assign to the cost involved what was judged to be the most probable value, the maximum value and the minimum value. As can be seen in Table 16, the consequential and the direct losses that were estimated to occur in the case of a major fire in one of the fire compartments are rather large. If the entire building should be destroyed by a fire, the damage costs would be in the order of 1 billion SEK. Since this is the sum, of course, of the consequential and the direct losses, the uncompensated losses for ABB could be smaller. Nevertheless, one can conclude that a large fire in this building would cause serious loss to the company.

Table 16 Damage costs associated with the destruction of a fire compartment. Values in million SEK.

Area		**Minimum**	**Most probable**	**Maximum**
New PK workshop	Consequential loss	115	160	210
	Direct loss	144	160	176
A workshop	Consequential loss	120	190	271
	Direct loss	108	120	132
Storage area	Consequential loss	-	-	-
	Direct loss	4	5	6
ABB Training Center	Consequential loss	-	-	-
	Direct loss	10	13	16
EMC	Consequential loss	-	-	-
	Direct loss	16	18	20
PS workshop	Consequential loss	59	82	109
	Direct loss	120	150	180
Office area	Consequential loss	-	-	-
	Direct loss	36	40	44
Old PK workshop	Consequential loss	180	250	330
	Direct loss	225	250	275

One can see in Table 16 that the single area that can cause the most damage is the old PK workshop, destruction of which can cause losses of some 500 million SEK. This old PK workshop is an area in which circuit cards are assembled in 7 production lines. A similar activity takes place in the new PK workshop, but in that area there are only 4 production lines.

Figure 29 indicates the relative positions of the different fire compartments.

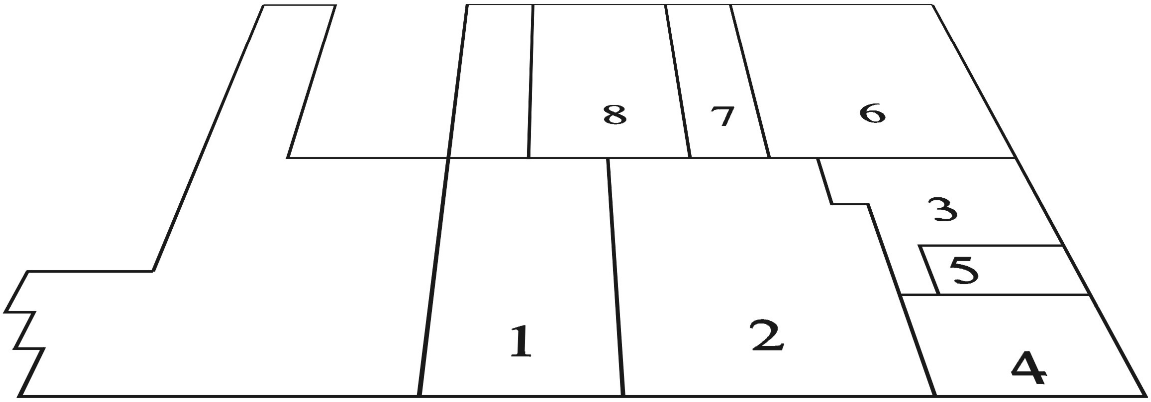

Figure 29 A sketch of the layout in building 358. The unmarked area is one that ABB Automation Products does not make use of.

Appendix C: Investment appraisal (ABB)

In this appendix each year which the investment appraisal of ABB concerns is presented. Note how the present value of the annual cash-flow decreases when cash-flows in the more distant future are considered. This indicates that it does not matter much if the economic lifetime of the system is assumed to be 40 or perhaps even 60 years, since the effect of the additional years involved would probably be only marginal.

	Year										
	0	1	2	3	4	5	6	7	8	9	10
Cost of sprinkler	-10										
Maintenance		-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10
Damage cost		6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36
Yearly surplus		6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26
Present value	-10,0	5,44	4,73	4,12	3,58	3,11	2,71	2,35	2,05	1,78	1,55
Discount rate	15%										
Net present value	31,1										

	Year										
	11	12	13	14	15	16	17	18	19	20	21
Maintenance	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10
Damage cost	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36
Yearly surplus	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26
Present value	1,35	1,17	1,02	0,88	0,77	0,67	0,58	0,51	0,44	0,38	0,33

	Year										
	22	23	24	25	26	27	28	29	30	31	32
Maintenance	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10
Damage cost	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36
Yearly surplus	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26
Present value	0,29	0,25	0,22	0,19	0,17	0,14	0,13	0,11	0,09	0,08	0,07

	Year							
	33	34	35	36	37	38	39	40
Maintenance	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10	-0,10
Damage cost	6,36	6,36	6,36	6,36	6,36	6,36	6,36	6,36
Yearly surplus	6,26	6,26	6,26	6,26	6,26	6,26	6,26	6,26
Present value	0,06	0,05	0,05	0,04	0,04	0,03	0,03	0,02

Appendix D: Damage costs in the Avesta Sheffield building

In the analysis of Avesta Sheffield's cold-rolling mill the costs resulting from the destruction of machinery were estimated. The column "Direct costs" presents the costs associated with the destruction of a particular machine and the column "Indirect costs" the costs due to interruption of business following the destruction of a machine. One should note, however, that if more than one machine is destroyed it is not possible to calculate the sum of the indirect costs by simply adding the indirect costs associated with the one machine to those associated with another. This is due to the destruction of some of the machines affecting production flow in the same way so that the indirect costs that result are the same, regardless of whether two machines are destroyed or only one. In contrast, in the case of direct costs one can simply add the direct costs of two or more machines to obtain the total direct costs.

Figure 30 indicates the relative positions of the different fire compartments.

Table 17 Direct and indirect costs due to the destruction of different machines in a cold-rolling mill. One can calculate the total direct costs of an area being destroyed by summing the direct costs associated with the individual machines. The indirect costs cannot be calculated in this way since many of the machines are part of one and the same production line. Values in million SEK.

	Direct costs			**Indirect costs**		
Machine	**Min**	**Most probable**	**Max**	**Min**	**Most probable**	**Max**
Area 1						
Smoothing roller	100	125	150	450	630	810
Cutter 1	32	40	48	225	315	405
Cutter 2	60	75	90	225	315	405
Area 2						
Cutter 3	80	90	100	450	630	810
Area 3						
Cold-rolling mill 1	162	180	198	90	126	162
Cold-rolling mill 2	225	250	275	360	504	648
Strip coiling machine	72	90	108	450	630	810
Area 4						
Production line 60						
Uncoiling capstan, weld, etc.	80	100	120	900	1260	1620
Cold-rolling mill	160	200	240	900	1260	1620
Oven and cooler	40	50	60	540	810	1080
Blaster	32	40	48	900	1260	1620
Pickling machine	40	50	60	900	1260	1620
Stretcher leveller	40	50	60	900	1260	1620
Cutter and coiling capstan	80	100	120	900	1260	1620
Other (filter, switch room etc.)	88	110	132	900	1260	1620
Production line 55						
Uncoiling capstan, weld, etc.	80	100	120	450	630	810
Oven and cooler	56	70	84	270	405	540
Pickling machine 1	56	70	84	450	630	810
Pickling machine 2 and dryer	56	70	84	450	630	810
Cutter and coiling capstan	80	100	120	450	630	810
Area 5						
Abrasive-belt grinder	120	150	180	180	252	324
Area 6, Engine room 1	56	70	84	90	126	162
Area 7, Machine shop 1	24	30	36	0	0	0
Area 8, Engine room 2	80	100	120	360	504	648
Area 9, Machine shop 1	24	30	36	0	0	0

Note: values are given in millions of Swedish crowns (kronor).

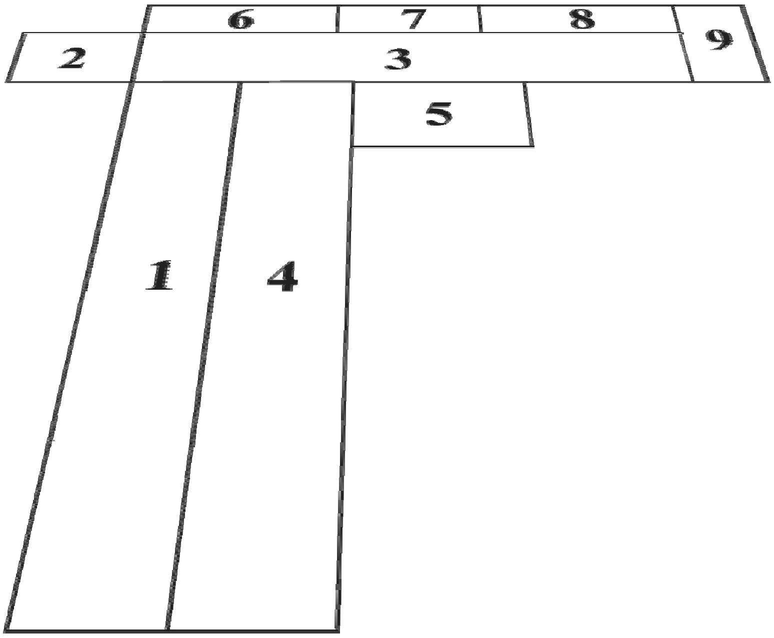

Figure 30 A sketch of the layout in the cold-rolling mill.

Appendix E: Investment appraisal (Avesta Sheffield)

In this appendix the complete cash-flow calculations concerning the Avesta Sheffield investment appraisal are presented.

	Year										
	0	1	2	3	4	5	6	7	8	9	10
Cost of sprinkler	-2,5										
Maintenance		-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05
Damage cost		31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80
Yearly surplus		31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75
Present value	-2,5	26,46	22,05	18,37	15,31	12,76	10,63	8,86	7,38	6,15	5,13
Discount rate	20%										
Net present value	155,6										

	Year										
	11	12	13	14	15	16	17	18	19	20	21
Maintenance	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05
Damage cost	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80
Yearly surplus	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75
Present value	4,27	3,56	2,97	2,47	2,06	1,72	1,43	1,19	0,99	0,83	0,69

	Year										
	22	23	24	25	26	27	28	29	30	31	32
Maintenance	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05
Damage cost	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80
Yearly surplus	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75
Present value	0,58	0,48	0,40	0,33	0,28	0,23	0,19	0,16	0,13	0,11	0,09

	Year							
	33	34	35	36	37	38	39	40
Maintenance	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05	-0,05
Damage cost	31,80	31,80	31,80	31,80	31,80	31,80	31,80	31,80
Yearly surplus	31,75	31,75	31,75	31,75	31,75	31,75	31,75	31,75
Present value	0,08	0,06	0,05	0,04	0,04	0,03	0,03	0,02

Appendix F: Using fire statistics to estimate the probability for different extents of fire spread in the Swedish industry

Introduction

In performing a quantitative risk analysis various probabilities needs to be estimated. These represent the uncertainties found in the problem in question. It is desirable to make such estimates in as objective a way as possible. In speaking of objective probabilities, what is usually meant are estimates based on a frequency interpretation of probability. The frequentistic probability of an event is the limiting value of the ratio of the number of trials on which the event occurred to the total number of trials performed. This is often regarded as the classical way of defining probability (see for example Benjamin & Cornell, 1970). However, it has one major disadvantage when used for quantitative risk analyses concerned with a *specific* factory or building. Many of the stochastic situations that one wishes to model have never occurred, there thus being no objective estimates of their probability. Instead, as both Apostolakis (1978,1986) and Apostolakis et al. (1980) have pointed out, subjective probabilities are needed here. A subjective probability is a representation of the belief of a particular person or group of persons concerning an event. One application of subjective probabilities is in Bayesian statistics, which is highly relevant to risk analysis, although subjective probabilities alone do not suffice, a combination of subjective and objective probabilities being needed. The problem discussed in the present paper is that of quantitative risk analysis concerned with the fire risk in a specific factory building, although the methodology suggested could be applied in quantitative risk analyses of other types as well.

Bayesian updating

In performing a quantitative risk analysis, use should be made, not simply of subjective judgement, but of information of three different types: general information, subjective judgements and objective measurements. In the present context, general information is information that has been gathered regarding a population of buildings to which the specific building of interest is assumed to belong. This building might, for example, be a steel mill and thus belong to the population of the industrial category Metalworking and machine industry. Information of this kind may involve estimates of the mean values of certain probabilities applying to this population. It can also consist of other kinds of general information about the type of buildings in question. Subjective judgements, in turn, are estimates made by a person (preferably an expert in the area involved). These subjective judgements depend both on the life experience of the person in question and on the information he or she has available. Part of the information could be general information about the mean value of some parameter or parameters of interest applying to some category of buildings in the industrial sector.

Suppose that the probability to be estimated is that a fire, if it occurred, in a particular building, for example in a steel mill, would spread outside the room of origin. Suppose too that the general information available included an estimate of the mean value for the industrial category Metalworking and machine industry that a fire in such a building would spread in this way. Assume that this general estimate is 0,25. The person who then made an estimate for the specific building of interest would use his knowledge of fire dynamics, of structural engineering and the like them, together with on-site observations, to adjust this general estimate to the conditions at hand. Suppose that the construction of this particular building, including the amount of combustible material inside, and other factors leads this person to consider the probability of such a fire spreading outside the room of origin in this particular building to be less than the mean value for steel industry as a whole, for example to be 0,2. This probability will be called the prior probability of the event, prior inasmuch as information the person receives or takes account of later may lead to this estimate being revised.

In making this probability estimate it can also be important that an estimate be made of how certain one is that one's estimate is correct. Consider an estimate of the probability that a football team will win a particular match. If the match is between two teams that one knows (call them team A and B), i.e. if one knows their players, their tactical strength and weaknesses, and the like, one's estimate of the probability that team A will win might be 0,5. If the match is between two teams one is completely unfamiliar with (call them teams C and D), one's estimate that team C will win may likewise be 0,5. Despite one's probability estimate being the same for team A as for team C, the estimates differ in how certain one is about them. In the first match, between two teams one know well, one may be rather certain that the estimate is a reasonable one. If one were asked whether the probability might not be 0,6 instead one would be likely to regard it as too high. In the second match, on the other hand, an estimate of 0,9, or just as well one at 0,1 might, as far as one could say, be a perfectly good estimate. The reason for one's estimating the probability to be 0,5 here is simply that one has no information about the teams. One can consider it wrong, under these circumstances, to favour either of the teams, 0,5 seeming to be a reasonable estimate for this reason.

Information on how certain one is about an estimate can be expressed in terms of a probability distribution, the median of this distribution representing the value one regards as the most likely one.

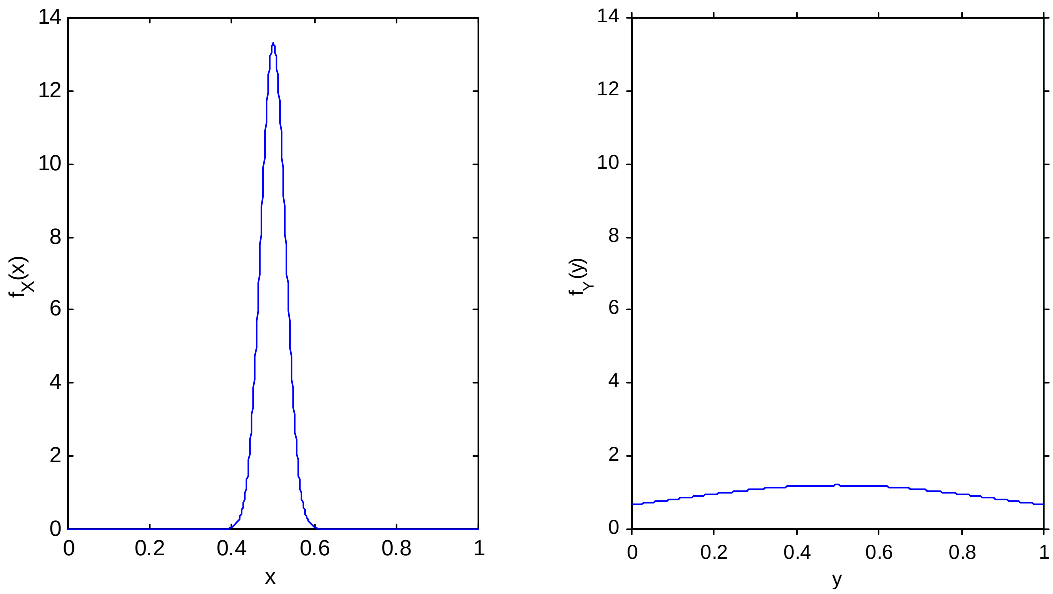

Figure 31 Distributions showing the uncertainty about the estimation of parameter X and Y.

In Figure 31 two probability distributions are shown, denoting how certain one is of one's estimate in the two football games just referred to. As can be seen, the distribution at the left (X) is very much concentrated around 0,5, meaning that one feels quite certain placing one's estimate there. The other distribution (Y), in contrast, is spread out fairly evenly over the entire range of probabilities, i.e. from 0 to 1. Nevertheless the mean values of the two distributions are the same, namely 0,5. Thus, the point estimate of the probability is the same in both cases, despite one's knowledge of the probability involved differing very markedly.

Expressing the value of a probability as a distribution is useful not only in communicating how certain one is about one's estimate but also for improving the estimates of probability that one makes in using Bayes theorem. If one assumes that the uncertainty about a probability is a discrete distribution, then Bayes theorem can be written as equation [3.13]. Discrete distributions are used because they are very simple to deal with in connection with Bayes theorem and because it is always possible to approximate a continuous probability distribution[10] by means of a discrete distribution if this is necessary or is desired. Bayes theorem represents a formal way of combining subjective judgements with objective measurements or evidence. In the case of the analysis of fire risk, Bayes theorem can be used not only for probability estimates of the spread of fire as such, but also for improving estimates of the probability of a variety of events, such as that a sprinkler system will extinguish the fire, that the staff will extinguish it, or whatever. The improvement can be achieved in probability estimates pertaining to the sprinkler system, for example, by first measuring the performance of the system in a number of fires (or even in just a single fire) and then using this information to revise one's prior probability estimate. The information thus gathered, which will be termed ε (evidence) could be, for example that the sprinkler system has been found to extinguish the fire in two out of three cases.

$$P(X = x_i \mid \varepsilon) = \frac{P(X = x_i)P(\varepsilon \mid X = x_i)}{\sum_{i=1}^{n} P(X = x_i)P(\varepsilon \mid X = x_i)} \qquad i = 1,2,...n \qquad [3.13]$$

In equation [3.13] $P(X=x_i)$ denotes the *prior* probability that $X = x_i$, where X could be the probability that the sprinkler system will extinguish the fire. The prior probability represents one's belief regarding the probability X *before* having obtained any added information. If one takes into account all possible values X might have, then $P(X=x_i)$ can be used to denote the prior distribution of X. This could take on the form of the distribution shown in Figure 33.

$P(\varepsilon|X=x_i)$ is a likelihood-function expressing the probability that the evidence ε would have been observed, *given* that the true value of the parameter X was x_i. For example, in the case in which X is the probability that the sprinkler system will extinguish the fire, $P(\varepsilon|X=x_i)$ is a binomial distribution. If one observes two fires and the sprinkler system extinguishes both of them, then $P(\varepsilon|X=x_i)$ is a binomial distribution expressing the probability for two separate fires to both be extinguished by the sprinkler system, given that the probability of the individual event of the fire being extinguished is x_i. The likelihood-function thus answers the question "What is the probability that the evidence ε would have been observed given that x_i was the correct parameter value?".

$P(X=x_i|\varepsilon)$ is the posterior distribution. This expresses one's belief regarding the parameter X after the evidence ε is known and thus after use has been made of Bayes theorem.

One difficulty one is faced with in analysing the fire risk in a particular factory, for example, is that in expressing one's belief about a particular parameter (such as a probability or a frequency), the information one has available is often of only a general character, applying to

[10] It is possible to use continuous distributions as well, but it is a bit more complicated unless so called conjugate distributions are used. In this paper only discrete distributions will be used, but more information about Bayesian updating of continuos and conjugate distributions can be found in Ang and Tang (1975).

factories of some given type but not necessarily to the factory building at hand. One way of solving this is to combine this general information with the judgement of experts, using this to create an adequate prior distribution for the parameter in question. This prior distribution can then be updated or improved using Bayes theorem in combination with observations, measurements or assessments made in the factory involved, so as to obtain the posterior distribution of the parameter. If one no longer wishes at this stage to express one's beliefs about the parameter as a distribution but prefers a point estimate (a single value) instead, one can use a so-called Bayesian estimator. If one employs a quadratic loss function here, the mean value of the posterior distribution is the Bayesian estimator (Ang and Tang, 1975).

An example serves to demonstrate the updating process which the use of Bayes theorem involves. Suppose that an event tree can be used to describe a fire in a particular building. At the first chance-node in this event tree, the probability that the staff will succeed in extinguishing the fire is indicated (p_{Staff} in Figure 32). One can assume that if the staff succeeds in extinguishing the fire (outcome A), no appreciable damage to the building and no injury to the people inside the building will occur, but if the staff is not able to extinguish the fire (outcome B) it will spread and could pose a serious threat to the building and to the people inside the building. How likely different fire scenarios are, given that the staff fails to extinguish the fire, is determined in the part of the event tree following outcome B. Since the events in the tree are assumed to be independent, this part of the tree has no effect on the estimation of p_{Staff}.

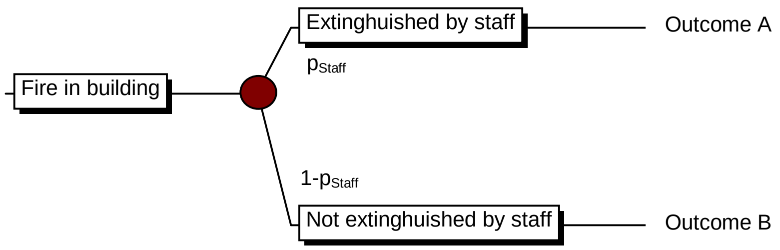

Figure 32 The first chance-node in an event tree used to describe different outcomes of a fire.

Assume that the person making the estimation believes that the staff would succeed in extinguishing the fire in 50% of all cases. The estimate the person makes of the parameter p_{Staff} is thus 0,5. If asked to express how certain he or she is of this estimate, the person could use a discrete distribution to represent this uncertainty. The distribution could be the one named “Prior” in Figure 33. In this example the likelihood-function has a binomial distribution, but if the parameter of interest had not been p_{Staff} but the frequency of fire, the likelihood-function would have been a Poisson distribution. These two distributions are probably the most useful ones in risk analysis connected with fire.

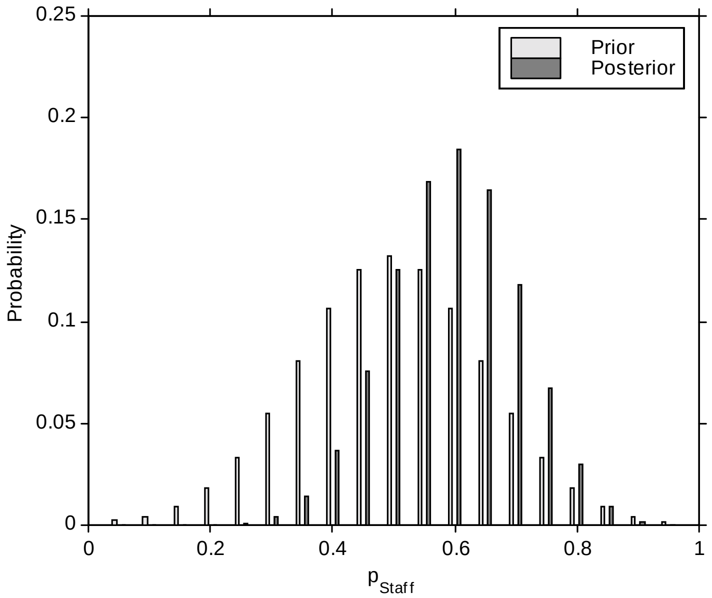

Figure 33 Discrete distributions representing the uncertainty of the parameter p_{Staff}.

In indicating the prior distribution here, the expert has expressed his or her belief concerning the parameter p_{Staff}, but has done so before knowing anything about how many times a fire has been extinguished by the staff and how many fires have occurred. By studying fire statistics for the building of interest the expert can obtain new information regarding the number of fires the staff has extinguished, seven fires out of ten for example. With the help of Bayes theorem, the expert can use this information to update his or her prior estimate of the parameter. If Bayes theorem is applied to the discrete distribution termed "Prior" in Figure 33, the resulting posterior distribution will be like the one termed "Posterior". Since the mean value of the posterior distribution is 0,59, this is the new estimate of the parameter p_{Staff}. One might have thought that the probability estimate would have been 0,7 (since seven out of ten fires were extinguished), but this presumes that no information on the parameter was available prior to the measurement. Here however, the expert had some knowledge of the parameter, namely the prior distribution, and since this distribution had a mean value of 0,5 it was natural that the posterior estimate lie somewhere between 0,5 and 0,7.

This example shows how subjective information can be improved using new statistical information pertaining to the building of interest, together with Bayes theorem. Additional examples of Bayesian updating can be found in Kaplan & Garrick (1979) where, among other things, the frequency of radioactivity release during the railroad transport of spent nuclear fuel is updated using Bayes theorem. In a study of the unreliability, or uncertain availability, of suppression systems in nuclear power plants, Apostolakis (1988) shows how information concerning nuclear power plants and indirect information obtained from tests and from non-nuclear applications can be combined using Bayes theorem.

In order to create good prior estimates information concerning the branch of industry involved can be utilized. The rest of this paper is concerned with the derivation of general probabilities for different extents of fire spread in a branch of industry which is applicable for creating a prior estimate. Since the focus is this branch of industry rather on a *specific* building within this branch, one can use here the classical frequency definition of probability.

Swedish fire statistics

In Sweden, each fire that is attended to by a local fire department is reported to the Swedish Rescue Services Agency. Before the statistics obtained are published there, the information on which it is based is reviewed to determine whether any of the information is inconsistent or appears questionable. If such is the case, the fire department responsible is asked whether it wishes to change anything or to add information. When all such questions have been dealt with, the statistics are published in an annual report.

These fire statistics can be used to determine the approximate extent to which each fire spread. This information can be used then to estimate probabilities for different extents of fire spread in a category of industries given that a fire has occurred.

Characterisation of a fire

To enable estimates to be made of the probabilities for different degrees of spread of a fire, some sort of model to represent the spread of a fire is needed. The model to be used here consists of an event tree describing different fire scenarios, given that a fire has occurred. The creation of the model involves a trade-off between accuracy and level of detail. A model with too many possible fire scenarios is problematical due to there being too few fires in the statistical material for meaningful estimates to be obtained. At the same time, a model with an insufficient number of possible fire scenarios tends to have little practical usefulness. The model used in the present paper is shown in Figure 34.

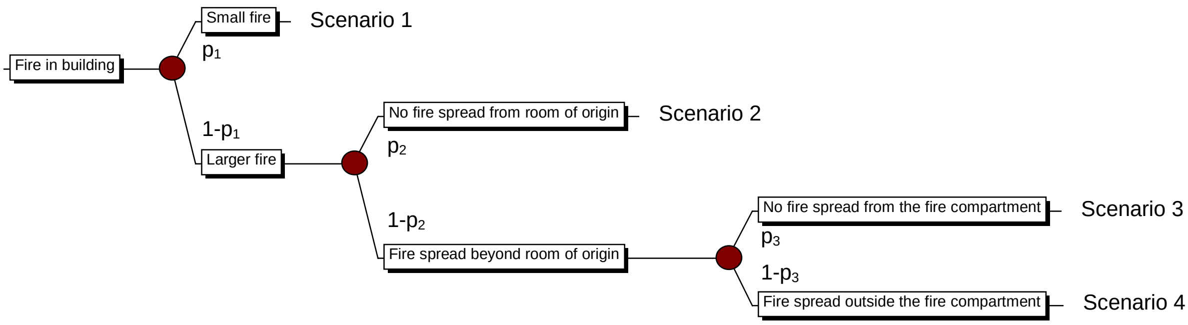

Figure 34 Description of the model used to characterise a fire.

The model is rather general, although it is detailed in the sense that it can distinguish between fires that pose no thereat to the room of origin, those that are contained in the room of origin, those that spread beyond the room of origin, and fires that spread beyond the fire compartment of origin.

The circles in Figure 34 are chance-nodes representing a stochastic situation. If a fire breaks out, there is a certain probability (p_1) that it will be limited to the object in which it started (small fire) a certain probability that it will grow but be limited to the room of origin, and so on.

On the basis of fire statistics, one can estimate the number of fires that resulted in only a small fire, those that were somewhat larger but were restricted to the room of origin, those that spread further but were extinguished in the fire compartment of origin, and fires that spread beyond the fire compartment of origin. These different extents of fire spread will be termed scenarios 1,2,3 and 4. The results for buildings with sprinklers and without in buildings of different industrial categories can be seen in Table 10.

Table 18 Extent of fire spread in buildings of different industrial categories in Sweden in 1996, 1997 and 1998.

	Scenario 1	Scenario 2	Scenario 3	Scenario 4	Total number of fires
Buildings without sprinklers					
Metalworking and machine industry	425	357	31	39	**852**
Chemical industry	124	106	3	18	**251**
Food manufacturing industry	91	81	4	12	**188**
Textile industry	26	20	2	6	**54**
Warehouses	39	79	5	35	**158**
Forest-product industry	220	283	24	78	**605**
Other branches of manufacturing	276	241	8	36	**561**
Repair shops	50	113	17	67	**247**
Buildings with sprinklers					
Metalworking and machine industry	94	40	3	2	**139**
Chemical industry	24	22	1	5	**52**
Food manufacturing industry	18	20	1	1	**40**
Textile industry	7	11	0	1	**19**
Forest-product industry	94	82	7	27	**210**
Other branches of manufacturing	115	59	12	7	**193**

The information shown in Table 10 can be used to estimate the parameters p_1, p_2 and p_3. These estimates provide information useful to an expert in creating a prior probability distribution for a specific building belonging to one of the building categories listed in Table 10. Note that the parameter values pertaining to a *specific* building (i.e. $p_{1,1}$, $p_{1,2}$…etc.) that belongs to the category in question can be very different from the mean value that applies to the category as a whole. It is also important to realise that the estimates of p_1, p_2 and p_3 are weighted mean values. They are weighted, in the case of p_1, with respect to fire frequency, in the case of p_2, with respect to fire frequency and p_1, and in the case of p_3, with respect to fire frequency, to p_1 and to p_2.

The weighting can be illustrated by the following: Consider an industrial category that consists of 100 buildings altogether. Assume that 50 of these buildings have a $p_{1,i}$ value (i is added to show that it is a parameter in the ith building) of 1 and the rest have a value of 0. The mean value of the category is thus 0,5. Assume in addition, that the first 50 buildings (those with a $p_{1,i}$ value of 1) likewise have a $p_{2,i}$ value of 1 and the rest a $p_{2,i}$ value of 0. The mean value of $p_{2,i}$ in the category is 0,5. When fire statistics for this industrial category is analysed, the estimate of p_1 will surely be close to 0,5 (assuming that the fire frequency is the same in all the buildings). This, however, is the same as the mean value of $p_{1,i}$. The estimate of p_2, however, will be close to 0,25, which is not the same as the mean value of $p_{2,i}$. The reason for this latter is that all fires that occurred in the first 50 buildings were small ones and in terms of Figure 34 were classified as belonging to Scenario 1 (since $p_{1,i} = 1$), whereas all fires in the other 50 buildings were classified as either belonging to Scenario 3 or Scenario 4. Although it might not be very probable in a real industrial category for 50% of the buildings to have a p_1 value of 1, the example nevertheless shows why the estimates of p_1, p_2, and p_3 must be considered as weighted mean values.

The weighting does not mean that the estimates are less useful for creating a prior-estimate pertaining to a specific building. On the contrary, this is often the kind of information one wants in creating an prior-distribution of a parameter, since it provides an answer to, for example, the question “What is the probability that a fire in building belonging to specific category will develop in accordance with scenario 2, *given* that it does *not* develop in accordance to scenario 1?”.

Estimation of parameters

To be able to estimate the different parameters in the stochastic model of a fire scenario (p_1, p_2, p_3) a Binomial distribution is required. This is a discrete distribution pertaining to an uncertain situation involving only two outcomes. One can see that in the model used to represent a fire scenario here (Figure 34) each chance-node represents an uncertain situation where there are only two outcomes. Thus, the Binomial distribution can be used to estimate the different parameters related to a given chance-node. The Binomial distribution can be written as in equation [3.14], where n is the total number of trials, k the number of trials in which the event of interest occurred, p is the probability that the event will occur on a given trial, q stands for $(1-p)$ and $p_X(k)$ is the probability that the event of interest occurred k times in n trials $(X=k)$.

$$p_X(k) = \binom{n}{k} p^k q^{n-k} \qquad (k=0,1,\ldots,n) \qquad [3.14]$$

To illustrate the Binomial distribution in this example, consider Figure 35, in which the first chance-node is shown along with the parameter p_1, denoting the probability that a fire will be restricted to the object of origin (small fire). X is the total number of fires in buildings of the industrial category in question in which the fire did not spread beyond the object of origin, and n is the total number of fires that occurred in buildings of this particular industrial category.

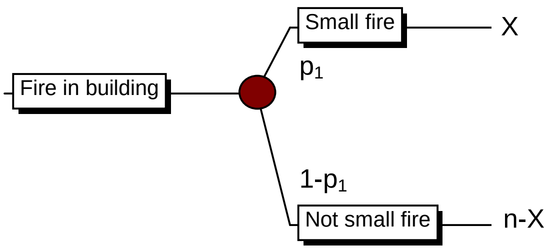

Figure 35 *Description of the first chance-node and the corresponding probability that a fire will be small* (p1)*, along with the number of fires that actually did remain small* (X) *and the total number of fires occurring* (n)*.*

In the case shown in Figure 35, X is binomially distributed, where $X \in Bin(n, p_1)$. For each branch of industry shown in Table 10, X and n are known. In producing an estimate of p_1, together with a confidence interval, the binomially distributed variable X can be approximated by a normal distribution (see for example Ang & Tang, 1975). Since this approximation is only valid if $np(1-p)$ is large (Blom, 1989, suggests larger that 10), one needs to check that equation [3.15] holds for each of the parameters within the branch of industry involved.

$$n \cdot p \cdot (1-p) > 10 \qquad [3.15]$$

If equation [3.15] holds, then the binomially distributed variable X can be approximated by the normal distribution in accordance with equation [3.16].

$$X \in N(np, \sqrt{np(1-p)}) \qquad [3.16]$$

The Maximum Likelihood estimate and the Least Squares estimate of the parameter p are found by dividing the number of trials in which the event of interest (e.g. a small fire) has occurred by the total number of trials, $p^* = x/n$ (* is used to denote p^* being an estimate of the parameter p). Since p_1, p_2 and p_3 are conditional probabilities, it is important to adjust the total number of trials (n) accordingly in producing an estimate of p_2, for example. In estimating p_1, the n employed is that shown in the column "Total number of fires" in Table 10, but in estimating p_2, n is the total number of fires *minus* all small fires (scenario 1).

Along with the estimation of the parameters, a 95% confidence interval needs to be obtained. Since $X \in Bin(n, p)$ can be approximated by equation [3.16], the *relative* frequency of positive outcomes at the chance-nodes in the model can be expressed in terms of equation [3.17].

$$X / n \in N(p, \sqrt{p(1-p)/n}) \quad [3.17]$$

Estimation of the probability (p^*) involves observing X/n. Since a 95% confidence interval is to be created for p, the standard deviation of p ($=X/n$) needs to also be calculated. The standard deviation of X/n can be estimated using equation [3.18].

$$d = \sqrt{p^*(1-p^*)/n} \quad [3.18]$$

The confidence interval for p is obtained then using equation [3.19], where $\lambda_{\alpha/2}$ is the $\alpha/2$-quantile in the standardized normal distribution.

$$I_P = p^* \pm \lambda_{\alpha/2} \cdot \sqrt{p^*(1-p^*)/n} \quad [3.19]$$

Results

Table 19 presents the results of those calculations. In the cases in which no confidence interval is given, equation [3.15] does not hold due to the number of fires having been too small for a confidence interval to be obtained. If no estimate is provided at all, the number of fires available for the estimate is less than 10.

Table 19 Summary of the probability estimations and 95% confidence intervals for the different industrial categories.

	$I_{p1, min}$	p_1	$I_{p1, max}$	$I_{p2, min}$	p_2	$I_{p2, max}$	$I_{p3, min}$	p_3	$I_{p3, max}$
Buildings without sprinklers									
Metalworking and machine industry	0,47	**0,50**	0,53	0,80	**0,84**	0,87	0,33	**0,44**	0,56
Chemical industry	0,43	**0,49**	0,56	0,77	**0,83**	0,90	-	**0,14**	-
Food manufacturing industry	0,41	**0,48**	0,56	0,76	**0,84**	0,91	-	**0,25**	-
Textile industry	0,35	**0,48**	0,61	-	**0,71**	-	-	**-**	-
Warehouses	0,18	**0,25**	0,31	0,58	**0,66**	0,75	-	**0,13**	-
Forest-product industry	0,33	**0,36**	0,40	0,69	**0,74**	0,78	0,15	**0,24**	0,32
Other branches of manufacturing	0,45	**0,49**	0,53	0,80	**0,85**	0,89	-	**0,18**	-
Repair shops	0,15	**0,20**	0,25	0,50	**0,57**	0,64	0,12	**0,20**	0,29
Buildings with sprinklers									
Metalworking and machine industry	0,60	**0,68**	0,75	-	**0,89**	-	-	**-**	-
Chemical industry	0,33	**0,46**	0,60	-	**0,79**	-	-	**-**	-
Food manufacturing industry	-	**0,45**	-	-	**0,91**	-	-	**-**	-
Textile industry	-	**0,37**	-	-	**0,92**	-	-	**-**	-
Forest-product industry	0,38	**0,45**	0,51	0,62	**0,71**	0,79	-	**0,21**	-
Other branches of manufacturing	0,53	**0,60**	0,67	0,66	**0,76**	0,85	-	**0,63**	-

The results (for p_1 and p_2) for the industrial categories without sprinklers are presented in graphical form in Figure 36 and Figure 37, where it can be seen that four of the categories here are very similar (Metalworking and machine industry, Chemical industry, Food manufacturing industry and Warehouses). The other groups (Forest-product industry, Other branches of manufacturing and Repair shops) appear to have lower values for the different probabilities.

As evident in Table 19, there are a considerable number of estimates in which, although all the estimates are uncertain, the confidence interval cannot be given, since the number of fires that occur are too few in number. Nevertheless, the results can be useful when similar parameters are to be estimated in a specific case and the quality of the results will surely improve, as fire statistics over a greater number of years become available.

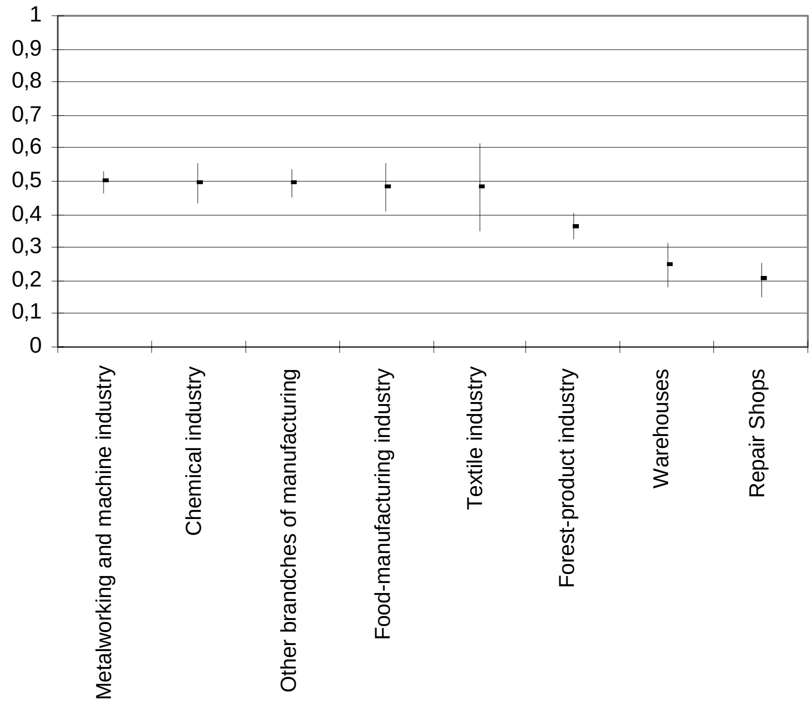

Figure 36 Graphical representation of the estimations of the probability p_1 and their 95% confidence interval.

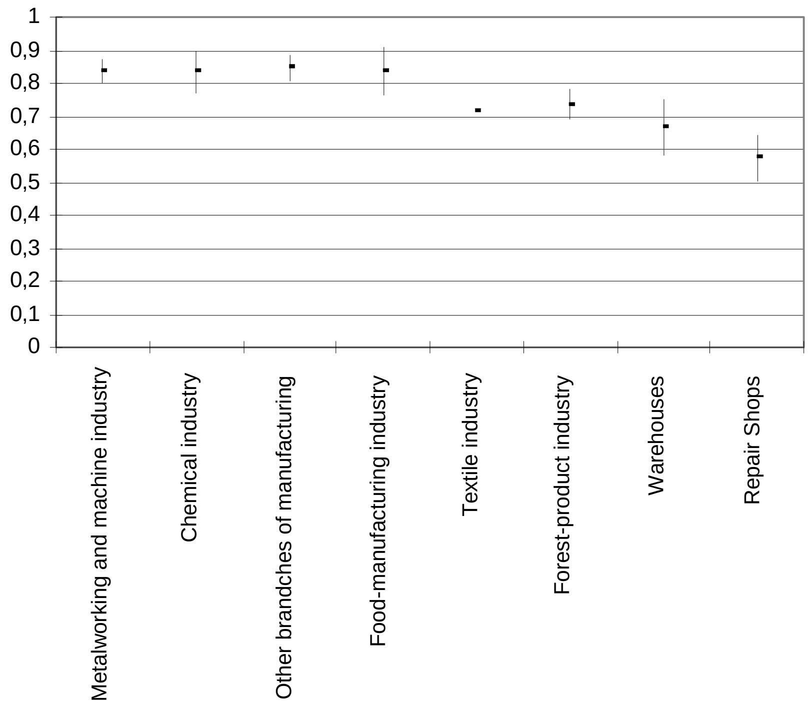

Figure 37 *Graphical representation of the estimations of the probability p_2 and their 95% confidence interval.*

Discussion

The possibility that not all fires that occurred were reported to the respective fire departments should be recognised. Since small industrial fires are surely more likely to not be reported than large fires are, it is reasonable to assume that estimates of the parameter p_1 are lower than they should be. Regarding estimates of the parameters p_2 and p_3 it would seem unlikely that a fire that had spread so as to involve the entire room of origin would not be brought to the attention of the local fire department. Thus, estimates of p_2 and p_3 should be more accurate than those of p_1.

Since the unknown numbers of small fires could affect estimates of the parameter p_1, it would be of interest to investigate how a moderate increase in the number of small fires would affect estimates of p_1. Figure 38 shows changes in the estimates of p_1 obtained for buildings without sprinklers as the number of small fires (scenario 1) increases. The maximum increase in small fires involved is 50%, which appears to result in an increase in p_1 of approximately 0,1, this varying somewhat with the industrial category involved. The results indicate that the estimates remain fairly stable with the occurrence of a reasonable increase in the total number of small fires through unreported fires being accounted for.

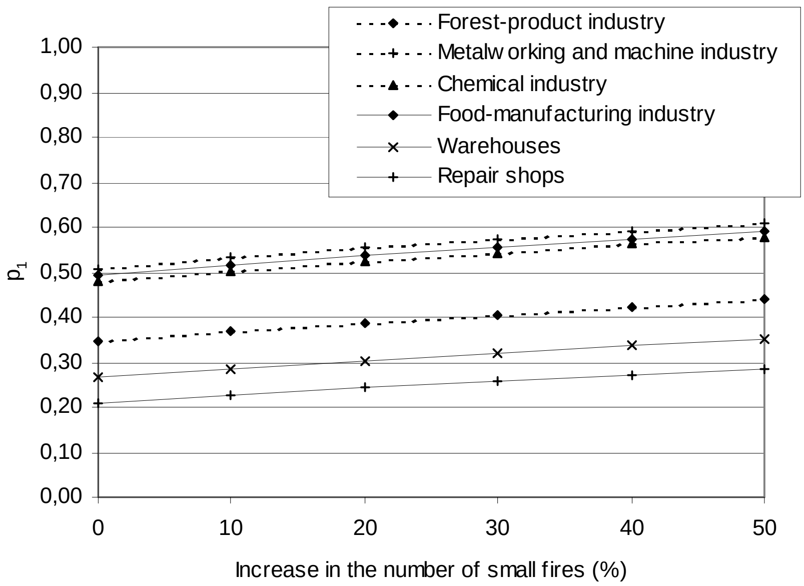

Figure 38 *Estimates of* p_1 *as a function of an increase in the number of small fires.*

Conclusions

The estimates of the probabilities for different extents of fire spread presented in this paper could be used in quantitative fire risk analysis as a help in assessing the probabilities of different fire scenarios, given that a fire has occurred. One must be aware of the fact, however, that the estimates of probabilities that are presented are estimates of the *mean values* for an industrial category. There can be and probably are large variations within such a category. The results presented here can be suitable for use in helping to create prior distributions for the parameters interest, distributions that can then be updated using Bayesian methods. The expert who creates such a prior distribution combines publicly available information concerning the category of buildings in question (Table 19) with his or her own knowledge of such buildings and the information he or she gathers concerning the building of interest. This prior distribution can be improved in the course of time by further information about this particular building being added so as to create a new posterior estimate, which can then be used in the final risk analysis.

Appendix G: Discussion of the use of second-order probabilities in decision making regarding fire protection

Pearl (1988) presents a concrete and illustrative example of how to express uncertainty regarding a probability. The example is about a coin which has been minted in the basement of a notoriously unscrupulous gambler.

Let E be the statement "The coin is about to turn up heads". If asked to estimate $P(E)$, one would probably have difficulties in making an exact estimate. More precisely, being uncertain about $P(E)$ means that one is aware of a set of contingencies such that, if one of them turned out to be true, this would substantially alter one's estimate of the probability. The contingencies might be the following, for example:

C_1 = "The coin is fair ($P(E)$ = 0,5)."
C_2 = "The coin is loaded for heads ($P(E)$ = 0,6)."
C_3 = "The coin is loaded for tails ($P(E)$ = 0,4)."

If one estimates the probability that the gambler has tampered with the coin to be 0,2 and, if he/she has done so, one considers the probability that he/she has loaded for tails to be equal to the probability that he/she has loaded for heads, the estimates one would make of the likelihood of the contingencies are as follows:

$P(C_1) = 0,8$
$P(C_2) = 0,1$
$P(C_3) = 0,1$

The situation at hand can be depicted in terms of the causal model shown in Figure 39 (from Pearl, 1988).

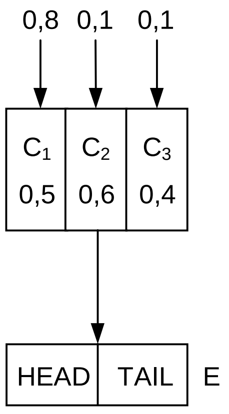

Figure 39 Causal model for the coin example.

Figure 39 should be interpreted as follows: the probabilities at the top of the figure are the person's subjective probabilities (estimated exactly) regarding which of the three contingencies is true. The next set of probabilities (0,5, 0,6 and 0,4) represents the respective probability that the coin is about to turn up heads. Thus, depending on which of the three contingencies is true, the probability that the coin will turn up heads differs. This is represented by the arrow pointing towards the boxes in which "HEAD" and "TAIL", respectively, appear.

Since *P(E)* depends on which of the contingencies is true, it is possible to create a *belief distribution* for *P(E)*. The belief distribution is shown in Figure 40, where $P(E|C_i)$ is the probability that the coin will turn up heads *given* that a particular contingency has turned out to be true. $P(C_i)$ is the probability that contingency *i* will turn out to be true.

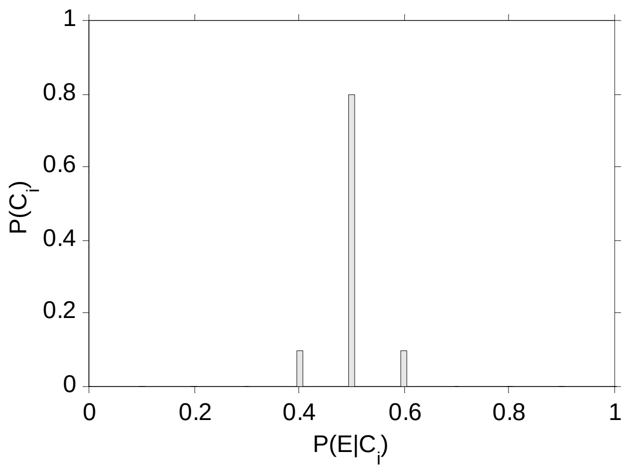

Figure 40 Belief distribution for the probability that the coin will end up heads.

This belief distribution is the same as a second-order probability distribution used in this thesis for *P(E)*. The difference between the second-order distribution and the belief distribution is that the belief distribution can be derived from the (subjective) probability estimates of a set of contingencies, whereas the second-order probability distribution is subjectively estimated directly. For fire protection problems, there appears, as mentioned above, to be little hope of being able to identify and estimate exactly in a practical way the probabilities of such contingencies. For example, what contingencies should be used in a case in which one wants to estimate the probability that the staff will extinguish a fire in a specific factory? One could perhaps take account of the number of people in the factory, the quality of the fire equipment, the time of the day that the fire occurs, and the like, but it nevertheless appears problematical to estimate precise probabilities for these factors that could influence the probability that the staff will extinguish the fire. Thus, the question arises of whether one should continue to search for contingencies that affect the contingencies that were identified as affecting the probability that the staff will extinguish the fire. This would appear to soon lead to a model that has lost all practical applicability. It is also doubtful that one would find such contingencies that one would feel satisfied in making an exact estimation of the probabilities of.

For *practical* decision making with respect to fire safety, it would seem better to use second-order probabilities to express one's uncertainty regarding specific probabilities in the model for fire spread one employs than to adopt a method involving contingencies such as those described above.